

Boosting Payment Channel Network Liquidity with Topology Optimization and Transaction Selection

Krishnendu Chatterjee ✉

ISTA, Austria

Jan Matyáš Kříšťan ✉

Faculty of Information Technology, Czech Technical University, Czechia

Stefan Schmid ✉

TU Berlin & Weizenbaum Institute, Germany

Jakub Svoboda ✉

ISTA, Austria

Michelle Yeo¹ ✉

National University of Singapore, Singapore

Abstract

Payment channel networks (PCNs) are a promising technology that alleviates blockchain scalability by shifting the transaction load from the blockchain to the PCN. Nevertheless, the network topology has to be carefully designed to maximise the transaction throughput in PCNs. Additionally, users in PCNs also have to make optimal decisions on which transactions to forward and which to reject to prolong the lifetime of their channels. In this work, we consider an input sequence of transactions over p parties. Each transaction consists of a transaction size, source, and target, and can be either accepted or rejected (entailing a cost). The goal is to design a PCN topology among the p cooperating parties, along with the channel capacities, and then output a decision for each transaction in the sequence to minimise the cost of creating and augmenting channels, as well as the cost of rejecting transactions. Our main contribution is an $\mathcal{O}(p)$ approximation algorithm for the problem with p parties. We further show that with some assumptions on the distribution of transactions, we can reduce the approximation ratio to $\mathcal{O}(\sqrt{p})$. We complement our theoretical analysis with an empirical study of our assumptions and approach in the context of the Lightning Network.

2012 ACM Subject Classification Networks → Network algorithms; Theory of computation → Graph algorithms analysis; Theory of computation → Approximation algorithms analysis

Keywords and phrases Blockchains, Cryptocurrencies, Payment Channel Networks, Throughput, Optimisation, Graph Algorithms, Approximation Algorithms

Digital Object Identifier 10.4230/LIPIcs.DISC.2025.4

Funding *Krishnendu Chatterjee*: European Research Council CoG 863818 (ForM-SMArt) and Austrian Science Fund 10.55776/COE12

Jan Matyáš Kříšťan: Czech Science Foundation Grant no. 24-12046S

Stefan Schmid: German Research Foundation (DFG) project ReNO (SPP 2378) from 2023-2027

Jakub Svoboda: European Research Council CoG 863818 (ForM-SMArt) and Austrian Science Fund 10.55776/COE12

Michelle Yeo: MOE-T2EP20122-0014 (Data-Driven Distributed Algorithms)

1 Introduction

Blockchain scalability is one of the key bottlenecks in the mainstream adoption of cryptocurrencies [8, 11, 23]. For instance, Bitcoin can only process an average of 7 transactions per

¹ Corresponding author

second which is paltry compared to Visa’s 47,000 [25]. This makes it impractical as of yet for the widespread usage of cryptocurrencies in everyday situations.

Some promising solutions to blockchain scalability are layer 2 solutions like payment channel networks (PCNs) [12, 13, 24, 25], which allow users to bypass the blockchain and transact directly with each other. To do so, two users have to first open a payment channel between themselves in an initial funding transaction on the blockchain, thereby locking some funds only to be used in the channel. Thereafter, the two users transact with each other simply by updating the distribution of funds on each end of the channel to reflect the payment amount. As compared to the tedious process of waiting for consensus on the blockchain, these payments can be finalised almost instantaneously since they only involve exchanging signatures between the two parties. As such, with only a constant number of blockchain transactions, any two users can make an unbounded number of off-chain transactions in the payment channel, consequently increasing the overall transaction throughput of the blockchain.

A PCN is a network of payment channels and users, where two users that are not directly connected to each other with a payment channel can still transact with each other in a multi-hop fashion over a path of payment channels. To incentivize intermediary nodes on payment paths to forward payments, PCNs allow these intermediary nodes to charge a small transaction forwarding fee that is typically linear in the transaction amount. Two of the most notable examples of PCNs are Bitcoin’s Lightning Network [25] and Ethereum’s Raiden [2].

Apart from joining a PCN in order to transact efficiently with other users, a secondary consideration for users in PCNs is to whom they should establish channels and at which capacities, as well as which transactions then to select for forwarding. Aside from the opportunity cost of rejecting to forward transactions and thus missing out on the revenue from transaction fees, users in PCNs also incur a cost whenever channels are created (the on-chain transaction fee) as well as when additional capacity is injected into the channel. As such, although it is tempting to accept to forward all transactions and thus profit from the transaction fees, doing so would quickly deplete the nodes’ capacity on their incident channels. A crucial optimisation problem that intermediary nodes in transaction paths would therefore have to consider is deciding which transactions to accept and which to reject, given a minimal injection of capacity into their channel.

In this work, we study the problem of both optimal topology design and transaction selection over a PCN. The input to the problem is a sequence of transactions and a set of p coordinating parties. Each transaction in the sequence is of a certain amount, and is associated with a source and target from the set of parties. The problem then proceeds in two stages: in the first stage, the goal is for the p parties to decide jointly on a PCN topology over themselves. Each created channel in this stage incurs both a fixed channel creation cost as well as a capacity cost that depends on the amount of funds injected into the channel. The second stage involves processing the transactions in the transaction sequence, either accepting or rejecting them (and in so doing incurring a rejection fee) along the created paths in the network. The goal is to create a network and inject enough capacity into each channel such as to maximise the number of transactions accepted while incurring the least amount of cost (see the cost model in Section 2 for more details).

1.1 Related work

To the best of our knowledge, our work is the first to consider the problem of designing an optimal PCN topology for optimising the problem of transaction selection in PCNs. Our work is closely related to the work of [27] which considers a similar problem over a single

payment channel and shows that the problem of deciding which transactions to accept or reject over a single channel is already NP-hard. The main open question of [27] is how to extend their proposed algorithm from a single channel to an entire network. Our work addresses this problem by proposing a solution for network design as well as algorithms for optimally accepting or rejecting transactions along paths in a network. Our work is also related to the works of [4, 7] which consider a similar problem of transaction selection but in the online setting and also limited to a single channel. Additionally, the works of [15, 16] also study minimal-cost PCN topologies, but do not consider the complementary problem of transaction selection.

Another related line of research is network creation games on cryptocurrency networks [5, 6, 14]. These works study how a node should connect to an existing cryptocurrency network in a stable and optimal way, optimising the utility of the single node in a unilateral fashion. In contrast, our work proposes a method of creating a network among any number of nodes such as to *jointly minimise* the cost (over all nodes) of network creation and processing transactions.

A crucial assumption in our model is coordination between the interested parties in creating the network during the network creation stage. We stress that this is not an impractical assumption and several works have made similar assumptions [5, 6, 7] where optimal behaviour is analysed in the setting where some or all involved parties cooperate and create certain network topologies together.

Further upfield but also related are works analysing the consensus number of a cryptocurrency [19, 21]. The transactions studied in our work, and in general the sequential nature of layer 2 transactions on PCNs, have consensus number 1, implying that expensive consensus protocols are unnecessary in such settings. Additionally, our work is also related to the classic works on theory of contention management in transactional memory [3, 17, 18].

More generally, our work is related to optimising flows and throughput in typical capacitated communication networks [9, 26]. However, we stress a crucial difference between our problem and problem over traditional communication networks: in traditional communication networks, the capacity is usually independent in the two directions of the channel [20]. In our case, the amount of transactions u sends to v in a channel (u, v) directly affects v 's capability to send transactions, as each transaction u sends to v increases v 's capacity on (u, v) .

1.2 Main challenges and our contributions

We generalise the model of processing transactions over a single channel proposed in [27] to an entire network. The first challenge in shifting from optimising processing transactions from a single channel to a network is deciding on the optimal network topology among the cooperating parties. Here, we show in Theorem 2 in Section 3.2 that the star graph of size p is a 2-approximation of the problem of creating an optimal PCN over p parties.

Once the optimal topology is decided, the next challenge would be to decide on which transactions to accept to forward, given that transactions have to be accepted along the entire payment path. In this setting, we leverage the star topology to generalise the existing single channel algorithm in [27] to get an algorithm for transaction selection optimisation in the star. Our main algorithm in this setting is described in Section 4 and we prove that it is an $\mathcal{O}(p)$ approximation of the weighted transaction selection problem in general graphs.

We describe another algorithm in Section 5 that has an improved approximation ratio of $\mathcal{O}(\sqrt{p})$ (see Theorem 10) for the transaction selection problem in general graphs. To achieve this approximation ratio, we adopt an additional realistic and reasonable assumption on the distribution of transactions in the transaction sequence. Namely, we assume that


136 transactions are clustered, following a stochastic block model.

137 Finally, we augment our theoretical analysis with a case study of the Lightning Network
 138 in Section 6. Using simulations, we find that the network statistics correspond to and thus
 139 justify our transaction distribution assumptions.

