

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/175234/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Corcoran, Pdraig and Lewis, Rhydian 2025. An analysis of the correctness and computational complexity of path planning in Payment Channel Networks. *The Journal of Financial Technology*

Publishers page:

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



---

# AN ANALYSIS OF THE CORRECTNESS AND COMPUTATIONAL COMPLEXITY OF PATH PLANNING IN PAYMENT CHANNEL NETWORKS

PADRAIG CORCORAN

*School of Computer Science & Informatics, Cardiff University, Cardiff, Wales, UK.*  
corcoranp@cardiff.ac.uk

RHYD LEWIS

*School of Mathematics, Cardiff University, Cardiff, Wales, UK.*  
lewisr9@cardiff.ac.uk

---

## Abstract

Payment Channel Networks (PCNs) are a method for improving the scaling and latency of cryptocurrency transactions. For a payment to be made between two peers in a PCN, a feasible low-fee path in the network must be planned. Many PCN path planning algorithms use a search algorithm that is a variant of Dijkstra's algorithm. In this article, we prove the correctness and computational complexity of this algorithm. Specifically, we show that, if the PCN satisfies a consistency property relating to the fees charged by payment channels, the algorithm is correct and has polynomial computational complexity. However, in the general case, the algorithm is not correct and the path planning problem is NP-hard. These newly developed results can be used to inform the development of new or existing PCNs amenable to path planning. For example, we show that the Lightning Network, which is the most widely used PCN and is built on the Bitcoin cryptocurrency, currently satisfies the above consistency property. As a second contribution, we demonstrate that a small modification to the above path planning algorithm which, although having the same asymptotic computational complexity, empirically shows better performance. This modification involves the use of a bidirectional search and is empirically evaluated by simulating transactions on the Lightning Network.

**Keywords.** payment channel networks; path planning; lightning network.

# 1 Introduction

A cryptocurrency is a digital currency that is decentralised in nature and not issued or controlled by a central authority. The most famous and successful example of a cryptocurrency is Bitcoin, invented by Satoshi Nakamoto [1]. Cryptocurrencies are enabled by the use of blockchain systems, which are a type of distributed system that maintains a ledger or database of digital currency transactions. A necessary component of any blockchain system is a *consensus algorithm*. This is an algorithm that allows all peers in the system to achieve agreement regarding the state of the ledger [2]. As a consequence of the need to ensure decentralisation and security, many consensus algorithms exhibit scalability issues relating to low transaction throughput (the number of transactions per second) and transaction confirmation latency (the time until a transaction can be considered settled). The former property can, in turn, lead to high transaction fees when large sets of transactions are competing to get into the blockchain. For example, Bitcoin can currently only process seven transactions per second and has a transaction confirmation latency of approximately ten minutes. This is assuming that the transaction is actually included in the very next block and is considered confirmed immediately when this happens. In contrast, the VISA payment system can process 65,000 transactions per second and has a transaction confirmation latency of just a few seconds. The low transaction throughput of Bitcoin means that, in times of high demand, transaction fees can be expensive and greatly exceed the cost of many everyday purchases.

To overcome the above scaling challenge, several potential solutions have been proposed [3]. One promising direction of research is *payment channel networks* (PCNs) that allow the consolidation of a larger set of transactions into a smaller set, where individual transactions can be confirmed almost instantly. The Lightning Network (LN) is the most widely used PCN and is designed to work with Bitcoin [4]. To illustrate the potential benefits of PCNs, consider the case where Alice is a merchant and Bob is a customer. Alice makes regular small payments to Bob, and sometimes Bob makes refunds to Alice. Instead of submitting each transaction to be processed immediately by the consensus algorithm, Alice and Bob can establish a payment channel between them that maintains a balance that is updated after each transaction. When Alice and Bob have completed their transactions, they submit the final balance to be processed by the consensus algorithm. Since only a single transaction is processed by the consensus algorithm, this can increase transaction throughput significantly. Furthermore, since each transaction is only processed locally, the transaction latency is reduced to the time required to send a small number of network messages between Alice and Bob. By establishing a network of payment channels, individuals can also make payments to others who they do not share a direct payment channel with. For example, consider the case where Alice wants to make a payment to Charlie but has not established a corresponding payment channel. If Alice has an existing payment channel to Bob, and Bob has an existing payment channel to Charlie, then Alice can make the payment to Bob and Bob can, in turn, forward the payment to Charlie.

