

Topological Evolution Analysis of Payment Channels in the Lightning Network

Gustavo F. Camilo¹, Gabriel Antonio F. Rebello^{1, 2}, Lucas Airam C. de Souza¹,
Maria Potop-Butucaru², Marcelo D. Amorim², Miguel Elias M. Campista¹, Luís Henrique M. K. Costa¹

¹Grupo de Teleinformática e Automação - Universidade Federal do Rio de Janeiro (UFRJ), Brazil

²Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

Abstract—Payment channel networks (PCN) offer a fast, secure, and distributed alternative payment method while avoiding slow consensus mechanisms of blockchains. Nonetheless, the PCN topology directly influences the performance, cost, and payment success rate. This paper¹ analyzes the evolution of the Lightning Network topology, which is currently the leading payment channel network. We reconstruct the network graph using real data from a set of channel announcement messages collected between January 2020 and August 2021. Our analysis uses typical graph metrics, such as transitivity, diameter, and degree centrality, to evaluate the state and evolution of the network. The results show a strong trend in resource and connectivity centralization. Only 0.38% of nodes concentrate 50% of the network capacity, exposing a vulnerability to targeted attacks. As with the Bitcoin cryptocurrency, the centralization of the Lightning PCN directly contrasts with the original goal of a fully-decentralized network. Moreover, the low network transitivity compromises channel rebalancing techniques, which contribute to the stability of the system. This trend evidences the need for new attachment policies prioritizing greater network decentralization and robustness.

Index Terms—layer-2 blockchain, payment channel networks, Lightning Network.

I. INTRODUCTION

The rapid growth of cryptocurrencies as an alternative payment method exposes the scalability problem of public blockchains. While traditional payment methods such as credit cards reach rates of 1,700 transactions per second, cryptocurrencies suffer from low throughput, averaging between 7 and 15 transactions per second on the Bitcoin and Ethereum networks, respectively [1]. As a consequence of the low performance, users pay expensive fees to miners to prioritize their transactions. In periods of high demand for payments on the network, average fees rise, reaching US\$ 62 in April 2021 [2]. Thus, the scalability problem of blockchains makes it impractical to use cryptocurrencies for routine small payments where rapid confirmation is required.

Payment channel networks (PCNs) present a solution to the cryptocurrency scalability problem [3]. By allowing transactions to occur off-chain in a secure and decentralized way, PCNs avoid slow consensus mechanisms and achieve high transaction throughput, charging low fees. Payment channel

networks have attracted the attention of both academia [4], [5] and the general public. PCNs are claimed to be one of the reasons for the recognition of Bitcoin as a legal currency in El Salvador [6]. Currently, the Lightning Network, Bitcoin’s payment channel network and main PCN, has over US\$ 130,542,759 allocated in over 86,000 public channels. It is worth noting that the Lightning Network allows channels to exist without being disclosed for privacy reasons. Thus, the actual number of channels in the network may exceed 86,000.

The great success of the Lightning Network and its consolidation as a leading PCN implementation motivates the analysis of its growth once the performance of routing and channel balancing proposals depends on the topological characteristics of the network [7], [8]. Even though the topology of the Lightning Network has been explored in other works [9], [10], [11], few of them analyze the evolution of network metrics. Furthermore, due to the rapid growth and high dynamism of the network, which allows easy opening and closing of channels, some analyses are already lagging behind the most recent state of the network [9], [10]. Thus, the analysis of network evolution and trends is essential to assess the feasibility of current and future routing proposals in PCN.

Contributions. This paper evaluates the evolution of the topological metrics of the Lightning Network, analyzing recent trends and implications for the future of the network. We use real information collected from gossip messages sent in the network to reconstruct the complete graph from January 2020 to August 2021. The results are split into two stages. First, we analyze the state of the network in August 2021, checking the distribution of resources and connectivity. The results show that, despite proposing fully-distributed transaction routing, the Lightning Network exhibits a high concentration of both resource and connectivity in a few nodes controlled by companies. Second, we analyze the time series of classical graph theory and complex network metrics to evaluate the evolution of the network structure. The analysis shows that the Lightning Network is becoming even more centralized and that some proposed channel rebalancing techniques are expensive or infeasible for most nodes in the network. The results also indicate the infeasibility of using current channel rebalancing techniques for much of the network. Furthermore, we show the vulnerability of the

¹A preliminary version of this paper was published in Portuguese and is available at <http://www.gta.ufrj.br/ftp/gta/TechReports/CRS22a.pdf>.

network by simulating targeted attacks. As a result, we find that an attack on the most central nodes in the network can result in 30% more expensive fees to the end user. Finally, the paper discusses solutions to the challenges presented.

