Quantitatively Analyzing Relay Networks in Bitcoin

Kai Otsuki, Ryohei Banno†, and Kazuyuki Shudo
Tokyo Institute of Technology
Tokyo, Japan
ootuki.k.aa@m.titech.ac.jp

***Abstract*—Bitcoin's throughput is far lower than those of centralized systems such as Visa and PayPal. To handle a comparable number of transactions, Bitcoin's throughput must be improved, which will require a reduction in the block propagation time. Relay networks have attracted attention as a method to improve throughput. However, the nature and strength of the effects of relay networks are not yet clear. In particular, the effects on individual nodes have received little attention. Thus, this study aims to investigate the effects of relay networks on the Bitcoin network and its nodes under the current network state in approximately 2019. We simulate a relay network under the current network state because simulation is virtually the only way to obtain data for all nodes. A major contribution of this study is that it is the first to investigate the relay network effect on the 90th percentile of block propagation time. Moreover, this study shows that relay networks benefit not only the system but also the nodes utilizing them. The utilizing nodes' blocks were more than six times more likely to be approved as valid blocks when nonutilizing and utilizing nodes generated blocks simultaneously. This finding can motivate the use of nodes in relay networks.**

***Index Terms*—Bitcoin, blockchain, relay network**

I. INTRODUCTION

Blockchain has generated extensive study and economic interest. The total value of Bitcoin, one of the best-known cryptocurrencies, exceeded 32 billion USD in 2017. From a technical perspective, blockchain technologies have solved some challenges; for example, they can manage information securely and make it almost impossible to tamper with, even when multiple malicious nodes participate in the network. In addition, blockchain applications do not need a single management entity. On the other hand, room for improvement remains. One problem with Bitcoin is low transaction throughput. Transaction throughput is defined by the number of transactions that a system can handle per second, and it is a performance indicator for system capacity. Compared to centralized systems, the throughput of Bitcoin and many other blockchain technologies is small. Bitcoin's throughput is approximately 7 transactions per second (TPS), whereas centralized systems such as Visa and PayPal have throughputs of 1776 TPS and 700 TPS, on average [1], [2].

Previous work [3], [4] has demonstrated that reducing the block propagation time can improve throughput while retaining the same level of security. One method for providing faster block propagation is utilizing a relay network. A relay network is an external network that offers faster block propagation, and the nodes utilizing it can receive and send blocks via the network's relay servers.

Although there are some measurements of relay networks, measurements of their effects are still needed. Relay network operators of Bitcoin [5], [6] make their data available to the public, but they cannot know the entire set of network data because their data are real data. Thus, it is impossible to know the relay network effects on the whole network and on individual nodes from their data. [7] examined the effects of a relay network when all nodes utilized it, but they did not consider the effects on individual nodes, such as the mining success rate, and they did not reveal the relay network effects when a portion of the nodes utilized it. Moreover, they simulated the Bitcoin network as of 2015 because their study was conducted in approximately that year. In this study, we designed a realistic relay network model and performed simulation on the latest Bitcoin network while changing the number of nodes utilizing the relay network to examine the effects of relay networks on the whole network and the individual nodes.

Our prior study [8] investigated the relay network effects by simulation. However, there are two significant differences between this work and the previous work; simulation accuracy and detailed analysis. In the previous work, the relay network model was quite simple and primitive, and it simulated the Bitcoin network as of 2015. Simulations of the latest Bitcoin network are needed because network parameters such as bandwidth and latency have improved in recent years. In addition, the simulations in the paper [8] were performed with uniform hash power of all nodes to investigate the mining stats fairly. However, the uniform hash power might have influenced the data, such as the block propagation times and the orphan block rate. In the present simulation, the hash power of each node is set based on the real hash power distribution used in the study [9]. The simulation parameters including the hash power are set to be realistic and up-to-date. The parameter settings are discussed in detail in Section 3. The relay network models in the current and previous study are shown in Figure 1. In the previous model, the utilizing nodes were able to send blocks to all the other utilizing nodes directly, and thus the relay network was somewhat stronger than the actual relay networks. This study uses a realistic relay network model that reflects the features of actual relay

