Calibrating the Performance and Security of Blockchains via Information Propagation Delays

Revisiting an old approach with a new perspective

Julius Fechner
Vrije Universiteit Amsterdam
juliusarf@gmail.com

Balakrishnan Chandrasekaran
Vrije Universiteit Amsterdam
b.chandrasekaran@vu.nl

Marc X. Makkes
Vrije Universiteit Amsterdam
m.x.makkes@vu.nl

ABSTRACT

Miners of a blockchain exchange information about blocks and transactions with one another via a peer-to-peer (P2P) network. The speed at which they learn of new blocks and transactions in the network determines the likelihood of forks in the chain, which in turn has implications for the efficiency as well as security of proof-of-work (PoW) blockchains. Despite the importance of information propagation delays in a blockchain's peer-to-peer network, little is known about them. The last known empirical study was conducted, for instance, by Decker and Wattenhofer in 2013 [11].

In this paper, we revisit the work of Decker and Wattenhofer on information propagation delays in Bitcoin. We update their measurement methodology to accommodate the changes made to the P2P network protocols since 2013. We also expand our measurement effort to include three other widely used blockchains, namely Bitcoin Cash, Litecoin, and Dogecoin. We reveal that block propagation delays have drastically reduced since 2013: The majority of peers in all four blockchains learn of a newly mined block within one second; the likelihood of forks is, consequently, low. Though blockchains networks have become quite efficient (i.e., have low delays), we observe that a significant number of nodes of these blockchains are present in cloud-provider networks and, more importantly, state-owned network providers; such deployments have crucial security implications for blockchains.

1 INTRODUCTION

In a permissionless, proof-of-work (PoW) blockchain, any miner can extend the chain by mining a new block on top of the most recently added block on the chain. Miners mine a block by compiling a set of transactions that are not included in a prior block and solving a cryptographic puzzle. Miners, hence, have to learn two pieces of information: (a) the most recent block added to the chain and (b) the transactions issued by users. They learn this information from one another by forming a peer-to-peer (P2P) network. Peers advertise and relay to others any new block or transaction that they recently learned. In this paper, we measure the speed at which this information propagates on the P2P network, since the propagation delays have crucial implications for the efficiency of blockchains.

SAC '22, April 25–29, 2022, Virtual Event

ACM ISBN 978-1-4503-8713-2/22/04.
https://doi.org/10.1145/3477314.3507003

A delay in learning about a newly added block to the chain translates to a likelihood that a miner will continue extending a "stale" chain: The miner may not append after the recently added block, since they are not yet aware of this block. Eventually the miner learns the missed update and all blocks on the stale chain are thrown away. Such disagreements between miners concerning what constitutes the chain at any point of time are referred to as *forks*. While forks cause the chain to branch, one of the branches eventually becomes the longest. Miners typically follow this longest chain (i.e., the one with most blocks), and all other branches are simply discarded. Although forks might occur due to mining being a stochastic process (e.g., more than one miner might mine a block at any given time), delays exacerbate the issue. Delays prolong the time for which the forks remain unresolved and, hence, likely increase the amount of "wasted" work [11].

Besides lowering the efficiency of blockchains, forks also have security implications [32]. Prior work on attacks on blockchains (e.g., [13]) typically assume that the miners have consensus. Forks imply, however, a lack of consensus among miners concerning the longest chain, and they reduce the fraction of miners with whom the attacker has to contend. Said differently, delays introduce forks that in turn decrease the threshold hash rate that an attacker needs for mounting a successful attack. Despite the implications of delays for the efficiency and security of blockchains, the topic has not received much attention from the community. The last empirical study, for instance, was conducted nearly a decade ago [11].

In this paper, we revisit the measurement study of Decker and Wattenhofer on information propagation delays in Bitcoin [11]. We updated their measurement methodology to account for the significant changes in the underlying (blockchain) P2P protocols since that study. We show that information propagation delays have substantially reduced since the decade-old study. We also expanded our study by including three more blockchains, namely Bitcoin Cash, Litecoin, and Dogecoin. These blockchains share a common code base with Bitcoin. Differences in observations across the blockchains should, therefore, primarily stem from differences in the infrastructure (e.g., where miners are located and how they peer with one another). We measure the information propagation delays across these blockchains and highlight that these networks are, typically, quite efficient (i.e., have low delays) today. We analyze potential factors that have enabled such efficient networks and find that some of the changes have crucial security implications. A non-trivial fraction of the peers (or *nodes*) are, for instance, deployed in networks owned or operated by nation states: They are, hence, highly vulnerable to regulatory changes (e.g., a ban on mining-related activities) enacted by nation states.

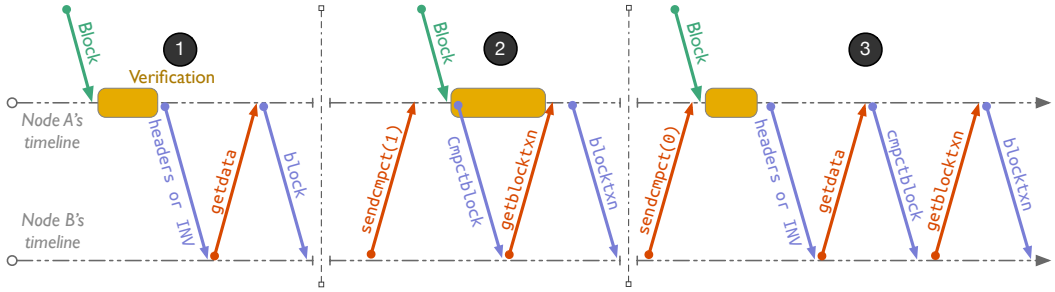

Figure 1: ***Illustration of the three different ways in which blocks may be propagated over a blockchain's P2P network.***

We summarize our contributions as follows.