II. PAYMENT CHANNEL NETWORKS

Payment channels are a blockchain-based technology that allows two users to make recurring payments to each other privately and in real time. The main assumption of this technology is that the consensus protocol, while being the primary mechanism that ensures the security of the blockchain, is also the leading cause of delays and excessive fees in the system. Thus, the basic idea of a payment channel is to perform as many transactions as possible directly between the parties involved, without publishing them in a block, and use the consensus protocol only when necessary. Off-chain transactions require only a signature from each party involved.

To open a payment channel, two users, A and B, must sign and publish a transaction that reserves an amount of tokens at a wallet address shared by both users. Once the initial funding transaction has been published, A and B can rebalance the funds in the address by privately exchanging commitment transactions with each other. Thus, both users can instantly issue payments directly, within the reserved limit, without paying fees to the blockchain miners. When they wish to close the channel, A or B can post the most recent commitment transaction to the blockchain and wait for the consensus protocol delay to recover the tokens they invested in the channel. This approach allows day-to-day recurring payments, such as a grocery purchase or monthly service payments, which require agility and are of low value.

Opening a payment channel allocates tokens that cannot be used in the blockchain. Thus, it becomes infeasible for a person to open payment channels with all other users in the network. Therefore, the idea of a PCN is that users employ pre-existing channels to route payments through intermediaries. The set of established payment channels among users forms a payment channel network (PCN), illustrated in Figure 1. Each link has a capacity, illustrated inside the rectangle on the edge of the figure. This capacity indicates the token amount allocated by the parties on that channel. This information is accessible and is stored in the funding transaction published in the blockchain. The balances of each party on the channel, represented by the numbers at each end of the edge, indicate the current state of the link. The balances are private to the involved parties for security and scalability reasons. On the one hand, for security, one could trace payments if the balances were public. On the other hand, for scalability, it would be necessary to flood the network with update messages with each payment.

The Bitcoin cryptocurrency has the first and largest implementation of PCN, in addition to pioneering the implementation of blockchain technology. The Lightning Network held its first transaction in 2017, reaching over 13,000 nodes and 86,000 payment channels distributed worldwide today.

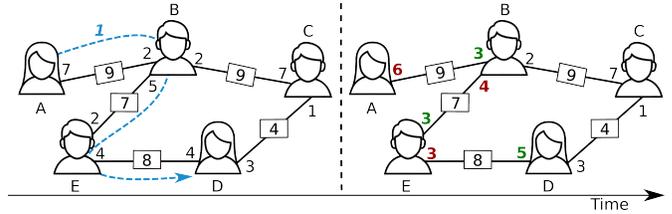


Fig. 1: A payment of 1 token occurring from user A to user D in a payment channel network. The payment travels along the path indicated by the sender, changing the balances on each channel on the path. The channel capacities, indicated in the rectangles, represent the total reserved by the two parties and are constant throughout the lifetime of the channel.

III. ANALYSIS OF THE LIGHTNING NETWORK

The default routing of the Lightning Network is source-based routing. Thus, Lightning requires that nodes know the full topology of the network to make payments, information which can be gathered from announcement messages. In the present paper, we use a dataset collected by Lightning Network developers of channel announcement messages transmitted using a gossip protocol for node synchronization in the Lightning Network [12]. The dataset contains three types of messages collected by a node using the Lightning Network implementation in C language, called *c-lightning*: (i) channel announcement messages, which make a channel public in the network; (ii) channel update messages, used to update the parameters of a given channel; and (iii) node announcement messages, used to inform the metadata of a node. To create the dataset, raw messages were stored in a file and a filter was used to remove duplicate messages. The file contains messages from the Lightning Network from August 2018 through August 2021. Given that the period 2018 and 2019 has already been covered by other papers [9], [10], we consider the period between January 2020 and August 2021. All messages have a timestamp which tells the time at which they were created using the UNIX system time base. Thus, using a time window, it is possible to reconstruct the topology of the Lightning Network at any instant of time. This work uses two-week time windows for topology analysis, i.e., it ignores messages prior to a two-week period when reconstructing the network at a given instant.