PCNs are typically implemented using smart contracts, which are pieces of code executed on a blockchain that ensure the security of the payments.

Despite the potential of PCNs, the technology still faces several challenges. This work focuses on path planning which concerns finding a useful path in the network of payment channels for making a given payment. There exists no universal method for performing path planning in most PCNs including the LN. Instead, different implementations use different methods for solving this problem [5]. Given that the size of PCNs can be very large, it is important that path planning is performed in a correct and computationally efficient manner. Most current path planning algorithms use a variant of Dijkstra’s algorithm. Dijkstra’s algorithm assumes that the network in question contains arcs with constant and non-negative weights. This assumption does not hold in the LN where arc weights correspond to fees charged for forwarding payments and vary as a function of the payment amount. Hence, existing correctness and computational complexity results for Dijkstra’s algorithm do not generalise to this variant. In fact, to the authors’ knowledge, there exists no analysis of this algorithm’s correctness and computational complexity in this setting. In this work, we address this research gap. These new results tell us that, if the PCN in question satisfies a consistency property relating to the fee charged for forwarding payments, the algorithm is correct and has polynomial computational complexity. Otherwise, in the general case, the algorithm is not correct and the path planning problem is NP-hard. In turn, we show that the LN satisfies this property. A second contribution made in this article is the proposal of a small modification to the above path planning algorithm which, though having the same asymptotic computational complexity, empirically shows better performance. This modification involves the use of a bidirectional search.

The remainder of this paper is structured as follows. In Section 2 we review related works on path planning in PCNs. In Section 3 we formally define the problem of path planning in the LN. This section also identifies a connection between LN path planning and a specific *recurrence relation*, which we solve. In Sections 4 and 5 we formally describe the variant of Dijkstra’s algorithm mentioned above and provide an analysis of its correctness and computational complexity. In Section 6 we then describe the proposed modification to this algorithm that uses bidirectional search and prove its correctness and computational complexity. In Section 7 we then demonstrate the computational benefits of this method empirically. Finally, in Section 8 we conclude this work.

## 2 Related Works

As discussed in the introduction to this article, a variant of Dijkstra’s algorithm is a path planning algorithm frequently used with PCNs. Given the importance of path planning, several other algorithms have also been proposed. The remainder of this section presents

a review of these algorithms followed by a survey of other relevant related topics. A more in-depth review can be found in [6].

Solutions to the path planning problem can broadly be categorised as centralised methods and decentralised methods. In centralised methods, a single peer independently computes the entire payment path. An application of centralised methods, commonly known as source-based routing, involves the payment sender computing a payment path and using onion routing to make the payment using this path. Onion routing ensures that each peer in the path only learns the identity of the peer immediately before and the peer immediately after them in the path. They do not learn the payment sender, receiver or amount. However, since the path planning task is performed exclusively by the payment sender, source-based routing requires that the sender has significant computational resources and knowledge of the current state of the network. A variant of Dijkstra's algorithm is the most commonly used centralised path planning method. In fact, this method is used by most LN node implementations including LND and Core Lightning which are the most commonly used implementations. Several authors have also proposed centralised path planning methods that compute a set of distinct paths between the source and destination vertices in PCNs [7, 8, 9]. This allows multi-part payments to be made, whereby a single payment is split into a set of smaller payments, and each is routed using a distinct path. The main motivations for multi-part payments are to allow larger payments to be made given the finite capacity of individual payment channels and to obfuscate the total amount being transferred.