⋆ We measure the information propagation delays for both blocks and transactions in the four blockchain-based applications (§3). Our measurements reveal that the majority of the peers in all four blockchains learn of a new block *within* a second.

⋆ We analyze each of the four P2P networks and share insights into the deployment of the peers (§4). We find that a non-trivial fraction of peers are (a) hosted on cloud service providers and (b) present in networks owned by state-owned network operators. Such deployment scenarios have crucial security implications for the blockchains and their users.

We have made the scripts and artifacts used in this study publicly available for supporting reproducible research [14]. Lastly, although we focus on cryptocurrencies, our measurement approach and inferences apply broadly to blockchains.

2 THE MECHANICS OF INFORMATION PROPAGATION

Below, we briefly discuss how the participants or nodes of a blockchain establish a P2P network and learn of blocks and transactions from one another. We use the term "node" to refer to the infrastructure run by a miner, or a *validator* (who only verifies the correctness of a block or transaction, but performs *no* mining), or a "casual" *observer* (who neither validates nor mines). The mechanics of information propagation in the P2P network are similar across all blockchains, except for some minor configuration details. We omit discussing these minor details to simplify the exposition.

The participants or nodes of a blockchain peer with one another to form a P2P network. When a node first joins the network, it does not know of any peers.[1] To find potential peers, the node either queries special DNS servers (or "seeds") that return a random set of other nodes or simply uses a set of predetermined or "hard coded" nodes in the network. The node may also connect to other reachable nodes, i.e., nodes that accept incoming connections, typically excluding those behind firewalls or network address translation (NAT) servers. After connecting with a peer, the node either proactively solicits the peer or passively listens to it for announcements of IP addresses of other peers. A node can peer, by default, with a maximum of 128 other nodes. Eight of these peerings are outgoing, initiated by the node itself, while the remainder are incoming. The manner in which peerings are established has no bearing on the interactions.

The peers communicate with one another via a P2P network protocol, which defines the message formats and interaction rules. Of these, we focus only on the subset of details pertaining to the exchange of blocks and transactions. The (Bitcoin) P2P protocol describes three messages for exchanging information on transactions and blocks: `INV`, `HEADERS`, and `CMPCTBLOCK`. When a node learns of a new block or transaction, it originally announced its hash in an `INV` message to its peers. Since 2017, the preferred method of announcing new blocks is, however, via `HEADERS` messages. The `HEADERS` are similar to `INV` messages, but contain only a list of block hashes, thereby separating the announcement of transactions and blocks. Lastly, `CMPCTBLOCK` carries a single block hash and a concise representation of the transactions contained in the block. The `INV` message is used primarily to relay transactions, while a small set of peers opt to use the `CMPCTBLOCK` messages.

Regardless of the message type used, the nodes follow a *gossip* protocol to disseminate information. When a node mines a block or creates a transaction, it will advertise this block or transaction to its peers. Upon reception a peer may request the full block or transaction (in case, it is the first time the peer learns of the item) and verify it. After verification, the process repeats: The peer re-advertises the block or transaction to other peers. Each peers works independently and maintains a *local* replica of the longest chain as well as pending transactions (i.e., those not yet included in any block), thus keeping the system decentralized.

Bitcoin, originally, defined only one method for block propagation (Part ❶ of Fig. 1) [28]. When a node receives a block that does not exist in its local replica of the blockchain, it would verify this block in its entirety before re-advertising it to its peers. Every node, working independently, will re-verify the block, prior to re-advertising it. This verification process, hence, adds delays to the propagation of the block.

The introduction of *compact* blocks aimed at reducing such delays. The key benefit of compact blocks over regular blocks is that in case of the former a peer uses the `CMPCTBLOCK` message (instead of the `BLOCK` message), which is much smaller in size than the `BLOCK` message. The receiving peer may have to issue, however, additional requests to retrieve details of transactions in the block. Compact blocks also introduce two new block propagation methods (Parts ❷ and ❸ of Fig. 1). When a node signals its intent to receive compact blocks to its peer, that peer will not wait until block verification is complete to re-advertise that block. This optimization reduces the

[1] In the rest of this paper, we use the terms 'peer' and 'node' interchangeably to refer collectively to both the user and their infrastructure.

Table 1: ***The sequence of blocks gathered and the number of reachable nodes observed in measurements obtained from four different blockchains.***

Blockchain	*Start block*	*End block*	*#reachable*
Bitcoin	685,000	691,500	9500
Bitcoin Cash	689,500	696,500	1000
Litecoin	2,060,000	2,085,000	1300
Dogecoin	3,745,000	3,800,000	1300

delays in block propagation at the risk of advertising invalid blocks. A node can, however, explicitly indicate to receive only verified (compact) blocks (❸ of Fig. 1) to mitigate such risks. Unsurprisingly, the downside of this approach is that it adds an extra round trip (for soliciting the compact block), which in turn increases the propagation delays.

Given these substantial changes in how peers exchange information about blocks or transactions, prior measurement efforts, notably Decker et al. [11], need to be extended to gather, parse, and analyze the new message types as well as to accomodate the protocol changes.

3 INFORMATION PROPAGATION DELAYS

We estimate the information propagation delays between nodes in a blockchain's P2P network (similar to [11]) as follows. Suppose that a new block b is first announced by a peer p_0 at some point in time t_0. Assume that t_0 is the earliest announcement of b in the network. When another peer p_i also advertises this block at time (t_i), we measure the time elapsed between t_0, when peer p_0 announced the block, and t_i, when peer p_i re-announced that block. This elapsed time then approximates the block propagation delay between the two concerned peers. By repeating this measurement for every subsequent peer that we observe re-announcing the block b, we estimate the block propagation delays over the P2P network.