Network state. The first set of experiments checks the most current state of the message dataset, from August 17 to August 31, 2021, to estimate the resource concentration and connectivity in the Lightning Network. Figure 2 shows the distribution of channel routing capacity in the network. The channel routing capacity is the total number of tokens allocated by the two parties during its creation. We note an intense concentration of funds in the network, where few channels have high capacity and multiple channels have low resources. The capacity of the channels is a decisive factor for transaction routing. The low amount of high-value channels, e.g., only 13 channels (0.02% of the network) with capacity

above US\$ 45,863.00 (≈ 2 BTC), restricts the amount of paths that high-value payments can use. Although the value seems high, payments of 2 BTC or more are not uncommon on the Bitcoin Network [13]. By concentrating high-capacity channels on a few paths, most of the network must resort to blockchain to make high-value payments, experiencing high confirmation time and high fees.

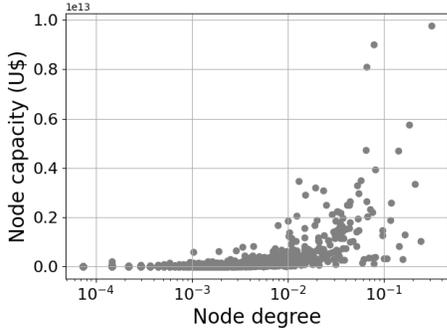


Fig. 2: Relationship between node degree and node capacity in the Lightning Network.

The concentration of routing power becomes even more evident by removing the higher-degree nodes. Figure 3 simulates an attack targeting the highest-degree nodes in the network, removing them, and illustrating the consequences of this attack. In the Lightning Network, this removal means closing all channels of a node or making it unavailable for routing for a time interval through a DoS attack. The removal of 49 nodes, approximately 0.38% of the network, is enough to cut the capacity of the network by half. A possible attack affects the entire network, increasing the probability of failure due to lack of available paths with sufficient capacity. Seres *et al.*'s similar experiment on the Lightning Network shows that in 2019, it was enough to remove 37 nodes to halve the network capacity [9]. It is worth noting, however, that in the period evaluated by Seres *et al.*, the network had only 2,344 nodes. Thus, we see that even after the number of nodes in the network have more than tripled by August 2021, the network is still extremely vulnerable to attacks targeted at central nodes. Furthermore, removing 100 nodes, approximately 0.74% of the network, increases the base fee paid by the end user by approximately 32%. The value of the average base fee paid by the user is calculated from the increase in the average number of hops when removing nodes assuming the base fee is homogeneous for all nodes. Since the base fee is fixed and paid at each hop, it is possible to verify the financial effect of this type of attack on the end user. The possibility of such attacks is not only theoretical. The Lightning Network has already suffered a DoS attack that affected 20% of the nodes in the network, making them unavailable for routing [14]. In this way, an attack targeting the highest degree nodes in the network directly affects all participants, who then pay more fees due to the increased average path length, as well as increasing the probability

of unsuccessful payment attempts due to lack of routing capacity.

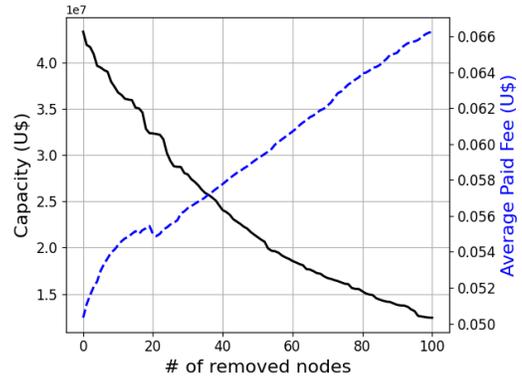


Fig. 3: Capacity and average base rate when removing highest degree nodes.

Network Evolution and Trends. The second set of experiments checks the evolution of network metrics to assess whether the degree of network concentration has been increasing. The experiments evaluate the network in the period between January 21, 2020 and August 31, 2021. Figure 4 shows the evolution of the Lightning Network metrics. All metrics consider only the largest component of the graph, except the number of components result. Although the number of components in the network grows from 6 to 53 during the evaluated period, as shown in Figure 4a, the number of nodes of the smallest components is negligible compared to the number of nodes of the largest component. As an example, the network snapshot with the largest number of components has 13,462 nodes in the largest component, while the smallest components have at most 3 nodes. Thus, using only the largest component does not affect the reliability of the results, given that the smallest components do not make payments with more than 1 hop and cannot meet the demands of most of the network. The increase in the number of components is natural due to the growth in the number of nodes in the network.