† Current affiliation as of October 2020: Kogakuin University

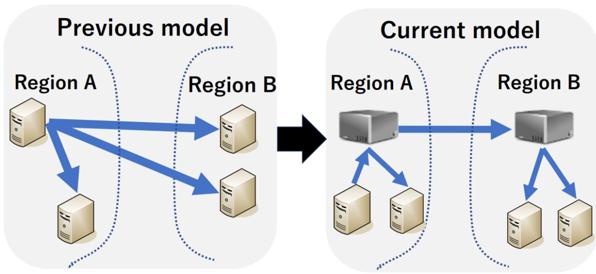

Fig. 1: The comparison of the previous model and the current one.

networks. Our current model is explained in detail in Section 3. In the detailed analysis, we quantitatively investigate the effects of the relay network and present not only the results but also their causes. For example, we thoroughly discuss the reason why the orphan block rates of the utilizing nodes are much smaller than those of the non-utilizing nodes in Section 4. Moreover, many new findings are obtained. For instance, this paper is the first to quantitatively investigate the effects on the 90% block propagation time. The relay network is found to have a larger influence on the 50% block propagation time than on the 90% block propagation time. Moreover, this paper compares the orphan block rates with those in a previous study by Gervais et al. [7] and demonstrates that the effect of relay networks on the orphan block rate is smaller in the present network. With respect to the effect on the individual nodes, although our previous study [8] suggested that the orphan block rate of utilizing nodes is smaller than that of non-utilizing nodes, this study ensures that the effect is maintained in the current realistic model under the latest network. Moreover, a more detailed analysis, which is shown in table VI, reveals that the utilizing nodes' blocks are more than six times more likely to be approved as valid blocks when non-utilizing and utilizing nodes generate blocks simultaneously. This is one of our new findings.

This paper reveal the effects of the relay network on the network and the individual nodes and provides a detailed numerical analysis of the effects.

II. Bitcoin and Relay Networks

This paper aims to identify the effects of relay networks, which are a recognized method for improving block propagation times. Herein, we describe block propagation in the Bitcoin network and the challenges it faces and then explain relay networks. These descriptions will be useful for understanding this study.

A. Transaction handling

Nodes will issue a transaction when they send coins. The coin transmission is completed when a block including the transaction is confirmed. When nodes issue transactions, they propagate them to the network through adjacent nodes. When nodes receive transactions, they save the transaction and then try to mine blocks including those transactions. For security reasons, block generation requires a great deal of computation. If nodes succeed in block generation and the blocks are confirmed, then they will receive mining rewards, which is the main incentive for nodes to participate in the Bitcoin network [10]. When nodes are able to mine blocks, they propagate these blocks to the network. When nodes obtain blocks, they confirm them, and the blocks will be appended to their ledger if they are valid. At this stage, the transactions included in the block are confirmed. The number of confirmed transactions per second is called the transaction throughput, and this number is used to evaluate system capacity.

B. Orphan blocks

The blocks that are not included in the longest chain of the ledger are called orphan blocks. Orphan blocks are important for both the system and its nodes. Prior works [3], [7] have reported that orphan blocks have a negative influence on security. Therefore, decreasing orphan blocks is important from a security perspective. For miners, orphan blocks are bad because they cannot receive mining rewards for orphan blocks; orphan blocks only waste their computation power. For both the whole system and its nodes, decreasing orphan blocks is significant. Previous works [3], [7] show that improving the block propagation time is necessary to decrease the number of orphan blocks.

C. Transaction throughput

Bitcoin's major challenge is to improve its transaction throughput. Indeed, its throughput falls behind that of centralized systems. Basically, there are two ways to improve throughput [4]: increase the block size or to reduce the block interval. However, improvement of the block propagation time is needed in both cases [3] because both cases are accompanied by an increase in orphan blocks. In other words, improving the block propagation time makes it possible to implement these methods while maintaining security.