To observe the block announcements and record timestamps of the announcements, we configured an observer node to join the P2P network, similar to prior work [11]. For each blockchain, we used the same software implementation for the observer that is typically used by other nodes in that blockchain. We use, for instance, *Bitcoin Core*, *Bitcoin Cash Node*, *Litecoin Core*, and *Dogecoin Core* implementations to study the propagation delays in Bitcoin, Bitcoin Cash, Litecoin, and Dogecoin, respectively.

We modified the observer node as follows. We changed the source code to ensure that the observer *cannot* perform any mining; it also does not validate any block or transaction. In each blockchain, the observer connects to as many peers as reachable on the P2P network; the only constraints are those defined by the underlying hardware and our network infrastructure. The observer logs each block or transaction that it receives along with the time it received that block or transaction and the IP address of the peer from whom it received that information. We processed these logs and persisted them in a simple relational database backend for further analyses. Unlike the prior work that collected only `INV` messages, we tailored our approach to gather, process, and analyze the `HEADERS` messages and `CMPCTBLOCK` messages. We also processed `INV` messages to estimate transaction propagation delays.

Table 2: ***Median and mean of block propagation delays (in seconds) observed in the four different blockchains.***

Blockchain	*Median delay*	*Mean delay*
Bitcoin	0.454	4.064
Bitcoin Cash	0.125	0.672
Litecoin	0.081	0.651
Dogecoin	0.181	1.693

3.1 Overview of Data Sets

We measured (via the observer nodes) the P2P networks of the four different blockchains for a period of 50 days. We observed a varying number of blocks across the four blockchains (Tab. 1), since the block generation times of these blockchains significantly differ from one another. Furthermore, since we customized the observer nodes to remove any connection limits, we observed a large number of peers over the course of the study period—much larger than that we would have observed had we limited ourselves to using the implementations with default configurations.

Of all the four networks, Bitcoin boasts the highest number of reachable peers, nearly 8-times than that observed in any other network. Despite there being a thousand or more reachable nodes, at any given point of time the observer nodes remained peered only with a subset of other nodes, presumably because of the high churn rate of nodes in the network. The observer nodes still simultaneously peered with a substantial fraction of the reachable nodes—varying from 26% in case of Bitcoin Cash to as high as 46% in case of Dogecoin—at many times during the course of this study.

Ethical concerns. Although we gather the IP addresses of peers in the P2P network, this information is publicly available to anyone joining the network. We also do not try to deanonymize the users behind the IP addresses. We do not, hence, raise any ethical issues.

3.2 Block propagation delays

Block propagation delays in Bitcoin have reduced substantially (Fig. 2) since the measurement effort of Decker et al [11]. We plot the normalized histograms of block propagation delays in Fig. 2, to approximate the probability distributions of the delays, in the four different blockchains. Per Fig. 2a, the median delay in Bitcoin is only 0.454 s, which is more than an order of magnitude lower than the 6.5 s delay observed in 2013 [11]. The low median delay in Bitcoin also matches the observations from another recent study [17]. The mean delay of approximately 4.1 s is, however, an order of magnitude higher than the median, indicating a long-tailed distribution: 5% of the nodes received new blocks, for instance, later than 15.8 s. The observed mean delay is still much lower than the mean of 12.6 s observed in the prior work.

The delays for propagating blocks between peers is also quite small in the other three blockchains (Fig. 2): The majority of nodes in all of these networks were able to learn of a new block within one second. The delays were lowest in Litecoin with a median of 0.08 s and 0.65 s in the mean. Of these three blockchains, Dogecoin experienced the longest delays with a median delay of 0.181 s, albeit that median delay is still more than twice smaller than that of Bitcoin. Unlike Bitcoin, the delay distributions for Litecoin and Bitcoin Cash do not exhibit significantly long tails: 97% and 95%

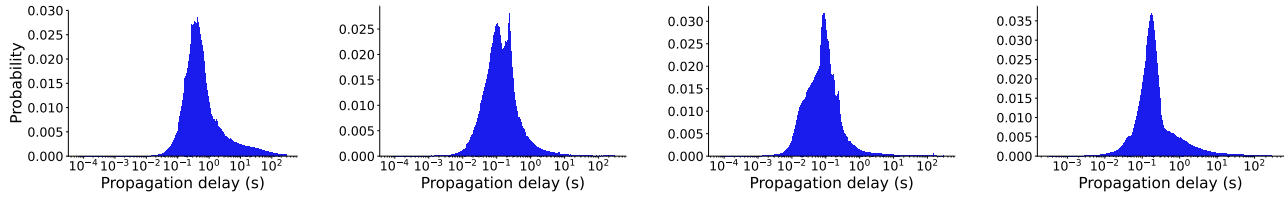

(a) *Bitcoin* (b) *Bitcoin Cash* (c) *Litecoin* (d) *Dogecoin*

Figure 2: ***Normalized histograms of block propagation delays in the P2P networks of different blockchains. The delays are generally quite low among all networks, and those in Bitcoin are substantially smaller than observed in prior work [11].***

of nodes in Litecoin and Bitcoin Cash, respectively, received new blocks within one second. Dogecoin is an outlier among these relatively new cryptocurrencies: In addition to having the largest median block propagation delay, it also exhibits a long-tailed delay distribution, with 5% of the nodes in its network taking more than 3.2 s to learn of a new block.

3.3 Musings on Delays

Blockchains have significantly evolved, since the work of the Decker et al. [11]. In this section, we discuss a few key changes or factors that have facilitated the low block propagation delays in the P2P networks of blockchains.