The assortativity r of a graph indicates the connection preference of nodes in the network. A disassortative graph has $r < 0$, indicating that low-degree nodes in the network prefer to connect to high-degree nodes, while an assortative graph has $r > 0$, indicating that high-degree nodes prefer to connect to high-degree nodes. The network's assortativity coefficient, in Figure 4b, shows that the Lightning Network is disassortative, indicating that low-degree nodes tend to connect to high-degree nodes. This is because some Lightning Network implementations, such as Lightning Daemon (LND), adopt policies that connects new nodes to more central nodes in the network [9]. This strategy is adopted because nodes that intend to make payments to more than one entity may prioritize connecting to more central nodes, seeking lower paths and fees. Despite the disassortativity of the network, we can see a growth tendency in this metric, pointing to

an assortative network in the future. This can be explained by the establishment of connection between the central hubs after they consolidate, increasing the assortativity. The upward trend in assortativity can be explained by the emergence of new hubs as a result of the rapid growth of the network.

The graph transitivity measures the fraction of cycles existing in the network relative to triads, the number of cycles that could exist in the network. Triads are measured by counting the number of edge-pairs (e_1, e_2) that share the same vertex v_1 . Thus, the transitivity measure only considers 3-node cycles and varies between zero and one. The transitivity T of a graph is given by $T = 3\frac{l_3}{t}$, where l_3 is the number of cycles between 3 nodes in the network and t is the number of triads. In PCNs, establishing cycles of few hops is crucial to the cheapness of some channel rebalancing proposals [7], [8]. These proposals take advantage of cycles so that a node can make payments to itself, rebalancing paths by moving money between channels. Keeping channels balanced is of utmost importance to increase the probability of a successful transaction [4]. Nevertheless, the low transitivity, presented in Figure 4c and its downward trend in the network point to the existence of few 3-node cycles, which makes proposals expensive and possibly unfeasible for most participants who have low financial power. This low transitivity is a natural consequence of the establishment of connections between low-degree nodes to high-degree nodes, generating several possibilities of cycles that are not formed. The centralized topology tends towards a core-periphery structure, in which most of the nodes are at the extremities, i.e., at the periphery. Thus, these nodes have only one neighbor, which does not contribute to the formation of triangles in the network, lowering transitivity.

The density of a graph is given by the fraction of the edges existing in the graph and the edges of a complete graph with the same number of nodes. Thus, the density of an undirected graph can be calculated by $d = \frac{m}{n(n-1)}$, where m is the number of edges in the network and n is the number of nodes. The low density observed in Figure 4d shows that the Lightning Network is an extremely sparse network. The downward trend in density can also be explained by the connection of new nodes to more central nodes. This type of behavior is common in scale-free networks, where a small portion of nodes concentrate a large part of the connections. The behavior of the Lightning Network as a scale-free network can also be observed in the drop in the average network degree and by the s-metric result, as shown in Figures 4e and 4f. The entry of one-connection only nodes contributes to the drop in the average network degree even as the degree of the most central nodes increases. The s-metric ranges from zero to one and measures how much a graph has a core of nodes that increases the connectivity of the network, resembling a hub. Despite the drop in s-metric, the Lightning Network still has a concentration of connectivity in a few nodes, as shown in Figure 4e.

Figures 4g and 4h expose the low diameter and average

shortest-path in the Lightning Network. The diameter of a graph is defined by the longest shortest path between any pair of nodes. The average shortest path is obtained through the sum of the minimum distance for each pair of vertices of the graph divided by the total number of vertices. Despite the number of nodes in the network tripling in the period, the two metrics present low variation if compared to the other metrics, varying by 28% in the case of the diameter and approximately 15% in the case of the shortest path.

The centralization trend becomes even clearer when analyzing Figure 5, which shows the resource concentration of the nodes over time. It can be seen that while in January 2020 approximately 2.0% and 0.9% of nodes concentrated 80% and 50% of the total network capacity, respectively, the numbers dropped to 1.1% and 0.3% in August 2021. The downward trend, albeit smooth, shows that a small set of nodes concentrate much of the allocated revenue and exposes future vulnerabilities and security challenges in the network.