In decentralised path planning methods, each peer that the payment passes through learns the payment receiver and uses this information to independently decide the next peer to forward the payment to until it is eventually arrives at the receiver. Decentralised methods do not require the payment sender to have significant computational resources or knowledge of the current state of the network. However, since each peer in the payment path learns the identity of the payment receiver, they provide relatively poor privacy. Hence decentralised methods are not used by most LN node implementations. Roos et al. [10] proposed a greedy decentralised path planning algorithm entitled SpeedyMurmurs. This algorithm assigns coordinates to each vertex in the network such that a path from a given source vertex can be obtained by iteratively selecting arcs to vertices closer to the destination vertex. Similarly, Prihodko et al. [11] proposed a decentralised path planning algorithm called Flare. In this method, each peer maintains a routing table that contains payment paths to both peers in their local neighbourhood and a small subset of other peers known as beacons. To determine a payment path, the routing tables corresponding to both the sender and receiver are considered. If both routing tables contain a common peer, a payment path is constructed by combining payment paths to and from this peer. Lin et al. [12] proposed a path planning method that partitions the PCN in a set of regions and elects a peer in each region to be a corresponding beacon. Using this method, path planning is performed by forwarding the payment in question to a beacon. This beacon then forwards the payment to the receiver if the receiver is contained in their

respective region. Otherwise, the beacon forwards the payment to the beacon corresponding to the region within which the receiver is contained.

To transfer a given amount along a given path in a PCN, all the channels in question must have sufficient balance or liquidity in the direction in question. To help overcome this challenge, several rebalancing methods have been proposed to adjust channel balances to make paths feasible. Broadly speaking, these methods involve peers making cyclic payments to themselves that move a portion of the balance on one of their channels to another of their channels [13]. For privacy reasons, the channel balances are not shared among peers in a PCN because by monitoring changes in channel balances, an adversary can infer the source, destination and amount of all payments. A consequence of this is that path planning reduces to a trial and error process whereby different paths are attempted until a feasible one with sufficient balance to perform the payment in question is found. Tang et al. [14] considered whether it is possible for channels to share noisy balances while still maintaining privacy and found this to be challenging. Dotan et al. [15] proposed a method that is robust to balance probing attacks that attempt to determine arc balances. This is achieved by adding noise to a channel's response to a request to forward a payment.

In a PCN, when a channel forwards a payment intended for a third party, they charge a fee for providing this service. Several methods have been proposed to automatically determine a suitable fee to charge, considering objectives such as maximising profit or maintaining a desired channel balance [16, 17, 18]. However, by monitoring changes in fees charged and combining this with knowledge of how fees are determined, an adversary can potentially infer information about the transactions being processed by different channels. Tochner et al. [19] demonstrated that, by exploiting the fact that most path planning algorithms select paths with the lowest fees, an adversary can establish channels with low fees that a significant percentage of payments use. These channels can then be used to perform several attacks, including a denial-of-service (DoS) attack.

Several works have also analysed the topological structure of PCNs. Martinazzi et al. [20] and Seres et al. [21] demonstrated that the LN tends to exhibit small-world network properties with a small proportion of nodes with high centrality that act as hubs. Rohrer et al. [22] examined the resilience of the LN with respect to random failures and targeted attacks. The authors also demonstrated that the LN exhibits small-world and scale-free network properties. Zabka et al. [23] performed a centrality analysis of the LN. Similarly, Kotzer et al. [24] demonstrated that Braess's paradox may exist in a PCN, whereby adding additional channels may negatively impact the ability to successfully make payments.

### 3 Problem Definition

In this work, we model a snapshot of the Lightning Network (LN) at a given moment in time as a directed graph  $G = (V, E)$ . We use the term “snapshot” because both the topology and channel properties of the LN are dynamic. The topology is dynamic in the sense that, over time, new channels can be created and existing ones can be deleted. LN channel properties are also dynamic in the sense that, over time, the properties described below of a given channel can change.