Fast *relay* networks are, undoubtedly, among the top enablers of low-latency (or delay) P2P networks. A (fast) relay network is an overlay network of nodes that is optimized for low-latency, by placing the nodes in strategic locations around the world and connecting them with one another via low-latency communication channels and/or protocols. A blockchain node can join such a relay to reduce the time it takes to learn of new blocks and transactions. The Fast Internet Bitcoin Relay Network (FIBRE) [7], for instance, is a relay for Bitcoin, and miners have been known to utilize this relay network for discovering new blocks as quickly as possible. Prior work has shown that such relay networks can substantially reduce the median propagation times [31].

The introduction of compact blocks in 2016, through the Bitcoin Improvement Proposal (BIP) 152 [8], was also aimed at reducing propagation delays. The proposal allowed nodes to re-advertise blocks *prior* to completing verification. Simulations using SimBlock [27] found that compact blocks may reduce propagation times by up to 90%, although more conservative estimates indicate a maximum reduction of 20% [26]. Most blockchains, including the four we studied, have adopted compact blocks.

Increases in network bandwidths also, unsurprisingly, reduces the block propagation delays. To estimate the increase in network bandwidths since 2013, we use the "State of the Internet" report from Akamai, one of the world's largest content delivery networks (CDNs). Per the Akamai report, the average Internet bandwidth in 2013 was around 8 to 9 Mbps in developed countries such as United States and Germany [4]. Such countries also host a large fraction of the nodes in different blockchains (discussed later in §4). The global average then was 3.6 Mbps. Today the average worldwide network bandwidth is approximately 100 Mbps, exhibiting a more than ten-fold increase compared to 2013, with an even larger average of 150 Mbps for some developed countries [2]. The average block sizes in Bitcoin have, however, increased nearly by a factor of three, from 500 KB in 2013 to about 1.5 MB in 2021. Size of blocks in other blockchains are, however, much smaller than that in Bitcoin, with Bitcoin Cash having the largest blocks at 600 KB. The large increases in network bandwidths coupled with relatively small block sizes may also explain the low propagation delays.

Table 3: ***Block generation times and forks observed in different blockchains compared with the analytically predicted values.***

Blockchain	#forks	%	P_b	Fork rate (%)
Bitcoin	1	0.015	705.759	0.576
Bitcoin Cash	1	0.014	629.365	0.106
Litecoin	8	0.032	156.151	0.417
Dogecoin	70	0.127	63.300	2.660

3.4 Blockchain forks

Forks occur when two blocks of the same height are mined and propagated across the network at the same time. For honest miners, forks should only occur if they continue mining a block B of height h while another valid block B' of same height is still being propagated across the network, and they are not yet aware of it. Once a miner becomes aware of a block that extends the blockchain, they will accept this block into their own local replica of the chain and begin mining atop this block. Therefore, a correlation exists between the propagation delays in the P2P network and the amount of forks experienced by the blockchain network. More precisely, the longer it takes for a block to disseminate through the network, the higher the chance that a miner mines another block of the same height.

In Tab. 3, we show the fork rates that we measured in the different blockchains via our observer nodes. Our nodes observed roughly 27%, 30%, 42%, and 46% of all reachable nodes in Bitcoin, Bitcoin Cash, Litecoin, and Dogecoin, respectively. While the lack of complete network coverage implies that the observed fork rates may not be completely accurate, they nevertheless offer a reasonable approximation of the current fork rates. Our estimates indicate that the number of forks is low (second column, Tab. 3), with fork rates between 0.014% and 0.127%. The low fork rates are also corroborated by prior work that examined the fork rates in Bitcoin [30], which also found a drastic drop in fork rates since 2017 and loosely correlate the observation to a reduction in the observed propagation delays. These low fork rates may certainly be due to the low propagation delays: With blocks being shared quickly across the network, there is little time for a miner to mine a conflicting block.

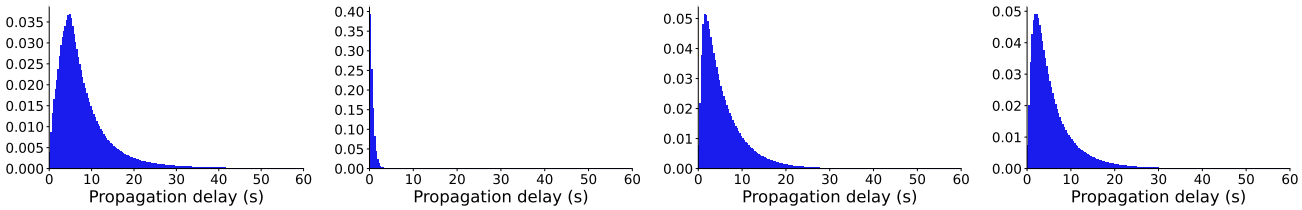

(a) ***Bitcoin*** **(b)** ***Bitcoin Cash*** **(c)** ***Litecoin*** **(d)** ***Dogecoin***

Figure 3: ***Normalized histogram of the transaction propagation delays for each blockchain. The delays observed here are considerably larger than the block propagation delays, taking several seconds for 50% of nodes to receive a transaction. (b) The delays in Bitcoin Cash are an anomaly, being considerably faster than those in other blockchain. Most nodes here receive a transaction within the first second.***

