2018 IEEE 24th International Conference on Parallel and Distributed Systems (ICPADS)

Optimal Fee Structure for Efficient Lightning Networks

Alvin Heng Jun Ren*†, Ling Feng*‡, Siew Ann Cheong†§ and Rick Siow Mong Goh*
*Computing Science Department
Institute of High Performance Computing, A*STAR Singapore 138632
†School of Physical and Mathematical Sciences
Nanyang Technological University, Singapore 637371
‡Department of Physics
National University of Singapore, Singapore 117551
Email: felney@gmail.com
§Complexity Institute
Nanyang Technological University, Singapore 637723

***Abstract*—Off-chain transaction handling like the Lightning Network (LN) is among the most promising solutions to solve the scaling challenges in blockchain technology like Bitcoin. At the same time, the LN faces its own challenges like transaction path lengths, centralization of channels (hubs), channel imbalances (depletion), etc. Here we study the effects of payment channel fees on these various factors. To get realistic insights, we based our study on empirical Bitcoin transaction patterns and existing LN structure, and apply a simple form of fee structure with only one tunable parameter, α, that is most influential on the transaction routing paths. We assume the transactions through the LN take the path of the lowest aggregate fees, and found as a consequence that one cannot have short average path lengths and low overall channel imbalances at the same time. A good compromise is to have fees proportional to the square root of the channel capacity, such that reasonably short path lengths and overall balanced channel capacities can be achieved that makes the operation of the LN more sustainable.**

I. Introduction

The emerging technology of blockchain has great promises in internationalizing many industries through mass decentralization. As the first successful example, Bitcoin (BTC) was designed as a decentralized digital currency by Satoshi Nakamoto in 2008 [1]. The Bitcoin network (BN) operates on a public ledger, known as the blockchain, which records every transaction that happens in the network, and verified by a public pool of 'miners'. The popularity, usage and price of BTC has skyrocketed since its inception, with growing awareness on the benefits of blockchain technology. However, its design is not able to meet the increasing number of transactions, which today results in long confirmation times and high transaction fees. Such challenge is not unique to Bitcoin, but universal for almost all public blockchain technologies at large.

Different solutions were proposed to increase the blockchain's scalability [2]–[4]. One of the solutions that have been implemented is the Lightning Network [4], an off-chain solution that proposes moving the bulk of BTC transactions off the blockchain ledger to a secondary network, where users create bi-directional committed channels that allows funds to be exchanged almost instantaneously and with minute fees. The public blockchain will serve as the settlement layer for opening and closing of channels, while transactions happen within the LN. Currently, the LN is still in its early experimental stages, with many potential challenges ahead.

In this work, we study the feasibility of the LN in being a global network that complements the BTC ledger, focusing on the interplay among three factors: transaction fees, transaction routing path lengths and channel imbalances. Transaction fee on each channel affects the optimal path a transaction takes, and in turn affects the imbalances of the channels on the path taken. An ideal LN protocol would have low transaction fees, short path lengths for transactions, and balanced channels for sustainable operations.

To carry out the study, we first examine the empirical features in the BN transactions, as well as the LN structure. Then we carry out realistic simulations of BTC transactions on LN based on these empirical features for different fee structures on channels. Some tools from network science [5], [6] are also employed in the study, in particular the analysis of centrality measures on the channels and their imbalances, and the overall resilience [7]. In the end, we propose an optimal fee structure for LN that strikes a balance between short transaction path and channel imbalances.

II. Empirical Observations of BTC and Lightning Network

We first carry out empirical analysis on Bitcoin transaction patterns, which will serve as the inputs to simulate realistic transaction flows on the LN.