In our model, each LN peer is modelled as a vertex  $v \in V$ . Since LN channels are bidirectional, each LN channel between the vertices  $v_i$  and  $v_j$  is modelled as a pair of arcs  $(v_i, v_j) \in E$  and  $(v_j, v_i) \in E$ . Each arc has a balance that equals the maximum payment amount that can be transferred along that arc and is modelled as a map  $b : E \rightarrow \mathbb{Z}^{\geq}$ . For a given LN channel, the sum of the corresponding pair of arc balances is known as the capacity of the channel in question. Although the capacity of a given channel remains constant, the corresponding arc balances change as amounts are transferred between the vertices in question. As an example, consider the pair of arcs  $(v_i, v_j) \in E$  and  $(v_j, v_i) \in E$  corresponding to a channel between the vertices  $v_i$  and  $v_j$ . If the amount  $a \in \mathbb{Z}^{\geq}$  is transferred from  $v_i$  to  $v_j$  using this channel, the balance  $b(v_i, v_j)$  reduces by  $a$  while the balance  $b(v_j, v_i)$  increases by  $a$ . To transfer the amount  $a$  along the arc  $(v_i, v_j)$ , the condition  $a \leq b(v_i, v_j)$  must therefore be satisfied. That is, the arc in question must have a sufficient balance.

If one wishes to transfer an amount  $a$  along an arc  $(v_i, v_j) \in E$ , the arc will also charge a fee for this service. This fee is charged by the arc operator, which is the source vertex  $v_i$ , and is parameterised by two values: a *base fee* and a *fee rate*. The base fee is a fixed value, while the fee rate is a proportion of the amount being transferred. We model the base fees and fee rates as maps  $f_b : E \rightarrow \mathbb{Z}^{\geq}$  and  $f_r : E \rightarrow \mathbb{R}^{\geq}$  respectively. In turn, we model the total fee charged by an arc  $e \in E$  to transfer an amount  $a \in \mathbb{Z}$  as the map  $f_{br} : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$ :

$$f_{br}(e, a) = f_b(e) + f_r(e) \times a. \quad (1)$$

To illustrate the above features of the LN, consider the toy LN displayed in Figure 1. This graph contains a single arc  $(v_i, v_j)$  that has a base fee of 2, a fee rate of 0.1 and a balance of 20. Consider the case where we wish to transfer an amount 10 from  $v_i$  to  $v_j$  along this arc. This amount is less than or equal to the arc balance making the transfer feasible. To perform this transfer, vertex  $v_i$  charges a fee of  $2 + 0.1 \times 10 = 3$ . An important point to highlight is that, to transfer an amount  $a$  from a vertex  $v_i$  to a vertex  $v_j$  using an arc  $(v_i, v_j)$ , we must first transfer  $a$  plus the fee in question to  $v_i$ . In this case, a value of  $10 + 3 = 13$  is therefore transferred to  $v_i$ .

At any given time, the LN will typically be a sparse graph with only a small proportion of vertex pairs being connected by an arc. If we wish to transfer an amount from a source vertex  $s$  to a destination vertex  $t$  it may be the case that no arc exists from  $s$  to  $t$ , or that this




Figure 1: An example LN containing two vertices and a single arc is displayed. The arc is labelled with a tuple containing the corresponding base fee, fee rate, and balance respectively.




Figure 2: An example LN containing three vertices and two arcs is displayed. Each arc is labelled with a tuple containing the corresponding base fee, fee rate and balance.

arc has an insufficient balance. On the other hand, it may be possible to transfer the amount along a sequence of arcs  $(v_1, v_2), (v_2, v_3), \dots, (v_{n-1}, v_n)$  where  $s = v_1$  and  $t = v_n$ . Such a sequence is known as a *path* from  $s$  to  $t$ . For example, consider the toy LN snapshot displayed in Figure 2 and the case where we wish to transfer a given amount from  $v_i$  to  $v_j$ . Here, there exists no arc from  $v_i$  to  $v_j$  but the path  $(v_i, v_k), (v_k, v_j)$  may potentially be used to perform the transfer in question.

To transfer an amount  $a$  from vertex  $s$  to vertex  $t$  along a path, one must transfer to each vertex in that path the amount  $a$  plus the sum of fees for subsequent arcs in the path. Consider again the LN displayed in Figure 2 and the case where we wish to transfer an amount  $a$  from  $v_i$  to  $v_j$  along the path  $e_1, e_2$  where  $e_1 = (v_i, v_k)$  and  $e_2 = (v_k, v_j)$ . To successfully transfer the amount  $a$  to  $v_j$ , we must transfer the amount  $a + f_{br}(e_2, a)$  to  $v_k$  and the amount  $a + f_{br}(e_2, a) + f_{br}(e_1, a + f_{br}(e_2, a))$  to  $v_i$ . Furthermore, the balance of the arc  $e_2$  must be greater than or equal to  $a$ , and the balance of the arc  $e_1$  must be greater than or equal to  $a + f_{br}(e_2, a)$ .