D. Relay networks

Several types of relay networks [5], [6], [11] have been used in the Bitcoin network. Relay networks have attracted attention as a method to improve the block propagation time, and their effectiveness has been discussed [7], [8], [12]. Four projects listed in Table I are well-known in Bitcoin. These relay networks differ from each other, but the basic structure is common to some extent across them. Figure 2 depicts this structure. Relay networks are external networks and are composed of relay servers as shown in Figure 2. These servers are deployed in each region. For example, Falcon [5] relay servers are placed in the US, Germany, Brazil, Japan, Australia and several other regions. The following provides an overview of block distribution by relay networks. First,

TABLE I: Summary of relay network projects.

Project	Year	# servers	Developer
BFRN [6]	2014	8	Matt Corallo et al.
Falcon [5]	2016	10	Cornell University research team
FIBRE [11]	2016	6	Matt Corallo et al.
bloXroute [13]	2018	Unknown	Uri Klarman et al.

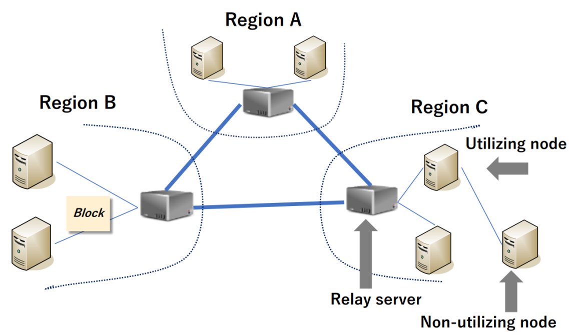

Fig. 2: Relay network structure.

utilizing nodes send blocks to relay servers in the same region when they discover the blocks. Then, relay servers propagate the blocks to other servers; block transmissions between servers are basically faster than normal transmissions. For instance, FIBRE [11] uses an UDP-based protocol, topology optimization and other methods to realize rapid transmission. Finally, the utilizing nodes can receive blocks from the relay servers with which they connect. Therefore, relay networks decrease propagation times between continents and can provide blocks with a large number of nodes simultaneously.

III. METHODOLOGY

This paper investigates relay network effects on the blockchain network and individual nodes when an ideal relay network is utilized in the network. Simulation is an appropriate and realistic approach for this examination because it is virtually impossible to obtain real data of individual nodes in peer-to-peer networks. This section describes our simulation settings and relay network model. At the end of this section, we explain how we analyze the simulation results.

A. Simulation

In this paper, SimBlock [14], which is a relatively new blockchain network simulator, was utilized. SimBlock has been used in several studies [8], [9], [15], and the simulator parameters, such as bandwidth and latency, have been updated [9]. Therefore, this simulator provides a simulation under the latest network state. The simulation parameter settings are shown in table II. The number of Bitcoin nodes was 10,000 based on data obtained from [16]. The block interval and block size followed those of Gervais et al. [7]. Bandwidth and latency were obtained from Nagayama et al. [9]; these values were the most recent internet values.

TABLE II: Parameter settings.

# of nodes	10000
Block interval	10 min
Block size	535 KB
Hash Power	Distribution according to Aoki et al. [14]
# of connections	Distribution according to Miller et al. [17]
Geographical distribution	Distribution according to Gervais et al. [7]
Bandwidth	6 regional bandwidth according to Nagayama et al. [9]
Delay	6 regional delay according to Nagayama et al. [9]

In this simulator, the transmission time is determined by the message size, latency and bandwidth. Bandwidth is chosen as the smaller of the upstream bandwidth of the sender or the downstream bandwidth of the receiver.

Simulations were conducted while changing the utilization rate of the relay network, that is, the rate of nodes utilizing the relay network. In detail, 100%, 50%, 25%, 12%, 6%, 3%, 1% and 0% were chosen as the utilization rates. Each simulation terminated when 1000 non-orphan blocks were generated.

B. Relay network model

For simulations, the relay network must be modeled, and this paper aimed to capture the effects of relay networks when they worked ideally. Basically, each relay network is quite different; specifically, their protocol, the number of deployed servers and their servers' capacities differ. Therefore, it is unrealistic to create a relay network model that reflects their traits identically. However, their ideal operation is basically common and obvious. Ideally, relay networks are expected to propagate blocks to nodes in different regions in three steps. First, nodes send blocks to one of the relay servers with which the nodes connect. The second step is transmission between regions; these connections are between relay servers and are faster than normal for the reasons mentioned in the previous section. Finally, the relay servers provide blocks with their connecting nodes in parallel. These three steps are basically the minimal steps for intercontinental transmission.