140 2 Model

141 *Payment channels.* A unique aspect of our problem is that we optimise the selection of
 142 weighted transactions over payment channels. Payment channels can be formally modeled as
 143 edges that satisfy the following three properties:

- 144 1. The capacity or weight of a channel (u, v) is the sum of the initial capacities b_u and b_v
 145 injected into the channel by users (or nodes) u and v . The capacities b_u and b_v correspond
 146 to the initial amount of funds that users u and v choose to inject into the channel during
 147 channel creation.
- 148 2. Given a capacitated channel (u, v) with weight w , the weight or capacity can be arbitrarily
 149 split between both ends of the channel depending on the number and weight of transactions
 150 processed by u and v . That is, b_u and b_v can be arbitrary as long as $b_u + b_v = w$ and
 151 $b_u, b_v \geq 0$. For instance, given a payment channel (u, v) with an initial capacity split
 152 of b_u and b_v , if u sends a payment of amount x to v , the capacity on u 's end of the
 153 channel drops to $b_u - x$ and the capacity of v 's end increases to $b_v + x$ after processing
 154 the transaction. See Figure 1 for more examples of how the capacities on each end of a
 155 payment channel can vary in the course of processing transactions.
- 156 3. The total capacity of a payment channel is invariant throughout the lifetime of the channel.
 157 That is, it is impossible for nodes to add to or remove any part of the capacity in the
 158 channel. In particular, if a node is incident to more than one channel in the network, the
 159 node cannot transfer part of its capacity in one channel to “top up” the capacity in the
 160 other one.



■ **Figure 1** The diagram on the top shows the outcome of processing a transaction of size 2 going from A to D along the path (A, B, C, D) in the graph. The diagram on the bottom shows the outcome of processing a transaction of size 5 from A to D . The transaction has to be rejected as the size of the transaction is larger than the sum of the capacities of C in the channel (C, D) and either A or E in the channel (A, F) or (E, D) respectively.

161 *Payment channel network.* We model a PCN as a weighted, undirected graph $G = (V, E)$
 162 where the vertices or nodes of G are users in the PCN and the edges of G are payment

channels between nodes. We denote the capacity of a channel $(u, v) \in E$ as $w_{u,v}$, and the capacity of nodes u and v on the channel (u, v) as $b_{u,v}$ and $b_{v,u}$ respectively.

Transaction sequence. Let V denote a set of nodes (parties). We denote a finite ordered sequence of transactions over V by $X_n = ((s_1, t_1, x_1), \dots, (s_n, t_n, x_n))$, where $s_i, t_i \in V$ represents the sender and target of the i th transaction respectively and $x_i \in \mathbb{R}^+$ represents the weight or size of the i th transaction.

Processing transactions in PCNs. Let $G = (V, E)$ be a connected PCN with some capacity over all of its channels. Let X_n be a sequence of transactions over V . Consider the i th element of the sequence (s_i, t_i, x_i) and suppose the transaction travels over some path $\pi_i = (s_i, \dots, t_i)$ from the i th sender to the i th target. Let us call the first node along each channel in the path π_i the *forwarding node*. Each forwarding node in the path π_i can choose to do the following to the transaction:

- **Accept transaction.** For a forwarding node u in a channel $(u, v) \in \pi_i$, node u can accept to forward the transaction if their capacity in (u, v) is at least the size of the transaction.
- **Reject transaction.** Node u can also reject the transaction. This could happen if u 's capacity is insufficient, or if accepting the transaction would incur a larger cost in the future. For a transaction of weight x , the cost of rejecting the transaction is $f \cdot x + m$ where $f, m \in \mathbb{R}^+$. Apart from incurring some cost, rejecting a transaction does not alter the capacity distribution on both ends of the channel (see the diagram on the bottom of Figure 1 for an example of rejecting a transaction).

We note that node v does not need to take any action nor incur any cost when transactions are going in the direction of u to v along a channel (u, v) . Finally, we stress that decisions on each channel in a transaction path need to be the same.

Problem definition. Suppose we have p coordinating parties that want to come together to create a PCN that allows the processing of transactions between themselves. The input to the problem is a transaction sequence over the p coordinating parties. The problem then proceeds in 2 stages. The first stage of the problem is the *network creation stage*, where the p parties have to firstly decide on a PCN topology over themselves. Once the PCN has been created, the second stage of the problem is *transaction selection*, where the network topology created in stage 1 determines the path the transactions travel over. In this stage of the problem, the parties need to decide whether to accept or reject each transaction in the sequence.

Hardness. We know from [27] that the transaction selection problem over a single channel is NP-hard. Our problem includes $p > 2$ parties and we also need to find the optimal network topology over these p parties (for two parties, it is only a channel). This means that our problem is at least as hard as the problem on a single channel, and hence NP-hard as well.

Objective and costs. We consider three types of costs faced by the p parties:

1. **Channel creation costs.** We assume creating each channel incurs a fixed auxiliary cost of $k > 0$.
2. **Capacity costs.** For each channel (u, v) created between parties u and v , the capacity cost of (u, v) is the sum of the capacities $b_{u,v}, b_{v,u}$ injected by u, v into the channel.
3. **Rejection costs.** These are costs incurred during transaction selection when a party rejects a transaction.

The *total cost* incurred by the parties is the sum (over all channels in the PCN and transaction decisions) of channel creation, capacity, and rejection costs. The objective of the p parties is to create a PCN over themselves and output the decisions for each transaction in the sequence such that the total cost is minimised.

3 Accepting all transactions

We begin our investigation of transaction selection over general graphs by first considering the problem a restricted setting where all transactions in the input transaction sequence need to be accepted. In this setting, the total cost consists only of channel creation and capacity costs. We also assume here that we know the optimal PCN topology among the p parties, which we denote by $G = (V, E)$.

3.1 Linear program for a set of channels

We present a linear program (LP) that, for a given set of channels, computes the optimal capacity that the parties need to lock into the channels of the network in order to accept all transactions in the sequence. Indeed, if the channel creation cost k is small or if we want to know a lower bound for our problem, we can assume the optimal graph G is complete. For a node $v \in G$, let $Ne(v)$ denote its neighbours. Let S_v denote the *total balance* of node v in the graph, which is simply the sum of its balances on all its incident channels. For a channel $(u, v) \in E$, we denote the balance of node v (resp. u) after processing the i th transaction in the sequence by $b_{v,u}^i$ (resp. $b_{u,v}^i$), and we also denote the total balance of v after processing the transaction by S_v^i . Let $AC(\mathcal{S}, G)$ denote the optimal capacity cost for the graph G where every transaction in X_t is accepted. That is, the total capacity injected into the channels of the graph G .

Before describing the LP, we define a *difference function* $\delta_v(s, t)$ between a source s and target t in a transaction path:

$$\delta_v(s, t) = \begin{cases} 1 & v = t \text{ and } v \neq s \\ -1 & v = s \text{ and } v \neq t \\ 0 & \text{otherwise} \end{cases}$$

This function will be used in the LP to maintain the invariant that processing a transaction can only affect the total balances of the source and target nodes in the transaction path, but the total balance of every other node in the graph must not change.

Given G and a transaction sequence $\mathcal{S} := ((\pi_1, x_1), \dots, (\pi_n, x_n))$ over G , we construct the following LP.

$$\begin{aligned} \text{minimise} \quad & \sum_{(v,u) \in E} w_{v,u} \\ \forall v, u, i : \quad & w_{v,u} = b_{v,u}^i + b_{u,v}^i \\ \forall v, i : \quad & S_v^i = \sum_{u \in Ne(v)} b_{v,u}^i \\ \forall v, i, (s_i, t_i, x_i) : \quad & S_v^i - \delta_v(s_i, t_i)x_i = S_v^{i-1} \end{aligned}$$

For the sequence of transactions \mathcal{S} and graph G , we denote the solution of the LP with the respective parameters by $LP(\mathcal{S}, G)$. The following lemma (with proof in Appendix A.1)

proves that the solution of the LP is a lower bound on the capacity cost of the problem in the setting where all transactions are to be accepted.

► **Lemma 1.** *Given the sequence of transactions \mathcal{S} , we have $LP(\mathcal{S}, K_p) \leq AC(\mathcal{S}, G)$.*

3.2 Star as a 2-approximation

Although the capacity cost for the setting where all transactions are accepted is lowest in the complete graph K_p , the complete graph has $\Theta(p^2)$ edges. Thus, depending on the auxiliary channel creation cost k , the cost of creating K_p as well as injecting sufficient capacity into all channels of K_p might be large.