A. Bitcoin Network

We use Bitcoin transaction data from the open-access ELTE Bitcoin project [8], whose raw data was retrieved directly from the Bitcoin Core client. Due the to sheer amount of transactions, we use only 1-day transactional data on 7 Feb 2018 and model it as a directed network. Each node in the network is a unique BTC address and each edge represents a

978-1-5386-7308-9/18/$31.00 ©2018 IEEE
DOI 10.1109/ICPADS.2018.00130

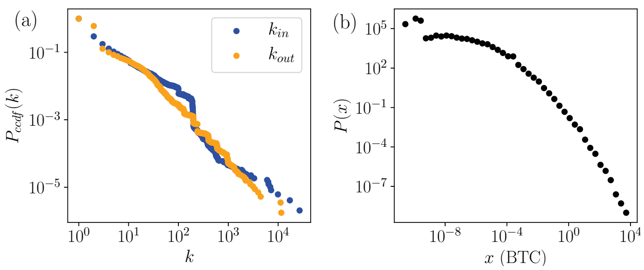

Fig. 1. (a) Complementary cumulative distribution functions (CCDF) of in degree (k_{in}) and out degrees (k_{out}) of the BN. Both display heavy-tailed power-law behavior $P(k) \sim k^{-\gamma}$, with $\gamma = 2.8$ and 2.4 respectively. (b) Probability density function of the transaction sizes, x, in BTC of the BN. Data points are log-binned. The tail above $x \sim 10^{-4}$ scales to a power-law with exponent = 1.6.

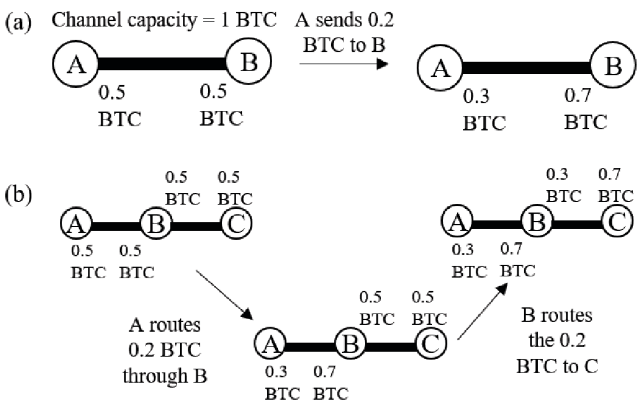

Fig. 2. (a) Illustration of a transaction in the LN between Alice (A) and Bob (B). Initially, both A and B commit 0.5 BTC into the channel totalling 1.0 BTC, and the channel imbalance $I_{cap} = 0$. After A transferred 0.2 BTC to B, the new channel imbalance is $I_{cap} = 0.4$. (b) Illustration of transaction routing of 0.2 BTC from A to Charlie (C) through B when A and C do not share a direct channel.

transaction between two nodes. Although the data is only from a single day, the network obtained consists of 814488 nodes and 2096150 edges, which is sufficient to obtain the statistical patterns of the transaction activities, and such patterns do not change significantly. Later, we will simulate artificial transactions based on the patterns extracted from the 1-day data.

Fig. 1(a) shows the in and out degree distributions of the BN. Here 'out degree' is the number of transactions a node has sent out, and 'in degree' is the number of transactions a node has received. We observe a power-law distribution that is characteristic of scale-free networks [9]. The distribution can be fitted to $P(k) \sim k^{-\gamma}$ with $\gamma = 2.8$ for the in-degree and $\gamma = 2.4$ for the out-degree, similar to that of the Internet [6]. Fig. 1(b) shows the distribution of transaction sizes in the BN in terms of BTC plotted as a probability density function. Again, a power-law behavior can be observed in the tail that scales with an exponent of 1.6.

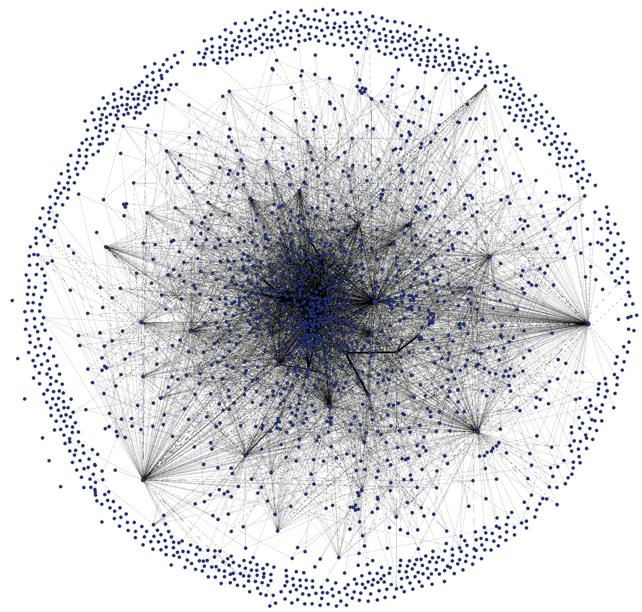

Fig. 3. The Lightning Network on 13 June 2018. Blue dots are nodes and black lines are channels. Width of channels are proportional to channel capacity. The nodes on the periphery are disconnected from the network, therefore we remove them in our simulation as they are not able to facilitate transactions. There are 1355 nodes in the connected network and 889 nodes in the periphery.