Discussion. Despite presenting itself as a fully decentralized alternative for fast payments in blockchains, the results show a clear centralization of resource and network connectivity. The resource concentration is evident when analyzing the distribution of capacity and channel degrees in Figure 2. Higher capacity nodes hold a monopoly on routing high-value payments in the network, getting richer by collecting more fees. These higher capacity nodes attract connections from new participants, creating hubs, as new participants seek to open channels in a location that can reach multiple destinations in a smaller number of hops. Thus, participants choose to connect to the more central nodes seeking to lower the routing fee paid when making a payment.

Low connectivity of nodes in the periphery also makes proposals that take advantage of cycles to rebalance channels unfeasible or expensive in practice. These proposals are restricted to central nodes, that have greater connectivity and are more likely to find paths with cycles for rebalancing. While nodes with higher connectivity can easily rebalance their own channels using the network topology, most network participants must resort to the blockchain. This phenomenon further accentuates network inequality, given that central nodes are less likely to experience periods of routing unavailability when waiting for transaction confirmation in the blockchain, nor do they pay high fees to rebalance their own channels.

Possible solutions to the centralization problem include network-attachment proposals that prioritize distributing connectivity when opening channels. In this case, the creation of alternative paths in the network ensures less dependence of peripheral nodes on centralized nodes and reduces the effects of split attacks on the network. Furthermore, connecting new nodes by creating cycles in the topology mitigates the impact of problems related to low transitivity, such as rebalancing channels through the network topology. Thus, generating cycles when opening channels ensures that nodes can rebalance their channels without having to resort to blockchain. This solution, however, presents a common trade-off. By

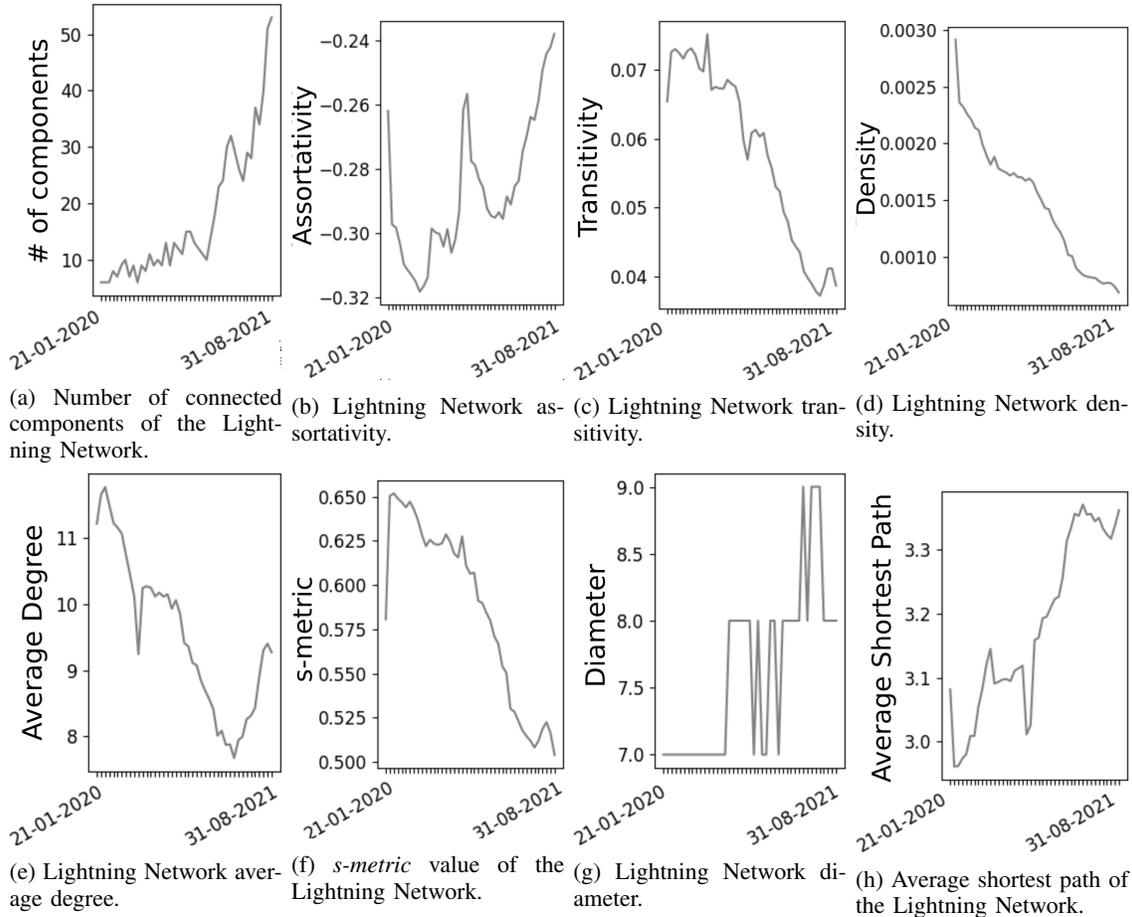


Fig. 4: Evolution of the graph metrics of the Lightning Network over the period from January 21, 2020 to August 31, 2021.