As a useful heuristic, we observe from the proof of Lemma 1 that the capacity cost is related to the diameter of the graph: the smaller the diameter, the lower the cost. Thus, a useful class of graphs to look at for the network creation stage of the problem are graphs that have both a small diameter and a small number of channels. An example of a graph that simultaneously satisfies both of these requirements is a star. In a star, one central node is connected by a channel to all other nodes. Thus, the star has a small diameter of 2, and $p - 1$ channels, which is minimal for a connected graph over p nodes. We denote the star on p nodes as $K_{1,p-1}$.

We first observe that for the star, it is easy to compute a lower bound on the amount of capacity required in the setting where all transactions are accepted. This is due to the fact that between any two nodes u, v in a star, there is only one unique path channeling u and v , rendering transactions unsplittable. The capacity on the channel between any node v and the centre of the star c , $w_{v,c}$, has to therefore be larger than $\max_{j,k \leq n} |\sum_{s_i, t_i, x_i | k \leq i \leq j} \delta_v(s_i, t_i) \cdot x_i|$, as this quantity captures the largest difference in the balance between both users of the channel.

We now show in Theorem 2 that we can upper bound the optimal capacity cost on the star. The proof of Theorem 2 is in Appendix A.2.

► **Theorem 2.** *For every graph G over p nodes, and any graph G' with diameter d , the following inequality holds*

$$AC(\mathcal{S}, G') \leq d \cdot AC(\mathcal{S}, G).$$

► **Corollary 3.** *For every graph G over p nodes, and graph $K_{1,p-1}$, the following inequality holds*

$$AC(\mathcal{S}, K_{1,p-1}) \leq 2 \cdot AC(\mathcal{S}, G).$$

With these insights, we restrict the network created among the p parties during the network creation stage of the weighted transaction problem to a star in Section 4, and a star of stars (of diameter 4, see Figure 3 for an example) in Section 5.

4 A general setting approximation algorithm

We now present an approximation algorithm for the general setting where transactions can be either accepted or rejected. We first reiterate that from Theorem 2 we can restrict ourselves to the star graph at the network creation stage of the problem, and in doing so incur twice as much capacity cost as the optimal capacity cost. Given that the channel creation cost for the star topology is fixed at $k(p - 1)$, the focus of our algorithm in this section is to come up with capacities that should be injected on both ends of each channel as well as a sequence of transaction decisions so as to minimise the capacity and rejection costs.

4.1 LP for a star

We describe a LP that minimises the capacity cost as well as rejection cost over transactions going between any two leaf nodes in the star. This LP is a relaxation of the general problem, where we allow some transactions to be accepted fractionally. That is, for a transaction of size x_i , the LP does not have to accept the full transaction amount, but can accept some portion y_i such that $0 \leq y_i \leq x_i$. Let us call transactions for which $y_i = x_i$ in the LP solution *fully accepted* transactions.

Let c denote the central node of the star. For a leaf node v and $i \in \mathbb{N}$, we use $C_{v,i}$ and $C'_{v,i}$ to denote the capacity at the channel (v, c) after processing the i -th transaction on the side of v and c respectively. Let C_v denote the total capacity of the channel (v, c) . We can now formulate the LP as follows:

$$\text{minimise } \sum_u C_u + \sum_i f(x_i - y_i) + \frac{m(x_i - y_i)}{x_i} \quad (1)$$

$$\forall i : 0 \leq y_i \leq x_i$$

$$\forall (v, u, x_i) : C_v, C_u \geq x_i$$

$$C_{v,i} = C_{v,i-1} - y_i$$

$$C'_{v,i} = C'_{v,i-1} + y_i$$

$$C'_{u,i} = C'_{u,i-1} - y_i$$

$$C_{u,i} = C_{u,i-1} + y_i$$

► **Lemma 4.** *The total cost of the star LP solution is a lower bound on the total cost of the optimal solution.*

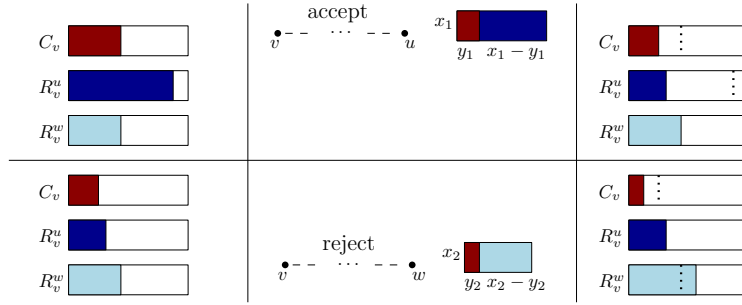
Proof. The proof follows from the observation that the optimal integer solution is also a feasible solution for the LP. ◀

► **Remark 5.** The solution of the LP assumes that the capacity cost and the transaction rejection cost are equally large, i.e., we cannot aggressively optimise for one objective while neglecting the other. An implication of this which motivates the design of our algorithms in the following sections is that our algorithms should infrequently reject fully accepted transactions, as these are transactions that do not incur any rejection cost and also change the distribution of the capacity along channels in a more optimal way.

4.2 An $\mathcal{O}(p)$ approximation algorithm

Algorithm for a single channel. We summarize an algorithm from [27] that solves the problem for a single channel, which we will call the *channel algorithm*. The channel algorithm first approximates the capacity to be injected into the channel. Next, the channel algorithm solves a LP (similar to Equation (1)). Given the solution of the LP, we denote by C_u and C_v the minimal capacity needed on each end of a channel (u, v) to mimic the (fractional) transport of transactions as dictated by the LP. In addition to C_u and C_v , the channel algorithm allocates additional capacities on each end of the channel, which we call *reserves*, and denote them by R_u and R_v . These reserves are mainly used to ensure that all transactions that are fully accepted by the LP are also accepted by the channel algorithm (a consequence of Remark 5).

326 *Extending the channel algorithm to a star.* We describe our algorithm that extends the
 327 channel algorithm over 2 nodes to a star $K_{1,p-1}$ with central node c in Algorithm 1. The
 328 key idea behind our algorithm is to first use the LP as described in Equation (1) to fix
 329 initial capacities on each end of the channel for every channel in the network (for instance
 330 the capacity C_v for v in (v, c) in Figure 2). We also provide each leaf with $p - 2$ additional
 331 reserves (to process transactions for each other leaf). For each leaf pair u and v , let M denote
 332 the smallest change in the capacity as specified by the output of the LP after all interactions
 333 between u and v . The additional reserve R_u^v of u for processing transactions of v would then
 334 be of size $\frac{\sqrt{3}}{2}M$ (similarly for v). We leave details of Algorithm 1 to Appendix B.1.



■ **Figure 2** Movement of capacities and reserves from the perspective of node v in a star with centre node c and 2 other leaf nodes u and w . C_v represents the capacity as specified by the output of the LP. R_v^u and R_v^w denote the reserves stored by v to process transactions from u and w respectively. The rightmost column represents the capacities and reserves of v after processing the transaction.

335 The next theorem shows that Algorithm 1 is an $\mathcal{O}(p)$ approximation of the optimal cost
 336 on a star. The proof of Theorem 6 is in Appendix A.3

337 ► **Theorem 6.** *Given the transaction sequence X_t , let \mathcal{C} be the optimal capacity cost and \mathcal{R}*
 338 *the optimal rejection cost on a star. Then Algorithm 1 run for every pair of nodes incurs a*
 339 *rejection cost of at most $(\sqrt{3} + 1)\mathcal{R}$ and capacity cost of at most $(1 + (p - 1) \cdot \sqrt{3})\mathcal{C}$.*

340 In the next section, we show that if we impose some reasonable assumptions on the
 341 distribution of the input transaction sequence in our model, we can design an algorithm with
 342 a significantly lower approximation ratio of $\mathcal{O}(\sqrt{p})$.

343 5 An improved algorithm using the stochastic block model

344 *Stochastic block model.* In reality, users typically interact with each other in a structured
 345 manner, resulting in spatial and temporal locality. For example, if two nodes that previously
 346 transacted with each other are more likely to transact again in the near future. Our model for
 347 how nodes interact and send transactions to each other is the stochastic block model [22], in
 348 particular, the planted partition model [10]. The stochastic block model is a generative model
 349 to create random graphs where nodes are partitioned into clusters (blocks). See Appendix C.1
 350 for more details on the stochastic block model.

351 *Clustering.* We now define the way nodes are partitioned into clusters given an input
 352 sequence of transactions. Given an input transaction sequence over a network, we say that the
 353 nodes are (m, k, t) -clustered if there exists a partition of nodes to m clusters C_1, C_2, \dots, C_m
 354 where each cluster has size at most k and the clusters satisfy the following *clustering*
 355 *conditions*:

■ **Algorithm 1** Modified channel algorithm

Input: nodes u and v , transaction sequence X_t , solution of LP indexed by $i: C_u, C'_u, C_v, C'_v$.