We aimed to simulate such a relay network when working ideally. Figure 3 shows the model. This model assumes that relay servers are deployed in each region and utilizing nodes receive/send blocks from/to the relay server in the same region. The model also assumes that the connections between these relay servers provide 10 times broader upward bandwidth than normal connections. In general, the connections between relay servers are faster than normal connections for the aforementioned reasons, and the model provides fast block propagation in the form of broader upload bandwidth. Next, when the relay servers receive blocks, they can send those blocks to all utilizing nodes in the same region simultaneously because the model simulates relay networks when they are working ideally. This simultaneous transmission is made possible by a server load distribution with a large bandwidth.

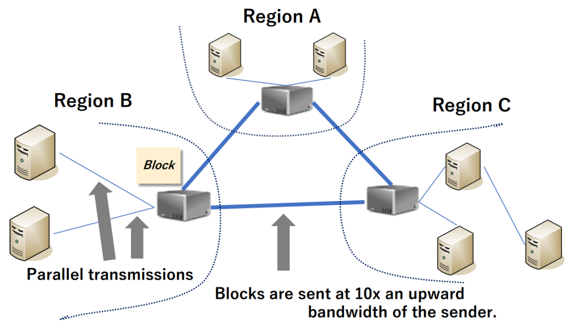

Fig. 3: Our relay network model.

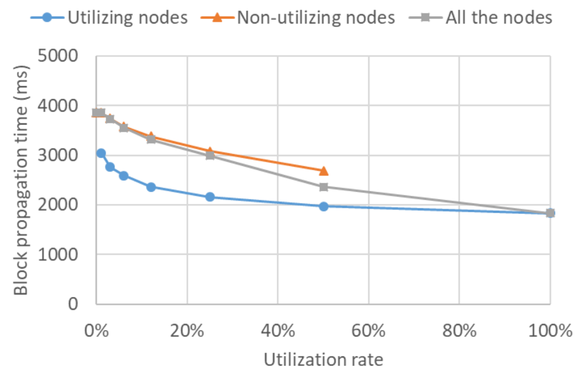

Fig. 4: 50th percentile of block propagation times.

C. Evaluation indices

Relay network effects on the Bitcoin network and individual nodes are examined in this paper. To evaluate the effects on the blockchain network, 50% and 90% of the block propagation time and orphan block rate were chosen, respectively. These variables are commonly used in many studies [3], [7], [8], [12], [15], [18], and they are appropriate parameters for evaluating effects. In particular, although the 90th percentile of block propagation time is important and has been investigated in many studies [12], [18], previous relay network measurements [7], [8] did not focus on it.

Moreover, this paper focuses on the number of generated blocks and the orphan block rate and compares them for utilizing nodes and -utilizing nodes to clarify the effects on individual nodes. The number of generated blocks is defined as the number of generated blocks per node, which is an important value for miners because they participate in the blockchain network to obtain mining rewards [10]. The other parameter, the orphan block rate of each node group, is the rate of each node's orphan blocks in all generated blocks; it is a significant value because nodes cannot obtain mining rewards for orphan blocks, and computational power is wasted in many blockchains, such as Bitcoin [10]. In other words, if relay networks decrease their orphan blocks, this would be a reason for nodes to utilize them.

IV. RESULT

This section examines the effects of relay networks based on the experimental results. As mentioned in the previous section, the block propagation time, orphan block rate and mining statistics were investigated. This section focuses on each individually.

A. Propagation time

Data obtained in previous studies [8], [12] indicate that relay networks decrease the block propagation time. Our results suggested two effects.

First, they showed an improvement in block propagation time, consistent with previous studies. Figure 4 and Figure 5 depict the 50th and 90th percentiles of block propagation time, respectively. In both figures, decreases in propagation time were observed in all three node groups. Regarding non-utilizing nodes and all nodes, block propagation times steadily decreased as the utilization rate rose. In fact, Table III and IV show that the 50th and 90th percentiles of propagation time decreased by approximately 100 ms and 200 ms, respectively, for both groups when comparing the 0% and 3% cases. Similar to the other node groups, an improvement in the utilizing nodes was observed. Therefore, relay networks improve the block propagation times of non-utilizing nodes and all nodes in addition to improving those of the utilizing nodes.