B. Lightning Network

The LN is a secondary layer beyond the public blockchain of Bitcoin. It consists of 'channels' between pairs of Bitcoin accounts, with each channel having a certain amount of BTC transferred into it. The total amount committed by the two accounts of a channel constitutes the total capacity of the channel. Any transaction between the two accounts smaller than the capacity can happen without going through the blockchain, since there is sufficient amount within the channel to facilitate the transaction.

Fig. 2(a) illustrates such a scheme in the LN. Assuming both Alice and Bob each commit 0.5 BTC to opening a channel, each of them initially has 0.5 BTC on their side of the channel. If Alice wishes to send 0.2 BTC to Bob, Alice's balance is reduced by 0.2 BTC, while Bob's balance is incremented by 0.2 BTC. This implies two properties of the LN. Firstly, transaction sizes are limited by the user's balances. For example, with only 0.3 BTC left in Alice's balance, she cannot send a transaction of more than 0.3 BTC to Bob. Secondly, LN channels can become imbalanced should transactions occur frequently in one direction. Should Alice continuously pay Bob, eventually Alice empties her balance and funds can no longer flow in the direction of Alice to Bob.

Fig. 2(b) illustrates the second important characteristic of the LN: transaction routing. If two nodes in the LN do not share a direct channel, transactions can still be sent by routing via intermediate nodes. In Fig. 2(b), Alice can route 0.2 BTC

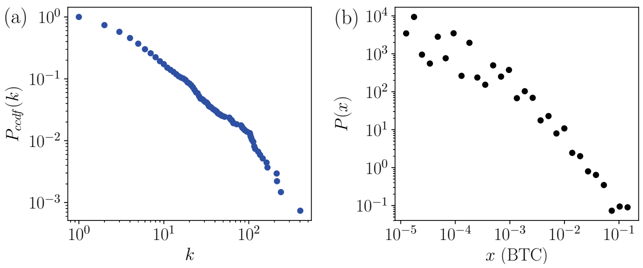

Fig. 4. (a) Complementary cumulative distribution function of the LN. Similar to the BN, it displays a heavy-tailed power-law behavior with $\gamma = 2.1$. (b) Probability density function of the channel capacities, x, in the LN. The tail beyond $x \sim 10^{-3}$ scales to a power-law with exponent = 1.7.

to Charlie through Bob despite not having a direct channel opened with Charlie. Note that this is a simplified illustration of the routing mechanism, and the proper implementation requires advanced cryptographic mechanisms that are being studied [10], [11].

Data of the LN is obtained by crawling an online LN explorer [12]. A snapshot of the LN is taken on 13 Jun 2018 and data is retrieved from the JSON contained in the explorer. Each node in the network represents a Lightning node, and each edge represents a Lightning channel between two nodes. Due to the bi-directional nature of LN channels, we model the LN as a network where a channel is two edges in opposite directions, and each edge has equal capacity to enable the flow of payments with half of the channel's total capacity.

A visualization of the real LN is shown in Fig. 3. The visualization shows a non-negligible number of 'Lightning hubs', which are nodes of large degrees that can serve as central hubs for transaction routing. On the periphery, many disconnected (0 degree) nodes are observed. One explanation for these disconnected nodes is that during the testing phase of the LN, many users run nodes for experimental purposes without committing funds to open channels. These 0 degree nodes thus do not contribute at all to the network dynamics, and consequently we remove them. After the removal, the network consists of 1355 nodes and 5592 edges.