Output: decisions to accept or reject for transactions between u and v

```

1:  $M \leftarrow \min(C_{u,0} + C'_{u,0}, C_{v,0} + C'_{v,0})$ 
2:  $R_u = \frac{\sqrt{3}}{2}M, R_v = \frac{\sqrt{3}}{2}M$ 
3:  $H \leftarrow \text{CREATEHEAP}()$ 
4: for  $i \in [t]$  do
5:   if  $x_i$  is transaction from  $u$  to  $v$  then
6:     if  $R_u - (x_i - y_i) \geq \frac{\sqrt{3}-1}{2}M$  then
7:        $x_i \leftarrow \text{Accept}$ 
8:        $R_u, R_v = R_u - (x_i - y_i), R_v + (x_i - y_i)$ 
9:     else if  $\frac{y_i}{x_i} < \frac{\sqrt{3}}{\sqrt{3}+1}$  then
10:       $x_i \leftarrow \text{Reject}$ 
11:       $R_u, R_v = R_u + y_i, R_v - y_i$ 
12:   else
13:      $x_i \leftarrow \text{Accept}$ 
14:      $R_u, R_v = R_u - (x_i - y_i), R_v + (x_i - y_i)$ 
15:      $\text{ADDTOHEAP}(H, x_i)$ 
16:     if  $R_u < 0$  then
17:       while  $R_u < \frac{\sqrt{3}-1}{2}M$  do
18:          $x_j \leftarrow \text{POPHEAP}(H)$ 
19:          $x_j \leftarrow \text{Reject}$ 
20:          $R_u, R_v = R_u + x_j, R_v - x_j$ 
21:   if  $R_u \geq \frac{\sqrt{3}-1}{2}M$  then
22:      $\text{EMPTYHEAP}(H)$ 

```

- 356 1. The volume (sum of sizes of transactions) of transactions inside the cluster is t -times the
 357 volume of transactions that go outside the cluster. That is,

$$358 \sum_{(s_j, t_j, x_j) | s_j, t_j \in C_i} x_j \geq t \cdot \sum_{\substack{(s_j, t_j, x_j) | s_j \in C_i \text{ and } t_j \notin C_i \\ \text{or } s_j \notin C_i \text{ and } t_j \in C_i}} x_j$$

- 359 2. For every $S \subset C_i$ such that $|S| \leq \frac{C_i}{2}$, the volume of transactions between S and $C_i \setminus S$ is
 360 at least $\frac{|S|}{4|C_i|}$ the volume in the cluster C_i . That is,

$$361 \frac{|S|}{4|C_i|} \sum_{x_j \text{ inside } C_i} x_j \leq \sum_{x_j \text{ between } C_i \setminus S \text{ and } S} x_j$$


- 362 3. For every $S \subset C_i$, the ratio of volumes of transactions from S and $C_i \setminus S$ is at least $\frac{1}{3}$.
 363 That is,

$$364 \frac{\sum_{x_j: C_i \setminus S \rightarrow S} x_j}{\sum_{x_j: S \rightarrow C_i \setminus S} x_j} \geq \frac{1}{3}.$$

365 These conditions need to be satisfied for every cluster C_i and any contiguous subsequence of
 366 transactions involving C_i (i.e., source or target in C_i) with volume at least $C_i \cdot x_{max}$, where
 367 x_{max} is the weight of the biggest transaction in the input transaction sequence. In the rest
 368 of this section, we assume that given an input sequence of transactions, the nodes can be

369 $(\sqrt{p}, \sqrt{p}, 24\sqrt{p})$ -clustered. Looking ahead, this allows us to design an algorithm that gives us
 370 an $\mathcal{O}(\sqrt{p})$ approximation of the optimal cost. In the general case, for an (m, k, t) -clustering,
 371 for any $t \geq 24\sqrt{p}$, we have an $\mathcal{O}(m + k)$ approximation (by a similar analysis).

372 *Double star.* We define a double star as a graph of diameter $d = 4$ with one central node,
 373 some middle nodes, and some leaf nodes. The central and middle nodes create a star and
 374 the leaf nodes are connected to only one middle node, forming a cluster (see Figure 3 for an
 375 example with 6 clusters). In each cluster, the leaf nodes are connected to the middle node
 376 in a star topology. In our setting, given our aforementioned assumption on $(\sqrt{p}, \sqrt{p}, 24\sqrt{p})$ -
 377 clustering, parties create a double star network topology during the network creation stage,
 378 where each cluster represents a group of nodes of size \sqrt{p} that interact frequently with each
 379 other.



■ **Figure 3** Double star graph with 6 clusters in the dashed boxes. The node labelled c is the centre node, and the middle nodes are the square nodes.

380 5.1 An $\mathcal{O}(\sqrt{p})$ approximation algorithm

381 *Main ideas.* We describe our algorithm that gives an $\mathcal{O}(\sqrt{p})$ approximation of the optimal
 382 cost of the problem in Algorithm 2. On a high level, Algorithm 2 consists of two main ideas.
 383 The first idea is to again formulate and solve an LP on a double star to get the amount of
 384 capacities nodes should inject on their ends of their incident links. This LP is similar to the
 385 LP for a star described in Equation (1) and we detail it in Appendix C.2. From Theorem 2,
 386 we know that the capacity cost of the double star is at worst a 4-approximation of the optimal
 387 capacity cost.

388 The second idea is to split the processing of transactions into processing transactions
 389 *between* clusters and processing transactions *within* clusters. Here, we use the specific double
 390 star topology to ensure that we can apply Algorithm 1 (for processing transactions on a
 391 single star) to process transactions between clusters as well as within each cluster. Finally,
 392 we use our clustering assumptions to allocate enough reserve capacity for each cluster to
 393 ensure that we can always shift the reserve capacities within a cluster from nodes with high
 394 to low reserves. This ensures that each node always has sufficient reserve capacity to align
 395 the decisions on their transactions (whether to accept or reject) as closely as possible to the
 396 ideal decisions as output from the linear program.

397 *Processing transactions between clusters.* After running the linear program for the double
 398 star, let us rewrite each transaction in the sequence as (s, t, x_i, y_i) , where s is the source node,
 399 t is the target node, x_i is the weight of the transaction, and y_i denotes the fractional amount
 400 of the transaction accepted by the solution of the linear program. For processing transactions
 401 between clusters, we look at the star created by the centre and middle nodes of the double
 402 star and treat these transactions as going only between these middle nodes. In Algorithm 2,

■ **Algorithm 2** Approximation algorithm for clustered nodes.

Input: Sequence of transactions X_t with LP solution, clustering $\mathcal{C} = \{C_1, C_2, \dots\}$

Output: Decisions to accept or reject transactions

```

1:  $X_i \leftarrow \text{BETWEENCLUSTERS}(X_t, \mathcal{C})$ 
2:  $S_i \leftarrow \text{SOLVE}(X_i)$ 
3: for  $C \in \mathcal{C}$  do
4:    $X_C \leftarrow \text{BALANCEINCLUSTER}(X_t, S_i, C)$ 
5:    $S_C \leftarrow \text{SOLVE}(X_C)$ 
6: return  $S_i \cup S_b \cup (\bigcup_{C \in \mathcal{C}} X_C)$ 

```

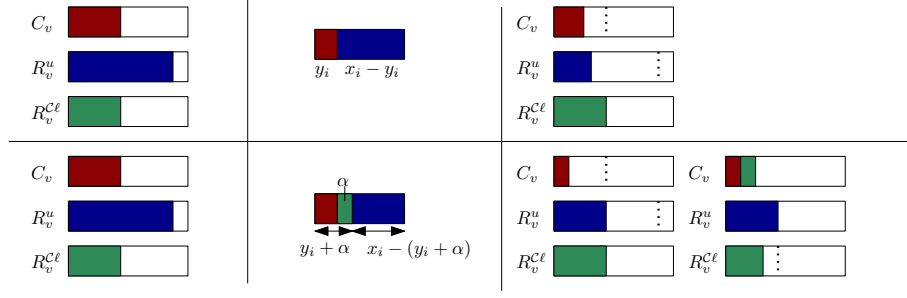
403 this is handled by the function `BETWEENCLUSTERS` which selects only transactions that go
 404 between clusters and treats clusters as a single node. As this becomes processing transactions
 405 on a star, this allows us to use Algorithm 1 to process transactions between different clusters.
 406 In Algorithm 2, this is handled by the function `SOLVE(X)` in line 2 in Algorithm 2 which
 407 outputs the decisions using Algorithm 1 for every pair of nodes. Recall that Algorithm 1
 408 also gives us an upper bound on the *reserve capacity* requirement for each cluster (middle
 409 node): for each middle node, we require reserves of size $\mathcal{O}(\sqrt{p}M)$, where M is the maximum
 410 of the capacity of the channel between the middle node and the centre of the double star c
 411 and the capacities of the channels between the middle node and the leaf nodes in its cluster.
 412 We call these reserves *cluster reserves*.