Let  $e_1, e_2, \dots, e_n$  be a path from  $s$  to  $t$  and let  $v_1, v_2, \dots, v_{n+1}$  be the sequence of  $n + 1$  vertices in this path. Also, let  $a_i$  be the amount that must be transferred to vertex  $v_i$  such that, ultimately, the amount  $a$  is successfully transferred to  $t$ . The sequence of values  $a_{n+1}, a_n, \dots, a_1$  is defined by the recurrence relation in Equation (2) with the condition  $a_{n+1} = a$  [25, Chapter 7]. In the appendix of this article, we solve this recurrence relation to give a closed-form expression for each  $a_i$  value.

$$a_{i-1} = a_i + f_{br}(e_i, a_i). \quad (2)$$

We refer to the problem of finding a path from  $s$  to  $t$  that minimises the total arc fees while satisfying the necessary balance constraints as the *path planning problem*. Most solutions to the path planning problem use *source-based path planning* whereby the payment source



performs the path planning task. Note that, many implementations of path planning in the LN consider factors relating to both arc fees and arc reputation [26, 27]. However, in this work, we do not consider factors relating to arc reputation because these cannot be determined without access to historical transaction data.

## 4 Variant of Dijkstra’s algorithm

Dijkstra’s algorithm is a well-known path planning algorithm that constructs a lowest cost path from a source vertex  $s$  to a destination vertex  $t$  [28]. It achieves this by maintaining a map  $c : V \rightarrow \mathbb{R}$  that stores an upper bound on the cost from  $s$  to the vertex in question. The algorithm visits vertices in the order of their cost from the vertex  $s$  such that when a vertex is visited the above bound becomes tight. Dijkstra’s algorithm terminates when the vertex  $t$  is visited. When this happens, we obtain a path of lowest cost from  $s$  to  $t$  by backtracking through the exploration steps.

In its basic form, Dijkstra’s algorithm cannot be directly applied to solve the LN path planning problem. This is because, to make a payment of amount  $a$  from  $s$  to  $t$ , it is first necessary to know the amount  $a' \geq a$  that must be transferred initially to  $s$  such that the amount  $a$  is transferred to  $t$  after all fees have been subtracted. Let  $a_v$  be the amount transferred to vertex  $v$  so that  $a$  can, in turn, be transferred to  $t$  using a lowest-fee path from  $v$  to  $t$ . Note that,  $a_t = a$ . If there exists no path from  $v$  to  $t$ , then  $a_v$  is set to  $\infty$ . Given an arc  $(v, v')$ , the value  $a_{v'}$  is related to the value  $a_v$  by the following, which is derived from Equation (2):

$$a_{v'} = a_v - f_{br}((v, v'), a_{v'}). \quad (3)$$

Initially,  $a_v$  is unknown for all  $v \in V$  apart from  $t$ . However, applying Dijkstra’s algorithm to the graph  $G$  to construct a lowest-fee path from  $s$  to  $t$  requires knowledge of  $a_s$ . To illustrate this, consider the LN displayed in Figure 3(a) and the case where we wish to make a payment of amount  $a = 10$  from  $s$  and to  $t$ . The vertex  $s$  has two adjacent arcs  $(s, i)$  and  $(s, j)$  which have equal base fees but different fee rates. The value  $a_t = 10$ ; however, the value  $a_s$  is initially unknown and, therefore, it is initially not possible to compute the values  $a_i$  and  $a_j$  by applying Equation (3) to the arcs  $(s, i)$  and  $(s, j)$  respectively. Using a calculation described later, it turns out that  $a_i = 13$  and  $a_j = 30$ . The higher latter value is a consequence of the high base fee and fee rate of the arc  $(j, t)$  which means that a large amount would need to be transferred to the vertex  $j$ .