Second, an interesting observation is that the graph shapes of the non-utilizing and utilizing nodes resemble each other, which seems to indicate that the degree of improvement in these two groups was approximately the same. In fact, the differences in value for non-utilizing and utilizing nodes were kept within approximately 800 ms to 1000 ms in all cases according to Table III. The differences in the 90th percentile were within approximately 1500 ms to 2000 ms at any utilization rate, as shown in Table IV. However, note that the times of the utilizing nodes were always better than those of the non-utilizing nodes. Thus, the utilizing nodes maintain their advantage, although the degree of improvement is similar to that of non-utilizing nodes.

Moreover, the comparison of the 50th and 90th percentiles of propagation time reveals an interesting phenomenon. The differences between the 50th percentiles of non-utilizing and utilizing nodes were within approximately 800 ms to 1000 ms in all cases according to Table III. On the other hand, the differences in the 90th percentile were within approximately 1500 ms to 2000 ms, as shown in Table IV. This seems to suggest that relay networks have a smaller influence on the 90th percentile than the 50th percentile. Relay networks quickly disseminate blocks to nodes that are in the vicinity of their relay servers, and nodes that are distant from them have to rely on the Bitcoin network to obtain blocks, which seems to contribute to the differences between the 90th percentile and the 50th percentile.

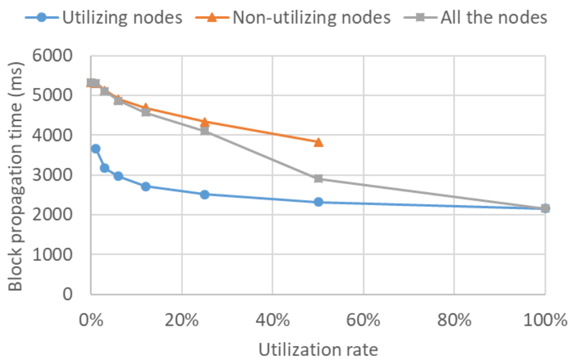

Fig. 5: 90th percentile of block propagation times.

TABLE III: 50th percentile of block propagation times (msec).

Group \ UR[1]	0%	1%	3%	6%	12%	25%	50%	100%
Utilizing	N/A	3050	2763	2590	2365	2160	1969	1837
Non-utilizing	3854	3853	3741	3576	3378	3081	2692	N/A
All	3854	3851	3730	3555	3316	2993	2366	1837

[1] utilization rate

B. Orphan block rate

Previous studies [7], [8] have indicated that relay networks decrease the orphan block rate. According to Gervais et al. [7], the orphan block rate was approximately 7.5% of the original value when all nodes participated in the relay network. Table V shows our results. We found that the improvement increased as the utilization rate rose, and the orphan block rate was approximately 36% of the original value when the utilization rate was 100%. In other words, our result indicates that the relay network improved the orphan block rate, but the improvement was smaller than the improvement observed in the study of Gervais et al. [7]. Two factors seem to contribute to this difference.

First, the number of Bitcoin nodes in our simulation was larger than that in the previous study. In general, the more nodes that participate in the network, the more time is required to propagate data. As a result, orphan blocks are more likely to emerge.

Second, the original orphan block rate of our study was much better than that used in the previous work. These experiments simulated the latest Bitcoin network, and the network is much faster than the network of the first study. As a result, the effect of relay networks decreased.

C. Number of generated blocks

The average numbers of generated blocks for the utilizing and non-utilizing nodes were compared to investigate the effect on individual nodes.

The results are shown in Figure 6. In our experiments, no significant difference was observed between utilizing nodes and non-utilizing nodes. In fact, the values of the utilizing nodes were larger in some cases and smaller in other cases than those of non-utilizing nodes regardless of the utilization rate. Additionally, the differences in the values were smaller than 0.01 in all cases.