Similar to the BN, the LN exhibits a power-law behavior as shown in Fig. 4(a). We obtain the exponent $\gamma = 2.1$ when fitting the distribution to the power-law function. This implies that the LN also classifies as a scale-free network, supporting the visual observation of the hubs in Fig. 3. The fact that the LN's degree distribution has exponent $\gamma = 2.1$ indicates that it is resilient against random failures of nodes, such that random removal of nodes would not significantly breakdown the network [7]. This is because majority of nodes in the LN are low degree nodes, whose removal will not significantly affect the network connectivity. However, targeted attacks on hubs can cause the network to quickly breakdown, with the giant component size rapidly disintegrating. This is due to the fact that in scale-free networks, the few high-degree nodes hold the network in place - their removal will destroy a large number of central links, rapidly reducing the network's connectivity.

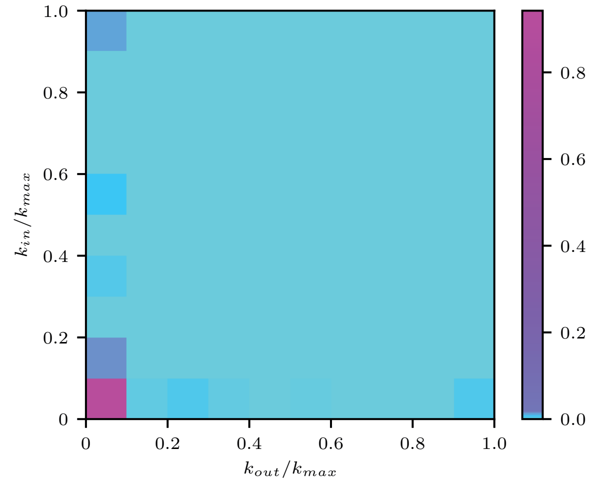

Fig. 5. 2D heatmap of the degree distribution of the nodes at the ends of the edges in the BN, i.e. k_{out} and k_{in} are degrees of the sending and receiving nodes of a transaction in BN. The degrees are normalized by the maximum degree in the BN, therefore the horizontal and vertical axis $k_{out}/k_{max}, k_{in}/k_{max} \in [0, 1]$. The value in each bin is the probability to find an edge connecting two nodes with the relevant degrees.

In Fig. 4(b), we plot the distribution of channel capacities in the LN. Firstly, we note that the channel capacities are very small, with the largest being $\sim 10^{-1}$ BTC - a far cry from the $\sim 10^{4}$ BTC transactions in the BN. This behavior is likely due to the experimental nature of the LN, such that users are unwilling to commit large amounts of BTC for a network that is still in testing. Secondly, the probability density function also scales as a power-law. However, the exponent of 1.7 is smaller than that in the BN, which suggest large channel capacities are relatively very common in the LN.

III. Simulation Methodology

A. Transaction sampling

There have been numerous studies on quantifying the BTC transaction graph [13]–[16], while studies quantifying the LN are few, if any. If the LN were to fulfil the vision of handling the majority of BTC transactions, with the main blockchain serving as the consensus layer for opening and closing channels, simulating realistic BTC transactions on the LN will provide us a first study on the feasibility of such a proposition.

To do so, we propose a way to draw a sample of BTC transactions from real historical data of the BN to simulate on the LN. There are two steps to this problem: 1) determining which two LN nodes to simulate a transaction between and 2) determining the size of the transaction.

For the first step, we probabilistically choose two LN nodes in accordance to the degree distributions of nodes in the BN. Fig. 5 shows a 2D histogram of the probability distribution of the transactions (links) for different degrees of sending node k_{out} and receiving node k_{in}. Each square is binned, whose value at the square is the probability of finding an edge in the BN, such that the two end nodes are of the degrees at the respective axes. From Fig. 1(a) and Fig. 4(a), we observe from the power-law tails that nodes in the BN have much larger degrees than nodes in the LN. To be able to make degree comparisons between the two networks, we normalize all degrees by the maximum, k_{max}, which guarantees $k/k_{max} \in [0, 1]$ for both the BN and LN. This allows us then to make a mapping between the nodes in the BN to LN.