This challenge can be overcome by constructing a lowest-fee path in the opposite direction from  $t$  to  $s$  instead of from  $s$  to  $t$ . This approach is motivated by the fact that we initially only know the value  $a_t$ . Given an arc  $(v, v')$ , the value  $a_v$  is related to the value  $a_{v'}$  by the following:

$$a_v = a_{v'} + f_{br}((v, v'), a_{v'}). \quad (4)$$




Figure 3: An example LN is displayed in (a) and the corresponding transpose, where arc directions are reversed, is displayed in (b).

By applying this equation recursively from the vertex  $t$ , we can now compute the value  $a_v$  for all vertices  $v$  where there exists a feasible path from  $v$  to  $t$ . For example, consider again the LN displayed in Figure 3(a). Using the fact that  $a_t = 10$ , we can compute the values  $a_i$  and  $a_j$  by applying Equation (4) to the arcs  $(i, t)$  and  $(j, t)$  respectively. These values in question are  $10 + 2 + 0.1 \times 10 = 13$  and  $10 + 15 + 0.5 \times 10 = 30$  respectively. Applying Equation (4) recursively to the arcs  $(s, i)$  and  $(s, j)$ , the value of  $a_s$  is computed to be the minimum of  $13 + 2 + 0.2 \times 13 = 17.6$  and  $30 + 2 + 0.1 \times 30 = 35$ , which equals 17.6. Since the vertex  $s$  has now been encountered, we have discovered a lowest-fee path from  $s$  to  $t$  which involves transferring a value of 17.6 to vertex  $s$ , next transferring a value of 13 to vertex  $i$  and finally transferring a value of 10 to vertex  $t$ .

The above path planning process corresponds to searching for a lowest-fee path from  $t$  to  $s$  in the transpose of the graph  $G$ , denoted by  $G^T$ . The transpose of a graph  $G$  equals the graph  $G$  in which arc directions have been reversed. Figure 3(b), for example, shows the transpose of the graph from Figure 3(a). In the following theorem, we prove that a lowest-fee path from  $t$  to  $s$  in  $G^T$  corresponds to a lowest-fee path from  $s$  to  $t$  in  $G$ . Hence, the problem of determining a path with the latter property can be reformulated as the problem of determining a path with the former property. In turn, searching for a lowest-fee path from  $t$  to  $s$  in  $G^T$  represents a correct approach.

**Theorem 1.** *Let  $G$  be a graph representation of the LN and let  $G^T$  be its transpose. Let  $(v_1, v_2), (v_2, v_3), \dots, (v_{n-1}, v_n)$  be a lowest-fee path in  $G^T$  from  $t$  to  $s$ , where the amount transferred to  $t$  is  $a$ . The transpose path  $(v_n, v_{n-1}), (v_{n-1}, v_{n-2}), \dots, (v_2, v_1)$  in  $G$  is a lowest-fee path from  $s$  to  $t$  where the amount transferred to  $t$  is  $a$ .*

*Proof.* Let  $(v, v')$  be an arc in  $G^T$  and  $(v', v)$  be the corresponding transposed arc in  $G$ . Also, consider the path  $(v, v')$  in  $G^T$  and the path  $(v', v)$  in  $G$  where both paths contain a single

arc. Consider the fee for the path  $(v, v')$  in  $G^T$  in the case where the amount transferred to  $v$  is  $a$  and fees are added to the amount transferred. This equals the fee for the path  $(v', v)$  in  $G$  where the amount transferred to  $v$  is  $a$  and fees are subtracted from the amount transferred. This result generalises from paths containing a single arc to paths containing multiple arcs and, in turn, proves the theorem.  $\square$