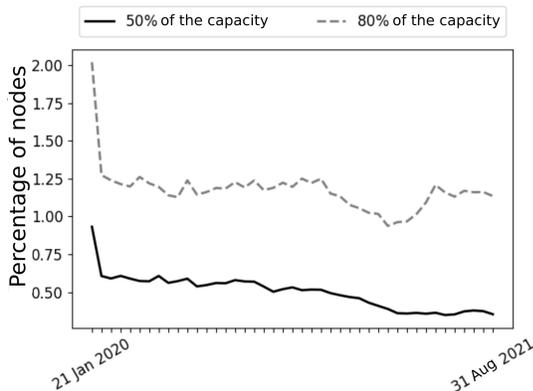


Fig. 5: Resource concentration in the Lightning Network from January 2020 to August 2021. The dashes indicate the percentage of nodes that concentrate 50% and 80% of the total network capacity in each two-week period.

modifying the connection preference of nodes to peripheral nodes, the network becomes more robust to some targeted attacks, but also grows in diameter, resulting in longer paths and more fees paid by users.

IV. RELATED WORK

The rapid growth of payment channel networks has motivated several research works. Most works seek to solve problems related to efficient transaction routing [4], security and privacy in PCNs [15], [5], and channel rebalancing [7], [8]. Due to the characteristics of transaction payment routing in a PCN, the proposal of efficient and secure routing and balancing algorithms strongly depends on the network topology. Thus, knowledge of the topological characteristics of a PCN, such as the Lightning Network, is indispensable to evaluate the practicality and feasibility of proposals.

Some papers investigate the Lightning Network topology [9], [10], [11], [16]. Seres *et al.* analyze the Lightning Network topology using graph metrics and a network snapshot from January 2019 [9]. From the analysis, the authors conclude that the Lightning Network exhibits behavior similar to scale-free networks and small-world networks, and exhibits robustness against random failure attack and vulnerabilities to rational attackers. The authors evaluate the state of the network in January 2019, without considering historical growth. Due to the rapid growth and high dynamicity of the network, the results are limited to the evaluated date and do not

consider the changing growth prospects of the network, as demonstrated in the present work. Moreover, Seres *et al.* analysis does not consider some financial parameters, such as the fees paid by users and the direct effect of attacks directed at the end user. Lin *et al.* evaluate the growth of the Lightning Network in relation to the resource concentration of the network [10]. The authors evaluate indices such as Gini, proximity centrality, and betweenness centrality of the network over the period between January 2018 and July 2019. The results show that the network exhibits core-periphery type structures where the core is formed by hubs and star-like substructures, verifying a strong centralization in the network around the highest capacity nodes. The authors expose that removing hubs would lead to partitioning of the network into multiple components, leaving the network more vulnerable to attacks. Similarly, Beres *et al.* analyze Lightning Network metrics, such as number of edges, average degree, and transitivity, over the period between January 2018 and May 2019 [16] to build a traffic simulator. The papers, however, are limited to the analysis of few centralization metrics, disregarding other metrics, such as evolution of density, components, s-metric and assortativity, and the evolution and distribution of channel characteristics. These metrics allow for a greater understanding of the state of the network and its prospects. Rohrer *et al.* focus on security analysis, quantifying the resilience of the Lightning Network to topology-based attacks [11]. The authors show that the network is susceptible to channel exhaustion and node isolation attacks, which affects the payment success rate and the average payment flow in the network. The result shows that attackers with available resources should follow the strategy of removing higher centrality nodes to perform effective attacks, while nodes with low resource should focus on low capacity channels whose removal partitions the graph.

Unlike above cited articles, this work evaluates the evolution of centralization of the Lightning Network based on graph-theory metrics and complex networks, as well as evaluating the result of targeted attacks and their effects on end users. The analysis uses real data collected from Lightning Network messages from January 2020 to August 2021 and evaluates the centralization trend of network connectivity and resource, and discusses the consequences and trends of the current growth model of the network for future proposals.