In the second step, we determine the transaction size based on real BN transactions. As we have seen from Fig. 1(b) and Fig. 4(b), transaction sizes in the BN vary up to $\sim 10^4$ BTC, while channel capacities in the LN only vary up to $\sim 10^{-1}$ BTC. Hence, like the first step, we normalize the transaction sizes in Fig. 1(b) by the maximum transaction size, so that $x/x_{max} \in [0, 1]$. We do the same for channel capacities in the LN. Now that both quantities lie in the same range, a direct comparison can be made. We can then draw normalized transaction sizes from the probability density function of Fig. 1(b) after normalizing by x_{max}.

B. Simulation setup

Data available on the LN consists only of the total channel capacities, not the exact balance at the two ends of the channel. Hence, without loss of generality, we initialize our LN such that every channel is equally balanced (each user has half the total capacity). We keep track of channel balances in our simulation as illustrated in Fig. 2. If a transaction occurs in a channel, the channel balances of the two users are updated accordingly.

As mentioned, routing of transactions is a defining characteristic of the LN, and we allow for routing in our simulations. The route is selected such that it is the weighted shortest path [17], [18], where every channel in the route has sufficient capacity to route the transaction. If no such path exists, the transaction is taken as failed. The weighted shortest path is defined as the lowest aggregate fees in the path, which we define next.

C. Fee structure

It is necessary for users to pay fees to routing nodes for routing their transactions, so as to create incentive for nodes to maintain connectivity in the LN. Hence, fee structure charged by the intermediate nodes will affect the route of transactions, which in turn affect the imbalance of the channels. Here we define the channel imbalance I_{cap} as:

$$I_{cap} = \frac{|C_1 - C_2|}{C_1 + C_2},$$

where C_1 and C_2 are the capacities of the two users in a channel. If a channel is equally balanced, then $I_{cap} = 0$, while a completely imbalanced channel results in $I_{cap} = 1$. For example, initially $I_{cap} = 0$ for the channel between Alice and Bob in Fig. 2(a), while $I_{cap} = 0.4$ after Alice sends the transaction to Bob.

Imbalance of Lightning channels is undesirable, as it renders the channel unusable in a certain direction. The channel either requires re-balancing, or be closed and re-opened by the two users, which incurs on-chain fees. Channels with large capacities are less likely to become imbalanced than channels with smaller capacities as it takes more transactions to cause an imbalance. As such, we can propose a fee structure whereby channels with large capacities charge smaller fees. Channels with smaller capacities charge larger fees, as the same transaction would cause a larger imbalance for smaller channels. A fee structure that minimizes I_{cap} is optimal, i.e. the channel can be rebalanced to avoid closure.

If we treat the fees as weights of the edges, we would like large capacity channels to have smaller weights than small capacity channels. Hence, a good choice of fee structure is to let the weight be $1/x$, where x is the channel capacity [19]. The aforementioned weighted shortest path used for routing then has a natural interpretation: the shortest path in a route is one that minimizes the fees incurred for the user.

It must be noted that such a fee structure might result in paths that incur many hops, as the goal is to minimize the weights of the path, not the hopping length. Realistically, this is not the most ideal, as longer path lengths carry higher risks of intermediate nodes defaulting on the transaction, meaning higher risk of transaction failure. We would like a compromise between minimal fees and low number of hops, and optimizing the fee structure to achieve this is the main goal of this study.

To study the influence of the fee structure, we further let the weight of each edge be $1/x^\alpha$, where α is a tunable parameter and $\alpha > 0$ [18]. As an example, in Fig. 2(a), after Alice has sent the 0.2 BTC to Bob, the fee for using the channel in the A-to-B direction is $1/0.3^\alpha$, while the fee for using the B-to-A direction is $1/0.7^\alpha$. We can see that regardless of α, the fee is cheaper to send a transaction in the B-to-A direction, as this helps to rebalance the channel. Increasing α has the effect of favoring more hops but lower fees, while decreasing α favors less hops but higher fees. We will determine the optimal value of α that gives the best compromise between shortest hops and more balanced channels.

IV. Simulation Results

In our simulations, we sample 50000 transactions from the BN, using the statistics obtained from Fig. 1(b) after normalizing and Fig. 5. Every simulation step involves running 1 transaction. Dijkstra's algorithm [17] is used to find the weighted shortest path between the source node and target node, with the condition that every channel in the path has sufficient capacity to route the transaction. If no such path exists, the transaction fails. The results for each α are averaged over 3 simulation runs.