We checked our observed rates against the model from [11]:

$$Pr[F \geq 1] = 1 - (1 - P_b)^{\int_0^\infty (1 - f(t))\, df}$$

where P_b is the time to find a block, $f(t)$ is the expected value of informed nodes (i.e., nodes that have learned of the new block or transaction) at time t, and F is a discrete random variable that counts the number of conflicting blocks being found. We used the block generation times observed during our study for P_b and the mean of the propagation delays to estimate how fast nodes learn about a block. The model reported larger fork rates than those we observed (Tab. 3). While the discrepancies may simply be due to the lack of complete network coverage, another likely factor might be that hashing power is not equally distributed across nodes in the network, which is a key assumption in the above model; computational resources are often clustered in mining pools [1]. Even if we use the model's fork rates as worst-case estimates, we find that they are generally quite low, ranging from 0.106% to 0.576% for Bitcoin, Bitcoin Cash, and Litecoin. The model reports, however, larger rates for Dogecoin at 2.660%. The high fork rate in Dogecoin likely stems from two key factors: (a) its low block generation time and (b) its slightly larger distribution of propagation delays (compared to the rest). The combination of these two factors makes it highly likely for a new block to be mined while another is still being propagated by nodes across the network. The empirically observed fork rates as well as those predicted by the model, regardless of the minor discrepancies, seem to suggest that the networks are relatively resistant to blockchain forks.

We utilize the block generation time and the fork rates to estimate, furthermore, how the delays reduce the *effective* computational or hashing power, $\hat{c}$, of the network as follows:

$$\hat{c} = 1 - \int_0^\infty (1 - f(t))\, df \cdot P_b^{-1}$$

We then halve the effective hashing power to compute a simple security threshold (of majority hashing power). Any miner or mining pool possessing more hashing power than this security threshold could have a far reaching influence on both their earnings as well as the revenue of other miners or mining pools. The existence of one or more miners or mining pools with such substantial hashing power has, therefore, crucial security and fairness implications (refer [24]). The security thresholds per our analyses are quite high for Bitcoin, Bitcoin Cash, and Litecoin, suggesting that the information propagation delays do not have a profound impact. Even Dogecoin, which has the lowest effective computational power out of all the blockchains we studied, lower than that observed in Bitcoin in 2013 [11], still has a security threshold of 48.663% hashing power.

Table 4: ***Median and mean transaction propagation delays (in seconds) observed in the four different blockchains.***

Cryptocurrency	Median delay (s)	Mean delay (s)
Bitcoin	6.356	10.062
Bitcoin Cash	0.461	0.892
Litecoin	4.268	6.116
Dogecoin	4.339	10.700

3.5 Transaction propagation delays

The propagation delays for transactions are considerably longer than those for blocks (Fig. 3). In Bitcoin, transactions require 6.356 s to reach 50% of the nodes in the network, compared to the corresponding value of only 0.454 s for blocks. The tail of the distribution is quite long, with 95% of nodes taking 24.983 s to receive a transaction. The mean delay is, consequently, 10.062 s. While Decker et al. [11] did not study transaction delays, data from other sources show that transaction delays were less than 2 s in the median in January 2016, and have increased over time [17].

Litecoin and Dogecoin also have large delays, with median delays of 4.268 s and 4.339 s, respectively, as shown in Tab. 4. Much like Bitcoin, their transaction propagation delay distributions feature long tails: It takes 16.052 s for 95% of nodes in Litecoin to receive a transaction, and the corresponding value for Dogecoin is even higher, at 20.493 s. Bitcoin Cash, in contrast to the other three blockchains, experiences much lower delays of only 0.461 s in the median and 0.892 s in the mean.

The P2P protocol defines only one method for transaction propagation: the legacy propagation method using `INV` messages (refer §2). Every advertisement of a new transaction, therefore, requires multiple round trips, depending on the size of the transaction, between the nodes exchanging the data. Nodes must also fully verify a transaction before they can re-advertise that transaction to others. The transaction volume is also quite high. Dogecoin, which features the lowest amount of transactions of all the four studied blockchains, still handles around 20,000 new transactions per

Table 5: ***Number of unique peer IP addresses and the number of ASes to which they belong.***

Network	*Unique IP addrs.*	*Unique ASes*
Bitcoin	8247	1435
Bitcoin Cash	503	118
Litecoin	3044	118
Dogecoin	14,856	1568

day, followed by Bitcoin Cash with about 80,000 new transactions per day. Litecoin and Bitcoin, in contrast to the other two, handle more than 130,000 transactions per day. The inefficient propagation method coupled with the high volume of transactions may help explain the large transaction propagation delays.

Another potential factor behind the large delays is churn in the P2P network. Bitcoin Cash had the most steady network, with most nodes remaining connected throughout the observation period. In addition, almost all nodes in the Bitcoin Cash network advertised a transaction as it was announced. The other three blockchains, in contrast, observed a high degree of churn: Many nodes were transient. Furthermore, many nodes did not participate in relaying transactions, which may also result in transactions taking a long time to be disseminated to all nodes in the network. These facts may both help explain the discrepancies between block and transaction propagation delays as well as the low transaction propagation delays for Bitcoin Cash.

4 A PEEK INTO THE P2P NETWORKS

In this section, we review the key enablers of the highly efficient (i.e., low-latency) blockchain (P2P) networks.

We observed a substantial number of IP addresses for peers in each of the four different blockchains (Tab. 5). Bitcoin Cash was an outlier with at least an order of magnitude smaller number of peers than the other networks. We observed that the number of stable nodes, i.e., those that remained connected for prolonged periods of time, varied significantly across the four networks. Bitcoin Cash and Litecoin had 86% and 74% of stable nodes, respectively, while Bitcoin and Dogecoin only had 47% and 32% of such nodes. Since high churn adversely affects propagation delays [20], the large fraction of stable nodes (i.e., fraction experiencing no churn) for Bitcoin Cash and Litecoin may have a key part in their comparatively low propagation delays.