The main reason for this result seems to be that the propagation time differences between non-utilizing and utilizing nodes were not sufficient. In general, hash power and time to start mining determine a great portion of mining success. In these experiments, we assumed that all nodes had the same hash power. Thus, it can be said that the difference in the block propagation time produced by the relay network was not sufficient. In fact, the differences between non-utilizing nodes and utilizing nodes were at most 1000 ms to 2000 ms at any utilization rate, as mentioned in the previous subsection, and it is obvious that the differences were too small in Bitcoin's block interval of 10 minutes. As a result, the differences were not observed at all utilization rates.

TABLE IV: 90th percentile of block propagation times (msec).

Group \ UR[1]	0%	1%	3%	6%	12%	25%	50%	100%
Utilizing	N/A	3660	3171	2972	2719	2511	2318	2154
Non-utilizing	5317	5307	5126	4905	4691	4338	3831	N/A
All	5317	5302	5106	4862	4573	4100	2903	2154

[1] utilization rate

TABLE V: Orphan block rate.

Utilization rate	0%	1%	3%	6%	12%	25%	50%	100%
Orphan block rate	0.95	0.93	0.89	0.79	0.76	0.63	0.43	0.34

D. Orphan block rate of each node group

In the previous subsection, the focus was on the orphan block rate to investigate the effect on the Bitcoin network. Herein, we compare the rates of the utilizing and non-utilizing nodes. The comparison is shown in figure 7. This result is consistent with the previous work [8] and indicates that utilizing nodes are significantly less likely to generate orphan blocks than non-utilizing nodes. In fact, the values of the utilizing nodes were always below those of the non-utilizing nodes, and the utilizing nodes generated no orphan blocks when the utilization rate was 1%. The effect is maintained in the realistic model.

Herein, two observations can be made. First, although the values of the utilizing nodes were always better than those of the non-utilizing nodes, the non-utilizing nodes' values consistently decreased as the utilization rate increased. This improvement is attributable to the propagation time improvement, as shown in Figure 4 and 5.

Second, the orphan block rate of utilizing nodes was not always proportional to the utilization rate. One reason is as follows. The utilizing nodes produced orphan blocks mainly when multiple utilizing nodes produced blocks at the same

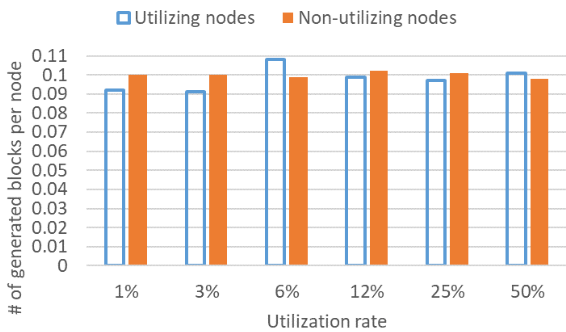

Fig. 6: Average number of discovered blocks.

TABLE VI: A breakdown of orphan blocks (%).

Case / UR[a]	case 1[1]	case 2[2]		case 3[3]	case 4[4]
1%	0	3	≫	0	97
3%	3	3	≫	0	94
6%	2	6	≫	1	91
12%	10	14	≫	0	76
25%	13	30	≫	4	53
50%	24	24	≫	4	48

[a] utilization rate.
[1] the possibility that both an orphan block and the main block are generated by utilizing nodes.
[2] the possibility that an orphan block and the main block are generated by non-utilizing and utilizing nodes, respectively.
[3] the possibility that an orphan block and the main block are generated by utilizing and non-utilizing nodes, respectively.
[4] the possibility that both an orphan block and the main block are generated by non-utilizing nodes.

time. Table VI, which shows the breakdown of how orphan blocks were produced, demonstrates that the possibilities of case 1 were larger than those of case 3. In other words, the majority of orphan blocks created by utilizing nodes were generated when an orphan block from a utilizing node was beaten by a block produced by another utilizing node. When the utilization rate is low, the orphan block rate of utilizing nodes is likely to be low because the number of utilizing nodes is small and plural blocks are rarely produced by the small group of nodes. By contrast, it is thought that when the utilization rate is high, the orphan block rate of utilizing nodes is likely to be higher because the numbers of utilizing nodes are larger and blocks are more likely to be produced at the same time compared to cases with low utilization rates. Therefore, it seemed that the values of utilizing nodes were not always proportional, and the values under low utilization rates were smaller those under high utilization rates.