Fig. 6(a) shows how the mean I_{cap}, denoted $\langle I_{cap} \rangle$, over all channels varies with simulation steps. We run the simulations for different fee structures, including a scenario where routing incurs no fees for users. Unsurprisingly, without a fee incentive, the channel imbalance is the largest. For all fee structures $1/x^\alpha$, a noticeable improvement in $\langle I_{cap} \rangle$ is observed. Larger α favors the usage of larger capacity channels, which are more resistant to imbalance. What is interesting is that beyond

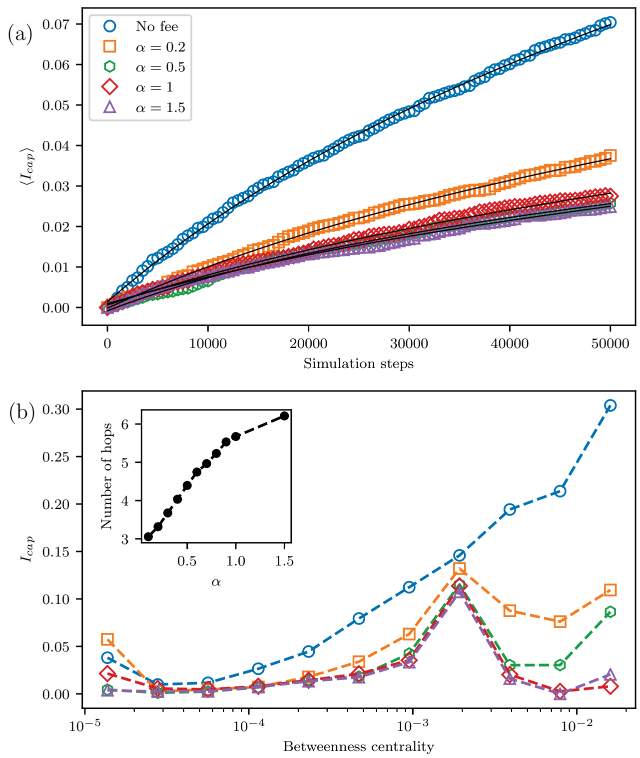

Fig. 6. (a) Plot of the mean channel imbalance $\langle I_{cap} \rangle$ over all channels in the LN against the simulation steps (number of transactions) for different fee structures determined by α. The solid black lines are fitted lines with log functions. As α increases, the LN is more balanced, and such a trend saturates at $\alpha = 0.5$. (b) Plot of I_{cap} against the channel betweenness centrality at the end of the simulation. Data points are log binned. The decrease in imbalances with increasing α similar to (a) is also observed for different centralities. Inset: Plot of the average number of hops in the simulation against the parameter α. As α increases, the average number of hops transactions take increases.

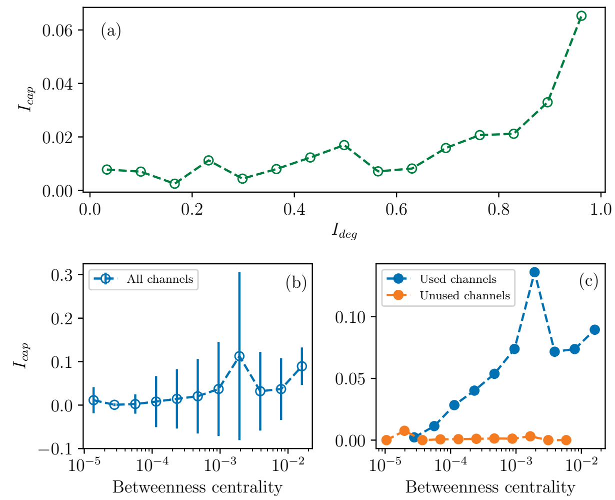

Fig. 7. (a) Plot of the channel imbalance I_{cap} against the channel's node imbalance I_{deg} at the end of the simulation for $\alpha = 0.5$. Data points are binned. We see that channels with a large imbalance in the end nodes have the largest imbalance in capacities. (b) Plot of I_{cap} against BC for $\alpha = 0.5$ with standard deviation bars. The peak has the largest standard deviation, indicating signs of a potential phase transition. (c) The plot of I_{cap} against BC for two types of channels for $\alpha = 0.5$: used channels (where there are transactions in the last 10000 simulation steps) and unused channels (where there are no transactions in the last 10000 simulation steps). Therefore, (b) is the average of the two plots in (c).