V. CONCLUSION

Payment channel networks present a fundamental mechanism for enabling fast, secure, and decentralized cryptocurrency payments in everyday life. Because of the way the routing of payments through payment channels is implemented, PCNs are heavily influenced by their topology. This paper analyzes the evolution of the topology of the Lightning Network, the main payment channel network, using graph theory metrics and real messages collected from the network. The results show that the connections and capacity allocated to the network are extremely concentrated in a few nodes, which

are mostly controlled by companies linked to blockchain development. This centralization is a natural consequence of how the entry of new participants takes place, seeking smaller paths to perform transactions and, consequently, pay lower fees. In addition, the paper finds that the results indicate an evolution towards a weakly transitive network, which makes channel rebalancing techniques unfeasible for most nodes.

VI. ACKNOWLEDGMENT

This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. This paper was also funded by CNPq, FAPERJ, and Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), Grants 2015/24494-8, 2018/23292-0, 2015/24485-9, 2014/50937-1.

REFERENCES

- [1] BitcoinWiki, “Bitcoin Scalability,” <https://en.bitcoin.it/wiki/Scalability>, 2019, last access: Jan 18th 2022.
- [2] BitInfoCharts, “Bitcoin Average Transaction Fee Chart,” <https://bitinfocharts.com/comparison/bitcoin-transactionfees.html>, 2022, last access: Feb 3rd 2022.
- [3] J. Poon and T. Dryja, “The bitcoin lightning network: Scalable off-chain instant payments,” 2016.
- [4] V. Sivaraman *et al.*, “High throughput cryptocurrency routing in payment channel networks,” in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, 2020, pp. 777–796.
- [5] G. A. F. Rebello, M. Potop-Butucaru, M. D. de Amorim, and O. C. M. B. Duarte, “Securing wireless payment channel networks with minimum lock time windows,” in *IEEE International Conference on Communications (ICC)*, 2022.
- [6] CloudTweaks, “How Bitcoin Brought The Lightning Network To El Salvador,” <https://cloudtweaks.com/2021/07/how-bitcoin-brought-lightning-network-el-salvador/>, 2021, last access: Feb 3rd 2022.
- [7] R. Pickhardt and M. Nowostawski, “Imbalance measure and proactive channel rebalancing algorithm for the lightning network,” in *2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 2020, pp. 1–5.
- [8] R. Khalil and A. Gervais, “Revive: Rebalancing off-blockchain payment networks,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’17. New York, NY, USA: Association for Computing Machinery, 2017, p. 439–453.
- [9] I. A. Seres, L. Gulyás, D. A. Nagy, and P. Burcsi, “Topological Analysis of Bitcoin’s Lightning Network,” *arXiv preprint arXiv:1901.04972*, 2019.
- [10] J.-H. Lin, K. Primicerio, T. Squartini, C. Decker, and C. J. Tessone, “Lightning network: a second path towards centralisation of the bitcoin economy,” *New Journal of Physics*, vol. 22, no. 8, p. 083022, 2020.
- [11] E. Rohrer, J. Malliaris, and F. Tschorsch, “Discharged payment channels: Quantifying the lightning network’s resilience to topology-based attacks,” in *IEEE EEuroS&PW*, 2019.
- [12] C. Decker, “Lightning network research; topology datasets,” <https://github.com/lnresearch/topology>, 2021, last access: Dec 29th 2021.
- [13] Blockchain.com, “Unconfirmed transactions,” <https://www.blockchain.com/btc/unconfirmed-transactions>, 2022, last access: Aug 5th 2022.
- [14] TrustNodes, “Lightning Network DDoS Sends 20% of Nodes Down,” <https://www.trustnodes.com/2018/03/21/lightning-network-ddos-sends-20-nodes>, 2018, last access: Feb 3rd 2022.
- [15] S. Roos, P. Moreno-Sanchez, A. Kate, and I. Goldberg, “Settling payments fast and private: Efficient decentralized routing for path-based transactions,” *arXiv preprint arXiv:1709.05748*, 2017.
- [16] F. Beres, I. A. Seres, and A. A. Benczur, “A cryptoeconomic traffic analysis of Bitcoin’s lightning network,” 2019. [Online]. Available: <https://arxiv.org/abs/1911.09432>