413 *Processing transactions within clusters.* Before processing transactions that go between
 414 nodes in a cluster, it is imperative that the cluster reserves are first shifted around within
 415 the cluster to ensure each leaf node in the cluster has enough reserves to process their
 416 between-cluster transactions (line 2 in Algorithm 2). This is handled by the function
 417 `BALANCEINCLUSTER` (described in Algorithm 5 in Appendix C.3), which modifies the
 418 solution of the double star LP such that cluster reserves are transferred from well-funded
 419 to depleted nodes. Finally, after balancing the reserve capacities, Algorithm 2 treats each
 420 cluster as an independent star graph and uses Algorithm 1 to process transactions within each
 421 cluster (function `SOLVE` in line 5). Using Algorithm 1 imposes a reserve capacity requirement
 422 of $\mathcal{O}(\sqrt{p}M')$ times the capacity for every leaf node, where M' is the capacity of the channel
 423 between the leaf node and the middle node.

424 *Properties of balancing reserve capacities.* We now prove several important properties of
 425 transporting reserve capacities within clusters as specified by `BALANCEINCLUSTERS`, ensuring
 426 the correctness of `BALANCEINCLUSTERS`. First, Lemma 7 states that we can modify the
 427 solution of the linear program (i.e., the fractional amount of transactions accepted by the
 428 linear program) to move the reserves capacities from the source s to target t , see Figure 4.
 429 From the statement of Lemma 7, we note that the only transactions that do not allow moving
 430 reserves are transactions that are fully accepted ($y_i = x_i$) or rejected ($y_i = 0$). Then, we
 431 quantify how much capacity we can move from one set of nodes to another in Lemma 8.
 432 Lastly, we show in Lemma 9 that we can always transport some reserves between any two
 433 sets of nodes. We defer the proofs of these lemmas to Appendices A.4–A.6.

434 ► **Lemma 7.** *If there is a transaction (s, t, x_i, y_i) , we can change it to*

- 435 ■ $(s, t, x_i, y_i + \alpha)$ to simulate moving α reserves from s to t , and
- 436 ■ $(s, t, x_i, y_i - \alpha)$ to simulate moving α reserves from t to s .



■ **Figure 4** Movement of capacities, reserves and cluster reserves when accepting a transaction x_i from v to u where $v, u \in C_i$ for some cluster C_i . The case on the top represents the typical movement of capacities and reserves when accepting x_i . The case on the bottom represents the treatment of the reserves (and cluster reserves) when the change in cluster reserves ($R_v^{C_\ell}$) is involved. The highlighted reserves are on v 's side of the channel between v and the middle node. Similar changes are made on u 's side of the channel between u and the middle node

437 ▶ **Lemma 8.** *Let S be a set of nodes that is connected to $C_i \setminus S$ for some cluster C_i by*
 438 *volume of at least $6|S| \cdot M$. Then at least $|S|M$ volume can be moved from S to $C_i \setminus S$.*

439 ▶ **Lemma 9.** *Let X'_t be a subsequence of transactions where transactions incident to C_i have*
 440 *volume at least $tr \geq M$, then Algorithm 5 can transport at least tr from any set of nodes*
 441 *$S \subset C_i$ to any other set $T \subset C_i$. The LP solution is modified by volume at most $\mathcal{O}(\sqrt{p} \cdot tr)$.*

442 We now have the requisite ingredients to state and prove (in Appendix A.7) the main
 443 theorem of this section.

444 ▶ **Theorem 10.** *Algorithm 2 is $\mathcal{O}(\sqrt{p})$ approximation of the problem where given a transaction*
 445 *sequence, the nodes are $(\sqrt{p}, \sqrt{p}, 24\sqrt{p})$ -clustered.*

446 We also get the following corollary as a direct consequence of the proof of Theorem 10.

447 ▶ **Corollary 11.** *Given (m, k, t) -clustered transaction sequence where $t \geq 24\sqrt{p}$, Algorithm 2*
 448 *is $\mathcal{O}(\max(m, k))$ -approximation of the problem.*

449 Finally, we end this section with two observations. The first observation effectively states
 450 that our algorithm just needs to focus on the total volume of partially accepted transactions
 451 and can ignore the volume of fully accepted and rejected transactions. The second observation
 452 allows us to further reduce the approximation ratio if we assume clustering of a larger depth.

453 *Practical relaxations.* We observe that the proof of Theorem 10 requires a condition
 454 on the partially accepted transactions (these are transactions which satisfy $0 < y_i < x_i$).
 455 However, we stress that the volume of fully accepted or rejected transactions can be arbitrary.
 456 Moreover, in every cluster, the middle node can interact with other clusters arbitrarily and
 457 this also does not affect the proof of Theorem 10.

458 *Hierarchical clustering.* We can use a similar idea to further reduce the approximation
 459 ratio if we assume that within each cluster the clusters themselves are also clustered. Namely,
 460 we first treat the top-level clusters as nodes and solve the problem as outlined by Algorithm 1.
 461 Then, we have some transactions pre-accepted or pre-rejected and we now shift the focus to
 462 the next level clusters. If these next-level clusters consist of other clusters, we treat them
 463 again as nodes. We do this recursively until we reach the lowest-level clusters and then
 464 use Algorithm 2 to solve the problem within the lowest-level clusters. Note that we require

the strength of the clustering to be bigger in order to achieve a similar approximation ratio in this setting. Nevertheless, assuming the nodes satisfy sufficiently strong clustering conditions, we can achieve an approximation ratio of $\mathcal{O}(p^{\frac{1}{k+1}})$ for two nested clusterings, and, in general, $\mathcal{O}(p^{\frac{1}{k+1}})$ for k nested clusterings (if the nodes in clusters are well distributed). We stress that although hierarchical clustering allows us to achieve a stronger approximation factor, it also comes with additional structural assumptions. However, we believe there are strong use cases of hierarchical clustering, for instance using a payment channel network to handle payments among users in a large company. The clusters would naturally follow the clustering of users into divisions and sub-divisions, with the head of each division being the centre of the star in their corresponding cluster.

6 Case study: the Lightning Network

To justify our transactions distribution assumption in Section 5, we perform an empirical case study of the largest and most commonly used PCN: Bitcoin’s Lightning Network. Specifically, the research question we want to investigate is the following:

Are transactions in the Lightning Network clustered according to the clustering conditions stated in Section 5?

To address this question, we obtained the latest snapshot (dated September 2023) of the Lightning Network from the Lightning Network gossip repository [1]. We further restrict our examination to the largest connected component which contains 20,348 users and 310,657 channels. Figure 5 shows the distribution of degrees in the largest connected component, where a large number of nodes ($\sim 10\%$) are of degree ≤ 2 .

Estimating the source and target of transactions. The main challenge in extracting transactions data from the Lightning Network is that the Lightning Network, by design, only provides information about the costs of transactions across channels, but does not provide information about transactions between users. Nevertheless, it is reasonable to assume that transactions between users are related to the channel capacities that connect them. Furthermore, we are mainly concerned with finding out if the clustering conditions hold for transactions between ordinary users with low degrees rather than between, e.g., large payment hubs with high degrees. We hence propose the following methodology: we first remove high degree nodes (i.e., nodes with degree ≥ 3) in the network and connect their neighbours to each other. Let us denote this resulting reduced network by G . In G , we assume nodes that are directly connected to each other transact with each other. Our methodology of generating the reduced network G is detailed in Algorithm 4.

Estimating the volume of transactions. Recall that a crucial clustering condition is that the volume of transactions that stay within a cluster is larger than the volume of the transactions that go between clusters. However, the privacy principles behind the design of the Lightning Network renders such transaction volume information hidden. Thus, a second challenge in our setting is extracting estimates of the size of transactions between users. In order to estimate transaction volume, we assume that the volume of transactions between two connected users in G is inversely proportional to the cost of a transaction. Namely, that the volume of transactions between users i and j is $1/(c_{i,j} + 1)$ where $c_{i,j}$ is the minimum cost of sending a transactions of 100 milli-satoshi from i to j .

Note that the volume of transactions is estimated under the simplifying assumptions that all transactions have the same fixed volume, and therefore the cost for a transaction between a pair of users is fixed and specified using $cost_H$ on input.

■ **Algorithm 4** Algorithm for estimating the communication network of regular users.

Input: Graph $H = (V, E)$, $cost_H : E \rightarrow \mathbb{R}_{\geq 0}$


Output: Graph G of users and vol assigning volume to its edges.