4.1 Underlying network infrastructure

We used Team Cymru's IP-to-ASN mapping [3] to identify the autonomous system (AS) associated with each IP address we gathered (from the peers) from the four different blockchain networks. Bitcoin and Dogecoin nodes exhibit a diverse geographical footprint, with nodes present in more than a thousand different ASes. These networks are, as a consequence, perhaps less susceptible to network-level attacks or outages. Bitcoin Cash and Litecoin, in contrast to the other two, have nodes deployed in relatively few ASes (Tab. 5). Litecoin's AS footprint is quite limited despite having a P2P network that is approximately six times bigger than that of Bitcoin Cash.

Table 6: ***Percentages of blockchain P2P nodes hosted on cloud-provider networks.***

ASN	Bitcoin	Bitcoin Cash	Litecoin	Dogecoin
Amazon	7.70	17.89	8.08	0.78
Digital Ocean	3.30	9.94	3.58	0.23
Google	3.78	1.99	1.28	0.24
OVH	2.24	2.58	3.55	0.44

Per Tab. 6, a substantial number of nodes are deployed in cloud-provider networks. The percentage of nodes on any one provider, however, is not significant, with the exception of Bitcoin Cash: We observed 17.89% of Bitcoin Cash nodes deployed on Amazon's cloud. There are also obvious overlaps in the use of cloud infrastructure across the four different blockchain networks. The cloud providers in total account for around 30% of all nodes in the different networks. Several factors may explain this reliance of blockchains on cloud providers and hosting services. One potential reason could be the size of the blockchains being too large to store on typical disks: Bitcoin requires, for instance, around 350 GB of storage. It might also be cheaper to deploy a node on cloud hosting services than to personally maintain a permenently online infrastructure.

We also find that a significant overlap between the geographic locations (i.e., centralization in physical deployment) of nodes in the different networks. To analyze the geographic footprint of the deployment of nodes in a network, we geolocated the IP addresses of the observed peers using Maxmind's Geolite2 City geolocation database [23]. Since IP geolocation databases are notoriously error-prone, we restrict our attention to only the country-level predictions of this database, which is well-known to have a high accuracy. At least 25% of the nodes across all networks are located in the United States, followed by Germany or China, as shown in Fig. 4. That a significant number of nodes use digital hosting services may explain the geographic deployment footprint (and the overlaps), since such digital hosting services have considerable infrastructure in the US and Germany.

4.2 On the lack of geographic diversity

The lack of geographic diversity may be partially responsible for the low block propagation delays. Even if physical proximity between endpoints does not always translate to low (network) latency between them, nodes that are physically close to one another may have a low latency connection. The physical proximity might two miners or mininig pools to establish a low-latency link between them, if one does not already exist, since the cost of such links should be a fraction of that already invested by the miners in their infrastructure.

One potential issue with the large amount of infrastructure sharing we observe across the different blockchains it that it may introduce single points of failure [16, 18, 19]. Should one of these digital hosting services, such as Amazon, experience outages, a considerable number of nodes across all the four blockchains may suffer connectivity issues; these nodes might not be able to continue mining successfully, if they become sequestered, for instance, from the majority of the nodes in the P2P network. Even if not all of these nodes are miners, the loss of connectivity experienced by

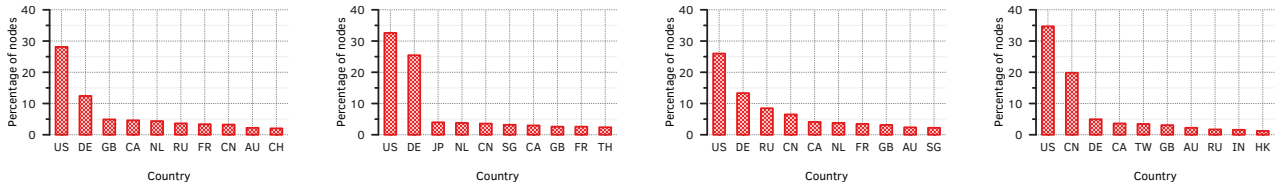

(a) ***Bitcoin*** **(b)** ***Bitcoin Cash*** **(c)** ***Litecoin*** **(d)** ***Dogecoin***

Figure 4: ***The top ten countries that host nodes for each blockchains. A large fraction of nodes in every network is located in the United States, followed by Germany, with the exception of Dogecoin, which has a large fraction of its nodes in China.***

Table 7: ***Percentages of blockchain P2P nodes located in networks of state-owned Internet operators.***

Bitcoin	*Bitcoin Cash*	*Litecoin*	*Dogecoin*
5.17	1.19	10.25	23.88

some nodes (because of failures) may have non-trivial implications, e.g., in terms of propagation delays, for the blockchains.

4.3 Susceptibility to nation-state attacks

A recent network-measurement effort identified the ASes of state-owned Internet operators in countries around the world [6]. We looked up the ASes of the IP addresses of peers we discovered in the different blockchains against this database of state-owned ASes. Our analyses revealed that a non-trivial fraction of the nodes of different blockchains are deployed in networks that are owned by state-owned network providers. Bitcoin, the largest cryptocurrency by market capitalization today, has about 5% of its nodes in state-owned networks. Relatively new blockchains such as Litecoin have twice as many nodes as Bitcoin deployed in state-owned networks. While Bitcoin Cash has comparatively few nodes in state-owned ASes, nearly a quarter of Dogecoin's network infrastructure is deployed in state-owned ASes. Such deployment footprints have crucial implications for the stability and security of blockchains. Sudden changes in laws, such as the Chinese government banning of cryptocurrency mining in June 2021, may significantly affect such nodes and lower a network's total hashing power.

5 LIMITATIONS

We used a one observer node for each of the four blockchains studied in this work, and all our observer nodes were physically colocated in Amsterdam, NL. The observer nodes, furthermore, only use IPv4 for peering, although reference implementations of nodes in most blockchains support IPv6. These two factors in combination may limit or bias our view of the P2P network of the different blockchains.