$\alpha = 0.5$, no noticeable improvements in $\langle I_{cap} \rangle$ is observed as the curves become almost indistinguishable. The curves fit to a log curve (solid lines), implying that it will take an exponentially long time for the LN to reach a state of total imbalance. In addition, larger α values leads to longer time of reaching imbalance, and better sustainability of the LN without replenishing new channels.

The inset of Fig. 6 plots the average hops in the transactions against α. In agreement with Ref. [18], a larger α favors longer path lengths. We now face the trade-off between two factors: larger α minimizes fees on the LN and minimizes $\langle I_{cap} \rangle$, but increases the routing path lengths. In light of the results of Fig. 6, we find that $\alpha = 0.5$ is a good compromise. Increasing beyond 0.5 yields no significant decrease in $\langle I_{cap} \rangle$, while going below 0.5 would yield significant increase in $\langle I_{cap} \rangle$. From our simulations, transactions hop on average 4.5 intermediate channels before reaching their target node when $\alpha = 0.5$.

As a final investigation, we study the heterogeneity of the imbalances for different types of links, in particular links of different shortest path betweenness centrality (BC) measures [19]–[22]. The shortest path BC ranks channels based on how many shortest paths of all pairs of nodes the channel lies on, which is a very natural definition for the LN due to its routing mechanism.

Fig. 6(b) shows imbalances for channels of different BC at the end of the simulation run. Echoing the results of Fig. 6(a), we find that no-fee structure results in the largest imbalance across channels of all BC. Increasing α decreases the imbalance, again with $\alpha = 0.5$ being the saturation point and any increase does not result in noticeable improvements. There is a general increasing trend in I_{cap} when BC increases to $\sim 10^{-3}$, which is a result of channels with higher BC being used more in routing transactions. This is followed by a peak in I_{cap} (for fee cases) at $\sim 10^{-3}$, before a decrease in I_{cap}.

It is unexpected that the nodes with BC values $\sim 10^{-3}$ have highest imbalances. In practice, it means it is the least favorable to open channels of intermediate BC values. This further acts as a barrier against preferential attachment, preventing small nodes from gaining increasing BC, possibly preventing centralization of the LN. In addition, at $\alpha = 0.5$, the channels with the highest BC values have relatively higher imbalances. Hence, these channels provide less favorable fees compared to those of lower BC values. This again acts as a penalty against hubs, possibly decreasing the extent of centralization.

To understand the general increasing trend between channel imbalance I_{cap} and BC of links in Fig. 6(b), we look at the

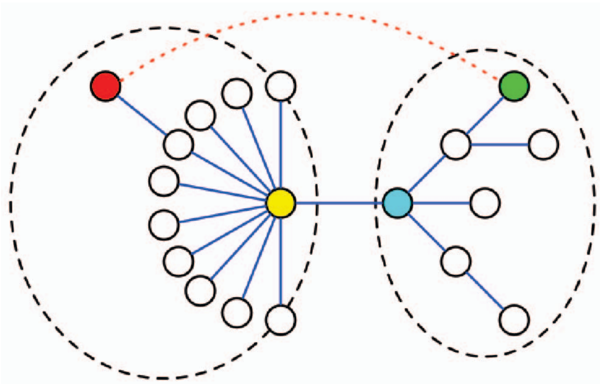

Fig. 8. An illustration of a hub-non-hub concept in a network. The channel with the yellow and cyan node is of high BC and I_{deg}. The link becomes easily imbalanced as transactions are biased from the left to the right direction.

relation between I_{cap} and degree imbalance I_{deg} defined as

$$I_{deg} = \frac{|k_1 - k_2|}{k_1 + k_2},$$

where k_1 and k_2 are the degrees of the two nodes at the ends of a link. As shown in Fig. 7(a), I_{cap} increases drastically with links of high I_{deg}. Due to high degree imbalance, there is an overall imbalance of transaction directions over such links, resulting in high channel imbalance. This effect is more prominent for links with high BC, since they are connected with hubs on one end with more neighbors than the node on the other end, as illustrated in Fig. 8. For the optimal fee structure $\alpha = 0.5$, we see from Fig. 7(b) that the links with largest I_{cap} also have the largest standard deviation among its channel imbalance values, suggestive of possible critical phase transition at the peak position. From Fig. 7(c) we see that I_{cap} for the used channels also peaks at the same position as all the channels, signifying that the peak is not a simple smooth cross over or local maximum, but potentially a critical phase transition phenomenon in physics, which could be interesting for further research.