- 1: Find the largest strongly connected component $C \subseteq H$
- 2: $users \leftarrow \{v \in C \mid 1 \leq \deg(v) \leq 2\}$
- 3: $hubs \leftarrow V \setminus users$
- 4: $E_{users} \leftarrow \{\{u, v\} \mid u, v \in users \text{ and } N[u] \cap N[v] \cap hubs \neq \emptyset\} \cup E(G[users])$
- 5: $G \leftarrow (users, E_{users})$
- 6: $cost_G(u, v) \leftarrow cost_H(u, v)$ for all $\{u, v\} \in E_{users}$.
- 7: $cost_G(u, v) \leftarrow \infty$ for all $u, v \in users$ and $\{u, v\} \notin E_{users}$.
- 8: **for all** $v \in hubs$ **do**
- 9: **for** $s_1, s_2 \in users \cap N(v)$ **do**
- 10: $cost_G(s_1, s_2) = \min(cost_G(s_1, s_2), cost_H(s_1, v) + cost_H(v, s_2))$
- 11: $vol(u, v) \leftarrow 1/(cost_G(u, v) + 1)$ for all $u, v \in S$.
- 12: **return** G and vol


510 We start with the largest strongly connected component of H as it is the largest set of
 511 users such that any user can send a transaction to any other. Also, to simplify the experiment,
 512 we assume that the volume of transactions from a user u to another user v is the same as
 513 the volume of transactions in the reverse order (from v to u).

514 This approach of computing costs for each pair of users runs in time linear with the
 515 number of edges of G . Note that the number of updates of a cost is at most the number of
 516 edges adjacent to the set $users$, which is bounded by $O(|users|)$ as the degree of each vertex
 517 is bounded.

518 *Results.* Figure 6 and Table 1 shows that we obtain 18 clusters with average within-cluster
 519 transaction volume of 826.1, and between-cluster transaction volume of 4.27. This gives
 520 us a ratio of 193.2 of within-cluster and between-cluster transactions, which confirms our
 521 experimental hypothesis.



■ **Figure 5** Degree distribution of the nodes in the largest connected component of the Lightning Network.



■ **Figure 6** Clustering of the nodes in the reduced network G . Each color represents a different cluster.

■ **Table 1** Details of individual clusters.

Cluster	Cluster Size	Inside Volume	Between Volume	Ratio
1	473	47.85	2.73	17.53
2	309	4.32	0.33	13.09
3	136	2.97	0.66	4.5
4	126	33.32	0.29	114.9
5	123	2.18	1.67	1.31
6	118	3.52	0.29	12.14
7	96	715.15	0.33	2167.1
8	76	0.79	0.99	0.8
9	67	0.18	0.12	1.5
10	64	0.12	0.07	1.7
11	59	4.20	0.00	∞
12	57	0.37	0.13	2.8
13	54	0.65	0.00	∞
14	54	3.09	0.71	4.35
15	31	1.06	0.04	26.5
16	30	0.22	0.11	2.0
17	25	6.07	0.04	151.8
18	19	0.04	0.04	1.0

522 **7 Conclusion**

523 In this work, we generalise a model from processing transactions over a single payment
524 channel to an entire PCN. To this end, we first define a network creation stage of the problem
525 which examines how to create a network with minimal capacity cost, and we show that the
526 star is a 2-approximation of the optimal capacity cost. We then propose two algorithms to
527 address weighted transaction processing over a network. The first algorithm assumes a star
528 topology over the p parties, but apart from that works without any further assumptions on
529 the transaction sequence, and gives an approximation ratio of $\mathcal{O}(p)$. The second algorithm
530 assumes a double star topology and a clustering of the transaction sequence, and gives an
531 approximation ratio of $\mathcal{O}(\sqrt{p})$. Finally, we also perform an empirical case study estimating
532 transaction information from the topology of the Lightning Network, and our estimates
533 suggest that our clustering assumptions regarding the input transaction sequence are close
534 to reality.

535 We believe our work is an important first step in the direction of designing optimal networks
536 and algorithms for transaction processing in PCN. We highlight two interesting directions
537 for future work. First, it would be interesting to extend our model and algorithms to the
538 online setting, which will also complement the analysis of [4, 7]. Secondly, developing better
539 estimates regarding transaction information in real world PCNs would also be paramount to
540 developing better optimisation algorithms and heuristics.

541 **References**

- 542 1 Lightning network topology datasets. URL: <https://github.com/lnresearch/topology>.
- 543 2 Raiden network. <https://raiden.network/>, 2017.

- 544 3 Hagit Attiya, Leah Epstein, Hadas Shachnai, and Tami Tamir. Transactional contention
545 management as a non-clairvoyant scheduling problem. *Algorithmica*, 57(1):44–61, 2010.
- 546 4 Georgia Avarikioti, Kenan Besic, Yuyi Wang, and Roger Wattenhofer. Online payment network
547 design. In *DPM/CBT@ESORICS*, volume 11737 of *Lecture Notes in Computer Science*, pages
548 307–320. Springer, 2019.
- 549 5 Zeta Avarikioti, Lioba Heimbach, Yuyi Wang, and Roger Wattenhofer. Ride the lightning:
550 The game theory of payment channels. In *Financial Cryptography*, volume 12059 of *Lecture*
551 *Notes in Computer Science*, pages 264–283. Springer, 2020.
- 552 6 Zeta Avarikioti, Tomasz Lizurej, Tomasz Michalak, and Michelle Yeo. Lightning creation
553 games. *CoRR*, abs/2306.16006, 2023.
- 554 7 Mahsa Bastankhah, Krishnendu Chatterjee, Mohammad Ali Maddah-Ali, Stefan Schmid,
555 Jakub Svoboda, and Michelle Yeo. R2: boosting liquidity in payment channel networks with
556 online admission control. In *FC (1)*, volume 13950 of *Lecture Notes in Computer Science*,
557 pages 309–325. Springer, 2023.
- 558 8 Anamika Chauhan, Om Prakash Malviya, Madhav Verma, and Tejinder Singh Mor. Blockchain
559 and scalability. In *QRS Companion*, pages 122–128. IEEE, 2018.
- 560 9 Chandra Chekuri, Sanjeev Khanna, and F Bruce Shepherd. The all-or-nothing multicommodity
561 flow problem. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*
562 (*STOC*), pages 156–165, 2004.
- 563 10 Anne Condon and Richard M. Karp. Algorithms for graph partitioning on the planted partition
564 model. *Random Struct. Algorithms*, 18(2):116–140, 2001.
- 565 11 Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed E. Kosba,
566 Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Watten-
567 hofer. On scaling decentralized blockchains - (A position paper). In *Financial Cryptography*
568 *Workshops*, volume 9604 of *Lecture Notes in Computer Science*, pages 106–125. Springer, 2016.
- 569 12 Christian Decker, Rusty Russell, and Olaoluwa Osuntokun. eltoo: A simple layer2 protocol
570 for bitcoin. <https://blockstream.com/eltoo.pdf>, 2018.
- 571 13 Christian Decker and Roger Wattenhofer. A fast and scalable payment network with bitcoin
572 duplex micropayment channels. In *Symposium on Self-Stabilizing Systems (SSS)*, pages 3–18.
573 Springer, 2015.
- 574 14 Oguzhan Ersoy, Stefanie Roos, and Zekeriya Erkin. How to profit from payments channels. In
575 *Financial Cryptography*, volume 12059 of *Lecture Notes in Computer Science*, pages 284–303.
576 Springer, 2020.
- 577 15 Paolo Guasoni, Gur Huberman, and Clara Shikhelman. Lightning network economics: Topology.
578 *SSRN Electronic Journal*, 2023. doi:10.2139/ssrn.4439190.
- 579 16 Paolo Guasoni, Gur Huberman, and Clara Shikhelman. Lightning network economics: Channels.
580 *Manag. Sci.*, 70(6):3827–3840, 2024.
- 581 17 Rachid Guerraoui, Maurice Herlihy, and Bastian Pochon. Polymorphic contention management.
582 In Pierre Fraigniaud, editor, *Distributed Computing*, pages 303–323, Berlin, Heidelberg, 2005.
583 Springer Berlin Heidelberg.
- 584 18 Rachid Guerraoui, Maurice Herlihy, and Bastian Pochon. Towards a theory of transactional
585 contention managers. In *PODC*, pages 316–317. ACM, 2006.
- 586 19 Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, and Dragos-Adrian Seredin-
587 schi. The consensus number of a cryptocurrency. In *PODC*, pages 307–316. ACM, 2019.
- 588 20 Piyush Kumar Gupta and Panganamala Ramana Kumar. The capacity of wireless networks.
589 *IEEE Trans. Inf. Theory*, 46:388–404, 2000.
- 590 21 Maurice Herlihy. Wait-free synchronization. *ACM Trans. Program. Lang. Syst.*, 13(1):124–149,
591 January 1991. doi:10.1145/114005.102808.
- 592 22 Paul Holland, Kathryn B. Laskey, and Samuel Leinhardt. Stochastic blockmodels: First
593 steps. *Social Networks*, 5:109–137, 1983. URL: [https://api.semanticscholar.org/CorpusID:](https://api.semanticscholar.org/CorpusID:34098453)
594 34098453.