The observer buffers all data received from a peer, prior to processing. Only when the data (i.e., block or transaction announcement) is drained from the buffer, the processing thread can timestamp the entry. Therefore, when the observer is congested, i.e., as a result of receiving an overwhelming number of updates from a large number of peers, the "receive times" of announcements might be inaccurate. Such issues should not, however, introduce errors beyond a few milliseconds.

Notably absent from this study is Ethereum, which is the second largest cryptocurrency by market capitalization. The implementation differences between Ethereum and the other Bitcoin-based blockchains require some more thought into designing and performing precise experiments and analyses. We leave the inclusion of Ethereum as well as a longitudinal measurement study of blockchains to future work.

6 RELATED WORK

Decker et al. studied the information propagation delays in Bitcoin in 2013 [11]. Bitcoin and its network have evolved substantially over the years since that study. Several optimizations aimed at improving propagation delays and transaction throughout have since been adopted. This work, hence, focussed on reappraising the delays in Bitcoin, by modifying the original approach to take into account the recent protocol changes and optimizations. Recently, Neudecker et al. studied various characteristics of the Bitcoin network [29]. Compared to both these prior work, we evaluated the effect of the recent optimizations on the performance and security aspects of Bitcoin as well as three other large PoW blockchains.

While Decker et al. [11] outline how the propagation delays of Bitcoin affect its security, Kovalchuk et al. [21] derive a formula to calculate the security threshold for any blockchain given the block generation time and the propagation delays in the network. Following these prior work, we use the empirically measured propagation delays for calculating the security thresholds of blockchains.

Mariem et al. investigated the state of Bitcoin's centralization in 2020 by crawling the network [22]. In a similar vein, we utilized the IP addresses of peers that we discovered in the different networks to study the state of decentralization (or the lack thereof) of not only Bitcoin, but also of three other large PoW blockchains. Our experiments are easy to replicate and all the required artifacts are publicly available [14].

Many prior efforts focussed on discovering the network topologies of the P2P networks of blockchains. AddressProbe and CoinScope, for instance, reveal the Bitcoin topology and show that despite having low propagation delays, the inclusion of a transaction into a block may be affected by influential nodes in the Bitcoin network [25]. Our work is orthogonal to this study: Assuming that miners are honest with benign intent, we simply study how fast information propagates in different networks, and investigate some of the factors that have resulted in low delays.

PoW blockchains use substantial amounts of electricity, and there is, unsurprisingly, rich literature on analyzing the energy

use of blockchains and their implications. De Vries [10] found that Bitcoin consumed at least 2.55 GW of energy, with an upper bound of up to 7.67 GW, in 2018. A more recent study, in 2020, found that Bitcoin used around 4.3 GW of electricity [15]. They also showed that blockchains other than Bitcoin added another 50% to the total electricity consumption, with Litecoin using 0.164 GW, Dogecoin using 0.157 GW, and Bitcoin Cash using 0.153 GW. Reducing the propagation delays will ensure that the invaluable energy already spent in mining a block is not wasted by the network. To this end, our study shows that the P2P networks of four large PoW blockchains are quite efficient, with small propagation delays.

7 CONCLUSION

In this paper, we analyzed the information propagation delays in four of the largest proof-of-work Blockchains: Bitcoin, Bitcoin Cash, Litecoin, and Dogecoin. We showed that the propagation delays in Bitcoin have reduced significantly—from 6.5 s to 0.45 s in the median—since they were first studied in 2013. Our measurements indicate that the propagation delays are rather small in all four blockchains.

If we reflect upon these delays and assume that all miners experience similar propagation delays, we can deduce the average cost of energy waste that they effect in the different blockchains. Using the current estimates of energy consumption of Bitcoin of around 132.5 TWh/year [9, 12], we estimate a median wastage of around 0.449 TWh/year (0.00339%) resulting purely from propagation delays. This estimate provides an upper bound of the energy waste stemming only from propagation delays. In calculating this estimate, we assume that any individual miner or node may mine a block, which is generally considered a rare case since the majority of miners participate in mining pools [5]. We estimate a lower bound of 0.376 TWh/year, if we assume the most powerful mining pool mines a block, which (as of this writing) would be a mining pool with 16.2% of mining power in the network. While these estimates indicate a substantial energy waste, the loss might be still limited given the concentration of nodes in a few geographic locations. We hope that this work provides a reference point for blockchain designers and researchers to understand (a) the current state of the networks and (b) the implications of delays on energy waste, and nudges them towards further optimizing these networks.

ACKNOWLEDGEMENTS

This work was partially funded by the Dutch Organisation for Scientific Research (NWO) under contract 629.009.014. Furthermore, we thank Marcel Mulder for contributing computational resources to the project.