Algorithm 1 provides a pseudocode description of the above path planning process. It uses a map  $c : V \rightarrow \mathbb{R}$  to model the total fee for the current lowest-fee path from  $t$  to each vertex in  $G^T$ . Initially, all vertices are determined to have an infinite fee apart from  $t$ . A priority queue  $Q$  is also used to order vertices by increasing values of  $c(v)$ . Initially, only  $t$  is present in  $Q$ . As shown, the algorithm processes vertices in order of increasing fee by removing the minimum element from  $Q$  at each step (lines 4 and 5). When the selected vertex  $v$  is explored, the algorithm considers each neighbouring vertex  $v'$  of  $v$ . The fee corresponding to the path from  $t$  to  $v'$  is the sum of the lowest-fee path from  $t$  to  $v$  plus the fee of the arc  $(v, v')$ . This is computed at line 9. If this fee is less than the fee of the current lowest-fee path from  $t$  to  $v'$  and the arc  $(v, v')$  has a sufficient balance (line 10), then the maps and priority queue are updated on lines 11 and 12. If  $v'$  is not yet present in  $Q$ , it is added to  $Q$  with a fee of  $c_{v'}$ . The algorithm terminates either when  $Q$  becomes empty (indicating that no path from  $t$  to  $s$  in  $G^T$  exists) or when a lowest-fee path from  $t$  to  $s$  in  $G^T$  is found (line 7). At termination,  $p$  can be used to determine the lowest-fee path from  $t$  to each vertex in  $G^T$ .

Observe that Algorithm 1 is a variant of Dijkstra's original algorithm due to the presence of two features. Firstly, in this case, the computed cost of an arc  $f_{br}((v, v'), a)$  (line 9) is dependent on the amount  $a$  being transferred and is therefore not a constant. In contrast, Dijkstra's original algorithm assumes that all arc costs are non-negative constants. Secondly, in the above algorithm, an arc is only explored if it has a balance greater than or equal to the amount being transferred (line 10). On the other hand, Dijkstra's original algorithm contains no such condition. These differing features mean that existing correctness and computational complexity results for Dijkstra's original algorithm are insufficient here.

## 5 Correctness and Computational Complexity Analysis

In this section, we present proofs of the correctness and computational complexity of the LN path planning process in Algorithm 1. We first introduce the definition of a *consistent* fee map. This is a generalisation of the definition of a consistent travel time map introduced by Kaufman et al. [29].

**Definition 1.** For a given LN graph  $G = (V, E)$ , a fee map  $f : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$  is consistent if for all  $e \in E$ ,  $a \in \mathbb{Z}$  and  $a' \in \mathbb{Z}$  where  $a \leq a'$ , the following condition holds:

$$a + f(e, a) \leq a' + f(e, a'). \quad (5)$$

**Algorithm 1:** Variant of Dijkstra's algorithm for LN

**Input:** A graph  $G^T = (V, E)$  that equals the transpose of the LN graph  $G$ ; a payment source  $s \in V$ ; a payment destination  $t \in V$ ; and a payment amount  $a \in \mathbb{Z}^>$  to be transferred to  $t$ .

**Output:** A map  $c : V \rightarrow \mathbb{R}$  that, for each vertex  $v$ , returns the fee for a lowest-fee path from  $t$  to  $v$  in  $G^T$  where the amount transferred to  $t$  is  $a$ ; a map  $p : V \rightarrow V$  that, for each vertex  $v$ , returns the previous vertex in a lowest-fee path from  $t$  to  $v$  in  $G^T$ .

```

1 For all  $v \in V$ , set  $c(v) = \infty$ 
2 Set  $c(t) = 0$  and insert the ordered pair  $(c(t), t)$  into a priority queue  $Q$ 
3 while  $|Q| > 0$  do
4   Let  $(c(v), v)$  be the element in  $Q$  with the minimum value for  $c(v)$ 
5   Remove the element  $(c(v), v)$  from  $Q$ 
6   if  $v = s$  then
7     | break
8   forall  $v' \in \{u : (v, u) \in E\}$  do
9     |  $c_{v'} = c(v) + f_{br}((v, v'), a)$ 
10    | if  $c_{v'} < c(v') \wedge c(v) + a \leq b(v, v')$  then
11    | | Add the element  $(c_{v'}, v')$  to  $Q$  and, if present, remove element  $(c(v'), v')$ 
12    | | from  $Q$ 
13    | | Set  $c(v') = c_{v'}$  and set  $p(v') = v$ 

```

Broadly speaking, this means that if a smaller amount is transferred along a given arc, the corresponding amount received will also be smaller. In the following theorem, we prove that the fee map  $f_{br}$  defined in Equation (1) and used within the LN is consistent.