- 595 23 Anurag Jain, Shoeb Siddiqui, and Sujit Gujar. We might walk together, but I run faster: Network fairness and scalability in blockchains. In *AAMAS*, pages 1539–1541. ACM, 2021.
- 596
- 597 24 Patrick McCorry, Surya Bakshi, Iddo Bentov, Sarah Meiklejohn, and Andrew Miller. Pisa: Arbitration outsourcing for state channels. In *AFT*, 2019. URL: 10.1145/3318041.3355461.
- 598
- 599 25 Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments. <https://lightning.network/lightning-network-paper.pdf>, 2015.
- 600
- 601 26 Prabhakar Raghavan and Clark D Thompson. Provably good routing in graphs: regular arrays. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 79–87, 1985.
- 602
- 603
- 604 27 Stefan Schmid, Jakub Svoboda, and Michelle Yeo. Weighted packet selection for rechargeable links in cryptocurrency networks: Complexity and approximation. In *SIROCCO*, volume 13892 of *Lecture Notes in Computer Science*, pages 576–594. Springer, 2023.
- 605
- 606

607 **A Omitted proofs**

608 **A.1 Proof of Lemma 1**

609 ► **Lemma 1.** *Given the sequence of transactions \mathcal{S} , we have $LP(\mathcal{S}, K_p) \leq AC(\mathcal{S}, G)$.*

610 **Proof.** We first observe that adding channels to a graph can only decrease the capacity cost (adding channels can only make the shortest paths shorter). This means that the capacity cost is the lowest for the complete graph, no matter what is the optimal graph topology (as computed by minimising the *total* cost of both capacity and channel creation). The LP computes the optimal capacity for the complete graph K_p , which is a superset of every graph, including the optimal graph. ◀

616 **A.2 Proof of Theorem 2**

617 ► **Theorem 2.** *For every graph G over p nodes, and any graph G' with diameter d , the following inequality holds*

618

$$619 \quad AC(\mathcal{S}, G') \leq d \cdot AC(\mathcal{S}, G).$$

620 **Proof.** Since for the complete graph K_n , we have $AC(\mathcal{S}, K_n) \leq AC(\mathcal{S}, G)$, we simply need to show that $AC(\mathcal{S}, G') \leq d \cdot AC(\mathcal{S}, G)$. We modify the solution of the LP for the complete graph by adding at most d -times the capacity and prove that this graph processes all transactions.

623 We start with graph G' with zero capacities. For all pairs of nodes v and u , let the capacity between them be $w_{v,u}$. We find the shortest path between them in G' , since the diameter is at most d , the length of the path is at most d . Then, we increase the capacity of every channel on the path by $w_{v,u}$. We denote this capacity to a special bucket that will be used to process transactions that the LP sends along u, v . At the start at edge k, l , where k is closer to v and l is to u , we increase the balance at the k 's side by $C_{v,u}^0$ (and l 's balance is $C_{u,v}^0$), the balance at the edge v, u at the start as given by the LP.

630 After we increase the capacities of G' for every pair of nodes, we process the transactions. Any time the transaction of size s was forwarded along some edge u, v , we can transfer s on the shortest path between u and v from the bucket on the edge denoted u, v . If the bucket for some edge becomes negative, the solution of LP is infeasible.

634 By this, we can process all transactions. Since the diameter is at most d , we know that the increase in capacity is at most d -times. ◀

635

A.3 Proof of Theorem 6

► **Theorem 6.** *Given the transaction sequence X_t , let \mathcal{C} be the optimal capacity cost and \mathcal{R} the optimal rejection cost on a star. Then Algorithm 1 run for every pair of nodes incurs a rejection cost of at most $(\sqrt{3} + 1)\mathcal{R}$ and capacity cost of at most $(1 + (p - 1) \cdot \sqrt{3})\mathcal{C}$.*

Proof. Every node has $p - 1$ connections while the size of the channel is viewed as the minimum of the two sizes as denoted by the LP. The reserves are $\sqrt{3}$ times that value. Moreover, every channel in a star has funds amount as dictated by the LP. Altogether it gives $1 + (p - 1) \cdot \sqrt{3}$ times of \mathcal{C} .

The algorithm ensures that the in every channel on the right side, there is at least the amount prescribed by the LP. This means that all fully accepted ($x_i = y_i$) transactions are accepted. Moreover, from [27], we know that the rejection cost between two nodes is at most $\sqrt{3} + 1$ times the rejection cost of the optimal rejection cost \mathcal{R} . Thus, Algorithm 1 incurs a rejection cost of at most $(\sqrt{3} + 1)\mathcal{R}$ and capacity cost of at most $(1 + (p - 1) \cdot \sqrt{3})\mathcal{C}$. ◀

A.4 Proof of Lemma 7

► **Lemma 7.** *If there is a transaction (s, t, x_i, y_i) , we can change it to*

■ $(s, t, x_i, y_i + \alpha)$ to simulate moving α reserves from s to t , and

■ $(s, t, x_i, y_i - \alpha)$ to simulate moving α reserves from t to s .

Proof. Every node has $p - 1$ connections while the size of the channel is viewed as the minimum of the two sizes as denoted by the LP. The reserves are $\sqrt{3}$ times that value. Moreover, every channel in a star has funds amount as dictated by the LP. Altogether it gives $1 + (p - 1) \cdot \sqrt{3}$ times of \mathcal{C} .

The algorithm ensures that the in every channel on the right side, there is at least the amount prescribed by the LP. This means that all fully accepted ($x_i = y_i$) transactions are accepted. Moreover, from [27], we know that the rejection cost between two nodes is at most $\sqrt{3} + 1$ times the rejection cost of the optimal rejection cost \mathcal{R} . Thus, Algorithm 1 incurs a rejection cost of at most $(\sqrt{3} + 1)\mathcal{R}$ and capacity cost of at most $(1 + (p - 1) \cdot \sqrt{3})\mathcal{C}$. ◀

A.5 Proof of Lemma 8

► **Lemma 8.** *Let S be a set of nodes that is connected to $C_i \setminus S$ for some cluster C_i by volume of at least $6|S| \cdot M$. Then at least $|S|M$ volume can be moved from S to $C_i \setminus S$.*

Proof. From the condition 3 of the clustering conditions, we know that at least $2|S|M$ volume moves from the cluster and $2|S|M$ volume to the cluster. For the sake of contradiction, suppose that we can move at most $m < |S|M$ from S to $C_i \setminus S$.

From Lemma 7, we know that we cannot move only on fully accepted transactions from S and rejected transactions from $C_i \setminus S$. Let m' be the incoming volume to S that can be moved. Then in the subsequence in the solution of LP , S receives m' volume and gives up $2|S|M - (m - m')$ volume. That means on average it lost at least $2|S|M - (m - m') - m' = 2|S|M - m > |S|M$. From the Dirichlet principle, at least one node lost more than M . This is impossible since M is the highest capacity in C_i and the solution of the LP ensures that for any subsequence the difference is at most M . ◀

675 A.6 Proof of Lemma 9

676 ► **Lemma 9.** *Let X'_t be a subsequence of transactions where transactions incident to C_i have*
 677 *volume at least $tr \geq M$, then Algorithm 5 can transport at least tr from any set of nodes*
 678 *$S \subset C_i$ to any other set $T \subset C_i$. The LP solution is modified by volume at most $\mathcal{O}(\sqrt{p} \cdot tr)$.*

679 **Proof.** Since there was tr transported, inside the cluster there was at least $24 \cdot tr \cdot |C_i|$
 680 transported, and next to S , there was at least $6tr|S|$ transported. From Lemma 8, that
 681 means that at least $tr|S|$ goes away from S to other nodes.

682 Let $S_0 = S$ and S_i be the nodes that we can transport nonzero amount from S_{i-1} .
 683 Similarly, let $T_0 = T$ and T_i be the nodes that we can transport nonzero amount to T_{i-1} .
 684 Since the clustering conditions, we have that the size of S_i and T_i is increasing if their size is
 685 at most half of the cluster. That means they intersect at some point.

686 We look at the minimum mn that can be transported and we change the transactions
 687 such that in v the reserves are decreased by mn and in u they are increased by mn . After
 688 the change, we know that the clustering conditions still hold, now we just need to decrease
 689 the inequalities by mn . Therefore, we can find the path again. ◀

690 A.7 Proof of Theorem 10

691 ► **Theorem 10.** *Algorithm 2 is $\mathcal{O}(\sqrt{p})$ approximation of the problem where given a transaction*
 692 *sequence, the nodes are $(\sqrt{p}, \sqrt{p}, 24\sqrt{p})$ -clustered.*

693 **Proof.** We are using a double star, which means it is at least a 4 approximation of the
 694 optimal topology from Theorem 2.