REFERENCES

[1] 2021. Bitcoin Mining Map. https://cbeci.org/mining_map
[2] 2021. Speedtest by Ookla - The Global Broadband Speed Test. https://www.speedtest.net/ (Accessed: 5-9-2021).
[3] 2021. Team Cymru: IP to ASN Mapping Service. (2021). https://team-cymru.com/community-services/ip-asn-mapping/ (Accessed: 10-14-2021).
[4] Akamai. 2013. The State of the Internet, 1st Quarter 2013 Report. (2013). https://www.akamai.com/us/en/resources/our-thinking/state-of-the-internet-report/global-state-of-the-internet-connectivity-reports.jsp
[5] BTC.com. 2021. Bitcoin Pool Statistics. (2021). https://btc.com/stats/pool (Accessed: 5-9-2021).
[6] Esteban Carisimo, Alexander Gamero-Garrido, Alex. C Snoeren, and Alberto Dainotti. 2021. Identifying ASes of State-Owned Internet Operators. In *Proceedings of the ACM Internet Measurement Conference (IMC '21)*.
[7] M. Corallo. 2019. FIBRE: Fast internet Bitcoin relay engine. (2019). https://github.com/bitcoinfibre/bitcoinfibre
[8] M. Corallo. 2020. Compact block relay. BIP 152. (March 2020). https://github.com/bitcoin/bips/blob/master/bip-0152.mediawiki
[9] Raynor de Best. 2021. Bitcoin energy consumption worldwide from February 2017 to June 27, 2021. (October 2021). https://www.statista.com/statistics/881472/worldwide-bitcoin-energy-consumption/
[10] Alex de Vries. 2018. Bitcoin's Growing Energy Problem. *Joule* 2, 5 (2018).
[11] C. Decker and R. Wattenhofer. 2013. Information propagation in the Bitcoin network. In *IEEE P2P 2013 Proceedings*.
[12] Digiconomist. 2021. Bitcoin Energy Consumption Index. (2021). https://digiconomist.net/bitcoin-energy-consumption/ (Accessed: 5-9-2021).
[13] Ittay Eyal and Emin Gün Sirer. 2014. Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In *Financial Cryptography and Data Security*.
[14] Julius Fechner, Balakrishnan Chandrasekaran, and Marc X. Makkes. 2021. Artifacts of the study titled "Calibrating the Performance and Security of Blockchains via Information Propagation Delays", published at ACM SAC '22. (2021). https://github.com/JuliusAF/info-prop-delays-artifacts
[15] Ulrich Gallersdörfer, Lena Klaaßen, and Christian Stoll. 2020. Energy Consumption of Cryptocurrencies Beyond Bitcoin. *Joule* 4, 9 (2020).
[16] Peter Garraghan, Renyu Yang, Zhenyu Wen, Alexander Romanovsky, Jie Xu, Rajkumar Buyya, and Rajiv Ranjan. 2018. Emergent failures: Rethinking cloud reliability at scale. *IEEE Cloud Computing* 5, 5 (2018).
[17] DSN Research Group. 2021. Bitcoin Network Monitor. https://www.dsn.kastel.kit.edu/bitcoin/
[18] Haryadi S Gunawi, Mingzhe Hao, Riza O Suminto, Agung Laksono, Anang D Satria, Jeffry Adityatama, and Kurnia J Eliazar. 2016. Why does the cloud stop computing? lessons from hundreds of service outages. In *Proceedings of the Seventh ACM Symposium on Cloud Computing*.
[19] Sebastian Hagen, Michael Seibold, and Alfons Kemper. 2012. Efficient verification of IT change operations or: How we could have prevented Amazon's cloud outage. In *2012 IEEE Network Operations and Management Symposium*.
[20] Muhammad Anas Imtiaz, David Starobinski, Ari Trachtenberg, and Nabeel Younis. 2021. Churn in the Bitcoin Network. *IEEE Transactions on Network and Service Management* (2021). https://doi.org/10.1109/TNSM.2021.3050428
[21] Lyudmila Kovalchuk, Dmytro Kaidalov, Andrii Nastenko, Mariia Rodinko, Oleksiy Shevtsov, and Roman Oliynykov. 2019. Decreasing Security Threshold Against Double Spend Attack in Networks with Slow Synchronization. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*.
[22] Sami Ben Mariem, Pedro Casas, Matteo Romiti, Benoit Donnet, Rainer Stütz, and Bernhard Haslhofer. 2020. All that Glitters is not Bitcoin – Unveiling the Centralized Nature of the BTC (IP) Network. In *NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium*.
[23] MaxMind. 2021. GeoLite2 Free Geolocation Data. (2021). https://dev.maxmind.com/geoip/geolite2-free-geolocation-data?lang=en (Accessed: 10-14-2021).
[24] Johnnatan Messias, Mohamed Alzayat, Balakrishnan Chandrasekaran, Krishna P. Gummadi, Patrick Loiseau, and Alan Mislove. 2021. Selfish & Opaque Transaction Ordering in the Bitcoin Blockchain: The Case for Chain Neutrality. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC)*.
[25] Andrew Miller, James Litton, Andrew Pachulski, Neal Gupta, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2015. Discovering Bitcoin's Public Topology and Influential Nodes. https://www.cs.umd.edu/projects/coinscope/coinscope.pdf. (2015).
[26] Jelena Mišić, Vojislav B. Mišić, and Xiaolin Chang. 2020. On the Benefits of Compact Blocks in Bitcoin. In *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*.
[27] Ryunosuke Nagayama, Ryohei Banno, and Kazuyuki Shudo. 2020. Identifying Impacts of Protocol and Internet Development on the Bitcoin Network. In *2020 IEEE Symposium on Computers and Communications (ISCC)*.
[28] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. https://bitcoin.org/bitcoin.pdf. (October 2008).
[29] Till Neudecker. 2019. *Characterization of the Bitcoin Peer-to-Peer Network (2015-2018)*. Technical Report 1. Karlsruher Institut für Technologie (KIT).
[30] Till Neudecker and Hannes Hartenstein. 2019. Short Paper: An Empirical Analysis of Blockchain Forks in Bitcoin. In *Financial Cryptography and Data Security*.
[31] Kai Otsuki, Yusuke Aoki, Ryohei Banno, and Kazuyuki Shudo. 2019. Effects of a Simple Relay Network on the Bitcoin Network. In *Proceedings of the Asian Internet Engineering Conference (AINTEC '19)*.
[32] Shengling Wang, Chenyu Wang, and Qin Hu. 2019. Corking by Forking: Vulnerability Analysis of Blockchain. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*.