Moreover, an important new discovery can be disclosed from Table VI. The possibilities of case 2 are more than 6 times larger than those of case 3 in all utilization rates, which indicates that the blocks provided by utilizing nodes were quite likely to become main blocks when the non-utilizing and utilizing nodes produced blocks at the same time. Surprisingly, the values of case 3 are 0 in some utilization rates. Even in the worst cases, the values in case 2 are 6 times as large as those in case3. This analysis is the first of its kind and shows an advantage of using relay networks.

Lastly, we will discuss the cause of this difference. It is conceivable that this difference stems from the characteristics of relay networks. Utilizing nodes can send blocks to other utilizing nodes simultaneously via a relay network, and blocks produced by utilizing nodes are shared with other utilizing nodes rapidly. Two outcomes result. First, the utilizing nodes were less likely to continue hash calculation after new blocks were discovered. Second, blocks created by utilizing nodes were likely to win over blocks created by non-utilizing nodes. Indeed, this can be observed from the comparison of case 2 and case 3 in Table 7. For example, the probability of case 2 was 7.5 times as large as that of case 3 at a utilization rate

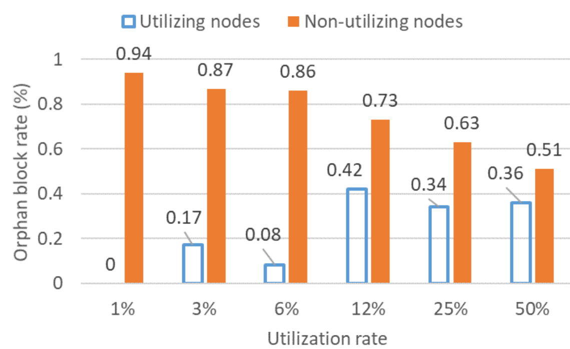

Fig. 7: Orphan block rates of utilizing nodes and non-utilizing nodes.

of 12%. These two factors were considered to contribute to the low orphan block rates of utilizing nodes.

V. Related Work

Prior work has investigated relay network effects. For example, Gervais et al. [7] reported that the orphan block rate decreased when all nodes utilized the relay network. However, this study did not consider the effects on individual nodes, and their simulation was performed only at utilization rates of 0% and 100%. Our prior study [8] reported that the relay network had good effects not only on the Bitcoin network but also on the utilizing nodes. However, the relay network model was quite simple and rough. Moreover, the hash power of each node was set uniformly in the simulation, which might have influenced the results. Thus, there is a need to confirm that the benefits that they found were maintained. In addition, some of the results were not thoroughly examined. For example, the factors responsible for the smaller orphan block rate of utilizing nodes compared with non-utilizing nodes were not sufficiently discussed. In addition, these previous studies [7], [8] performed simulations under the network state from approximately 2016, and the network state at that time, including the latency and bandwidth, was quite different from the current one. In this study, we used a more realistic relay network model and performed exact and in-depth analyses of each relay network effect in the latest network.

VI. Conclusion

In this study, we used a realistic relay network model and performed exact and in-depth analyses of the relay network effects in the latest network. The investigation of relay network effects under the most current network state is one of the major contributions of this study. In terms of the effects on the Bitcoin network, we observed a decrease in the orphan block rate and in the 50th percentile of block propagation time, as indicated in previous studies. However, a new finding is that the improvement of the orphan block rate by the relay network became smaller as the Internet speed increased. We also found that the relay network improved the 90th percentile of block propagation time, and this paper is the first to investigate the relay network effect on the 90th percentile of block propagation time. Moreover, the comparison of the 90th and 50th percentiles of block propagation time revealed that relay networks have a larger influence on the 50th percentile. Furthermore, the degree of block propagation time improvements were approximately the same for utilizing and non-utilizing nodes. However, the propagation times of utilizing nodes were always better.