695 The algorithm uses Algorithm 1 for every pair of clusters (between middle nodes) to
 696 accept and reject transactions between them. This incurs cost for additional capacity equal
 697 $\mathcal{O}(\sqrt{p})$ times the cost of the capacity cost of the LP for the capacity between clusters. The
 698 rejection cost is at most $(\sqrt{3} + 1)$ times higher than the rejection cost between clusters. All
 699 of this follows from Theorem 6.

700 Next, from Lemma 9, we know that we can transport reserves such that everything
 701 between clusters is accepted or rejected from the previous algorithm. It changes the values
 702 of the LP by at most $\mathcal{O}(\sqrt{p})$ times the volume of rejected transactions between clusters.

703 Finally, again from Theorem 6, we can solve the problem inside the clusters. This incurs
 704 capacity cost $\mathcal{O}(\sqrt{p})$ times the capacity cost of the solution inside the clusters. Moreover,
 705 the cost for rejection is $(\sqrt{3} + 1)$ times the rejection cost in the modified solution of LP (X'_t).
 706 The modified rejection cost of solution of LP is the rejection cost of the original solution plus
 707 at most $\mathcal{O}(\sqrt{p})$ times the rejection cost of transactions between clusters (from Lemma 9).
 708 Altogether, we get that the algorithm is $\mathcal{O}(\sqrt{p})$ approximation algorithm. ◀

709 B Omitted details of star topology solution

710 B.1 Modified channel algorithm details

711 We now describe the decision making process in Algorithm 1 for transactions going from u to
 712 v (the other direction is symmetric). Algorithm 1 accepts a transaction as long as the reserve
 713 R_u is above a threshold (i.e., $R_u \geq \frac{\sqrt{3}-1}{2}M$, line 6). Accepting a transaction decreases the
 714 capacities C_u and R_u^v of u by y_i and $x_i - y_i$ respectively (refer to the “accept” case in Figure 2).
 715 If accepting the transaction would make R_u decrease by too much and the LP demands to
 716 accept the transaction *weakly* (i.e., $\frac{y_i}{x_i} < \frac{\sqrt{3}}{\sqrt{3}+1}$, line 9), the algorithm rejects the transaction.

Rejecting a transaction from say u to v decreases both C_u and R_u^v of by y_i (refer to the “reject” case in Figure 2). If the reserve on u ’s side after accepting the transaction would be too small (i.e., $R_u < \frac{\sqrt{3}-1}{2}M$, line 17), the algorithm looks at all subsequent transactions between u and v in the transaction sequence. It accepts all transactions that increases R_u while rejecting weakly accepted transactions from u to v (rejecting them also increases R_u). From *strongly* accepted transactions (i.e., $\frac{y_i}{x_i} \geq \frac{\sqrt{3}}{\sqrt{3}+1}$) from u to v , the algorithm tries to accept all of them while simulating the movement of capacities between the reserves until either $R_u > \frac{\sqrt{3}-1}{2}M$ or $R_u < 0$. If $R_u > \frac{\sqrt{3}-1}{2}M$, then all transactions can be accepted. If $R_u < 0$, the algorithm rejects strongly accepted transactions until $R_u > \frac{\sqrt{3}-1}{2}M$. The proof in [27] shows that doing so maintains the approximation ratio at $\sqrt{3} + 1$ for the rejection cost.

Algorithm details. The simulating of the reserves is implemented by a heap. The heap is empty while $R_u \geq \frac{\sqrt{3}-1}{2}M$, and when $R_u < \frac{\sqrt{3}-1}{2}M$ it starts to fill up. When $R_u < 0$, then the decision is changed for transactions from the top of the heap until $R_u \geq \frac{\sqrt{3}-1}{2}M$ again.

Observe that for Algorithm 1 the funds on both sides can shift and the reserves are not affected. This allows us to compose the algorithm from pairwise reserves with shared funds denoted by the linear program.

The algorithm describes decisions for the star, so a transaction goes from u to c (star center) and then from c to v . That means it needs to traverse two channels. The algorithm sees it as only pairwise interaction. Since the decisions for the two channels are always the same for both channels (a transaction is either accepted or rejected for both), we can have the reserves stored twice, once between u and c , the second time between v and c . The value R_u is then the reserves at u in channel u, c and at c in the channel c, v . Both of these reserves move in sync.

C Omitted details of clustering algorithm

C.1 Stochastic Block Model details

The planted partition model creates a random graph where nodes are partitioned into clusters. Inside a cluster, two nodes are connected with probability q_1 , otherwise, they are connected with probability q_2 . The planted partition model creates an unweighted and undirected graph. Our graph created by the packets is directed, weighted, and temporal. These changes call for modification of the planted partition model.

The biggest change to the model is brought by the temporality. Some edges follow after others, we require that all of the conditions hold for some continuous subsequence of edges incident to some cluster of bigger size.

To make the graph weighted, we view q_1 and q_2 as ratios between volumes of the packets between the clusters and inside the cluster. This view is expressed by the condition 1 of the clustering conditions.

To make the graph directed, we allow for arbitrary orientation of the edges with the condition that at least $\frac{1}{3}$ of the volume needs to be directed in both directions. It is expressed by condition 2 of the clustering conditions. Note that $\frac{1}{3}$ of the volume can be oriented in any direction, no matter the other packets.

Finally, the condition 2 of the clustering conditions ensures that some volume is incident to every node.

Note that all of the conditions would be satisfied with high probability in the stochastic block model. Moreover, there are other conditions that would be satisfied (for instance any

two clusters would send packets between each other at the same rate) that we do not require. From the above the graph generated by the planted partition model with \sqrt{p} partitions of size \sqrt{p} and $q_2 = \frac{1}{p}$ and $q_1 = 1 - q_2$ satisfies all clustering conditions with high probability.

C.2 Linear program for double star

Let the set of leaf nodes be L and the set of middle nodes be M , and c denote the central node of the star. For a leaf node u , let $M(u)$ be the middle node corresponding to u 's cluster.

For a leaf node v and natural number i , we use $C_{v,i}$ and $C'_{v,i}$ to denote the capacity at the channel $(v, M(v))$ after processing the i -th transaction on the side of v and $M(v)$ respectively. Moreover let C_v denote the total capacity of the channel $(v, M(v))$, i.e., $C_v = C_{v,i} + C'_{v,i}$. For a middle node v and i , we use $M_{v,i}$ and $M'_{v,i}$ to denote the capacity at the channel (v, c) after processing the i -th transaction on the side of v and c respectively. Moreover, let M_v denote the total capacity of the channel (v, c) .

$$\text{minimise } \sum_{u \in L} C_u + \sum_{u \in M} M_u + \quad (2)$$

$$\sum_i f(x_i - y_i) + \frac{m(x_i - y_i)}{x_i} \quad (3)$$

$$\forall i : 0 \leq y_i \leq x_i$$

$$\forall (v, u, x_i) : C_v, C_u \geq x_i$$

$$C_{v,i} = C_{v,i-1} - y_i$$

$$C'_{v,i} = C'_{v,i-1} + y_i$$

$$C'_{u,i} = C'_{u,i-1} - y_i$$

$$C_{u,i} = C_{u,i-1} + y_i$$

$$\forall (v, u, x_i) \text{ between clusters: } M_{M(u)}, M_{M(v)} \geq x_i$$

$$M_{M(v),i} = M_{M(v),i-1} - y_i$$

$$M'_{M(v),i} = M'_{M(v),i-1} + y_i$$

$$M'_{M(u),i} = M'_{M(u),i-1} - y_i$$

$$M_{M(u),i} = M_{M(u),i-1} + y_i$$

C.3 Algorithm to transport cluster reserves

Here we detail Algorithm 5, which transfers cluster reserves from well-funded to depleted nodes *within* a cluster, so that the cluster nodes have sufficient reserves to handle transactions that go *between* the clusters.

■ **Algorithm 5** Transporting reserves between nodes.

Input: Subsequence of transactions X_t inside one cluster, two nodes v, u , amount A .

Output: Modified X_t such that A in reserves is transported from v to u .

```

1: procedure TRANSPORT( $X_t, v, u, A$ )
2:    $M \leftarrow \text{PATHBETWEEN}(v, u)$ 
3:    $mn \leftarrow \text{MINIMUMONPATH}(M)$ 
4:    $mn \leftarrow \min(mn, A)$ 
5:   for  $(dir, x, y, x_i, y_i) \in M$  do
6:     if  $dir$  from  $v$  to  $u$  then
7:        $y_i \leftarrow y_i - mn$ 
8:     else
9:        $y_i \leftarrow y_i + mn$ 
10:  if  $A - mn > 0$  then
11:    TRANSPORT( $X_t, v, u, A - mn$ )

```