V. CONCLUSIONS

In this paper, we have provided a first analysis of the network properties of the Lightning Network, and found that like its parent the Bitcoin network, it exhibits scale-free behavior with a power-law degree distribution. We proposed a simple fee model for the LN and ran simulations with transactions sampled from real-world Bitcoin transactions. We found that three important characteristics of the LN, namely fees, channel imbalances and routing path lengths, compete with each other and it is impossible to optimize all three simultaneously. Through simulations, we found a compromise at $\alpha \sim 0.5$ for the simple fee model $\propto 1/x^{\alpha}$, which provides the best trade-off of the three characteristics. It is worth noting that simplifications have been made in this study, in particular on the routing protocols. In practice the routing protocol may not try to find the lowest transaction fee possible, but we assume it may not deviate too far from it. Future studies can be focused on more realistic routing protocols as well as more optimal network structures, in conjunction with routing fee structures to construct more optimal Lightning Networks.

ACKNOWLEDGMENT

The authors would like to thank Dániel Kondor for the help in processing the raw Bitcoin data from the ELTE project. The computational work for this paper was done on resources of the National Supercomputing Computer, Singapore.

REFERENCES

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
[2] 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.
[3] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 17–30.
[4] J. Poon and T. Dryja, "The bitcoin lightning network: Scalable off-chain instant payments," *draft version 0.5*, vol. 9, p. 14, 2016.
[5] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, no. 1, p. 47, 2002.
[6] R. Albert, H. Jeong, and A.-L. Barabási, "Internet: Diameter of the world-wide web," *Nature*, vol. 401, no. 6749, p. 130, 1999.
[7] R. Albert, H. Jeong, and A.-L. Barabási, "Error and attack tolerance of complex networks," *nature*, vol. 406, no. 6794, p. 378, 2000.
[8] ELTE bitcoin project. http://www.vo.elte.hu/bitcoin/.
[9] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
[10] P. Prihodko, S. Zhigulin, M. Sahno, A. Ostrovskiy, and O. Osuntokun, "Flare: An approach to routing in lightning network," 2016.
[11] C. Grunspan and R. Pérez-Marco, "Ant routing algorithm for the lightning network," *arXiv preprint arXiv:1807.00151*, 2018.
[12] Lightning network visualizer. https://graph.lndexplorer.com/.
[13] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on*. IEEE, 2013, pp. 1–10.
[14] D. Kondor, M. Pósfai, I. Csabai, and G. Vattay, "Do the rich get richer? an empirical analysis of the bitcoin transaction network," *PlOS One*, vol. 9, no. 2, p. e86197, 2014.
[15] M. Lischke and B. Fabian, "Analyzing the bitcoin network: The first four years," *Future Internet*, vol. 8, no. 1, p. 7, 2016.
[16] M. Ober, S. Katzenbeisser, and K. Hamacher, "Structure and anonymity of the bitcoin transaction graph," *Future Internet*, vol. 5, no. 2, pp. 237–250, 2013.
[17] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
[18] T. Opsahl, F. Agneessens, and J. Skvoretz, "Node centrality in weighted networks: Generalizing degree and shortest paths," *Social Networks*, vol. 32, no. 3, pp. 245–251, 2010.
[19] M. E. Newman, "Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality," *Physical Review E*, vol. 64, no. 1, p. 016132, 2001.
[20] L. C. Freeman, "A set of measures of centrality based on betweenness," *Sociometry*, pp. 35–41, 1977.
[21] L. C. Freeman, "Centrality in social networks conceptual clarification," *Social Networks*, vol. 1, no. 3, pp. 215–239, 1978.
[22] U. Brandes, "A faster algorithm for betweenness centrality," *Journal of Mathematical Sociology*, vol. 25, no. 2, pp. 163–177, 2001.