Regarding the effects on individual nodes, we confirmed that blocks of utilizing nodes were quite unlikely to become orphan blocks even in the current fast network. Moreover, we investigated the outcome when non-utilizing and utilizing nodes produced blocks simultaneously. The result confirmed that blocks of utilizing nodes are far more likely to become main blocks than those of non-utilizing nodes, which seems to be a significant benefit of the use of relay networks by utilizing nodes. This finding is also a major contribution of this study. As this study and other studies have indicated, when more nodes utilize a relay network, the effects of the relay network are greater. Hence, it seems desirable to encourage more nodes to utilize the relay network, and our results may provide a motive for these nodes to do so. In terms of the mining success rate, we observed that the relay network did not exert a significant influence: the differences between utilizing and non-utilizing nodes were below 0.1 at any utilization rate.

Some limitations are worth noting. Although our results found effects of relay networks when they worked ideally, their effects when they are not working ideally, for example, because of heavy load, are not clear. Therefore, future work should include a follow-up investigation of the effects when relay networks do not work well. In addition, the effects on other blockchain networks, such as Ethereum, should be examined.

Acknowledgment

This work was supported by SECOM Science and Technology Foundation.

References

[1] "VisaNet — electronic payments network," https://usa.visa.com/about-visa/visanet.html (accessed Feb. 13, 2020).
[2] "PayPal reports second quarter 2018 results — paypal, inc." https://investor.paypal-corp.com/news-releases/news-release-details/paypal-reports-second-quarter-2018-results?ReleaseID=1072972 (accessed Feb. 13, 2020).
[3] C. Decker and R. Wattenhofer, "Information propagation in the Bitcoin network," in *Proc. IEEE P2P 2013*. IEEE, 2013, pp. 1–10.
[4] Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in bitcoin," in *International Conference on Financial Cryptography and Data Security*. Springer, 2015, pp. 507–527.
[5] "Falcon - A Fast Bitcoin Backbone," https://www.falcon-net.org/ (accessed Feb. 13, 2020).
[6] "Bitcoin relay network," http://bitcoinrelaynetwork.org/ (accessed Feb. 13, 2020).
[7] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," in *Proc. ACM CCS 2016*. ACM, 2016, pp. 3–16.
[8] K. Otsuki, Y. Aoki, R. Banno, and K. Shudo, "Effects of a simple relay network on the bitcoin network," in *Proceedings of the 15th Asian Internet Engineering Conference (AINTEC 2019)*. ACM, 2019, pp. 41–46.
[9] R. Nagayama, K. Shudo, and R. Banno, "Identifying impacts of protocol and internet development on the bitcoin network," in *Proc. 25th IEEE Symposium on Computers and Communications (IEEE ISCC 2020)*. IEEE, 2020.
[10] S. Nakamoto *et al.*, "Bitcoin: A peer-to-peer electronic cash system," 2008.
[11] "FIBRE Fast Internet Bitcoin Relay Engine," www.bitcoinfibre.org/ (accessed Feb. 13, 2020).
[12] T. Neudecker, "Security and anonymity aspects of the network layer of permissionless blockchains," Ph.D. dissertation, Ph. D. thesis, KIT, 2019.
[13] U. Klarman, S. Basu, A. Kuzmanovic, and E. G. Sirer, "bloxroute: A scalable trustless blockchain distribution network whitepaper," *IEEE Internet of Things Journal*, 2018.
[14] Y. Aoki, K. Otsuki, T. Kaneko, R. Banno, and K. Shudo, "SimBlock: A blockchain network simulator," in *Proc. CryBlock 2019 (in conj. with INFOCOM 2019)*, 2019.
[15] Y. Aoki and K. Shudo, "Proximity neighbor selection in blockchain networks," in *2019 IEEE International Conference on Blockchain (IEEE Blockchain 2019)*. IEEE, 2019, pp. 52–58.
[16] "Bitcoin Nodes Distribution - Bitnodes," https://bitnodes.earn.com/ (accessed Feb. 13, 2020).
[17] A. Miller, J. Litton, A. Pachulski, N. Gupta, D. Levin, N. Spring, and B. Bhattacharjee, "Discovering bitcoin's public topology and influential nodes," *et al*, 2015.
[18] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. G. Sirer *et al.*, "On scaling decentralized blockchains," in *International Conference on Financial Cryptography and Data Security*. Springer, 2016, pp. 106–125.