**Theorem 2.** *The fee map  $f_{br}$  defined in Equation 1 is a consistent fee map.*

*Proof.* As noted by Kaufman et al. [29], it follows from the definition of a consistent fee map that, a fee map is consistent if and only if its derivative with respect to  $a$  is greater than or equal to  $-1$ . The derivative of  $f_{br}(e, a) = f_b(e) + f_r(e) \times a$  with respect to  $a$  is  $f_r(e)$ . Since this value is constrained to be greater than or equal to 0, the result follows.  $\square$

Consider the fee map  $f_{br}^\infty : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$  defined below. For a given arc  $e$  and payment amount  $a$ , this map returns the value  $f_{br}(e, a)$  if  $e$  has a sufficient balance to transfer the amount  $a$ . Otherwise, this map returns  $\infty$ . This map is related to the concept of a barrier function in the field of optimisation [30]. In the following lemma, we prove that the fee map

$f_{br}^\infty$  is consistent.

$$f_{br}^\infty(e, a) = \begin{cases} f_{br}(e, a) & \text{if } a \leq b(e). \\ \infty & \text{otherwise.} \end{cases} \quad (6)$$

**Lemma 3.** *The fee map  $f_{br}^\infty$  defined in Equation 6 is a consistent fee map.*

*Proof.* To prove that the map  $f_{br}^\infty$  is consistent we must prove that the statement  $a + f_{br}^\infty(e, a) \leq a' + f_{br}^\infty(e, a')$  is true for all  $e \in E$ ,  $a \in \mathbb{Z}$  and  $a' \in \mathbb{Z}$  where  $a \leq a'$ . The space of all variable combinations can be partitioned into two sets:  $a \leq b(e)$  and  $a > b(e)$ . On the first set  $a \leq b(e)$ ,  $f_{br}^\infty(e, a)$  evaluates to  $f_{br}(e, a)$ . In this case, from Theorem 2, the statement  $a + f_{br}^\infty(e, a) \leq a' + f_{br}^\infty(e, a')$  is true. On the second set  $a > b(e)$ ,  $f_{br}^\infty(e, a)$  evaluates to  $\infty$ . In this case, the statement  $a + f_{br}^\infty(e, a) \leq a' + f_{br}^\infty(e, a')$  evaluates to  $\infty \leq \infty$  and is also true.  $\square$

Toward proving the correctness and computational complexity of Algorithm 1, we first show that the problem of determining a lowest-fee path in an LN with fee map  $f_{br}^\infty$  is equivalent to the problem of determining a lowest time path in a time-dependent network (TDN). A TDN is a type of network where the time it takes to traverse a given connection is a function of the time this action is performed [29, 31]. For example, a street network can be modelled as a TDN, since the time it takes an agent to traverse a given street is a function of the time this action is performed. In general, a TDN is modelled as a directed graph, where arcs model transportation connections, plus a travel time map from arc and time pairs to travel times. Consider a TDN  $G = (V, E)$  with travel time map  $f_d : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$ . Now consider the path  $e_1, e_2$  in this TDN where  $e_1 = (v_i, v_k)$ ,  $e_2 = (v_k, v_j)$  and the agent in question departs  $v_i$  at time  $a$ . The agent arrives at  $v_k$  at time  $a + f_d(e_1, a)$ . The agent in turn arrives at  $v_j$  at time  $a + f_d(e_1, a) + f_d(e_2, a + f_d(e_1, a))$ . This computation is analogous to the LN payment computations described earlier in the previous section. We formalise this relationship in the following lemma.

**Lemma 4.** *Consider a LN graph  $G = (V, E)$  with fee map  $f_{br}^\infty : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$  and a TDN graph  $G' = (V', E')$  with time map  $f_d : E' \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$ . The problem of determining a lowest-fee path in  $G$  from  $s$  to  $t$  (where the amount transferred to  $t$  is  $a$ ) is equivalent to the problem of determining a minimum time path in  $G'$  from  $t'$  to  $s'$  (where the time the agent departs  $t'$  is  $a'$ ).*

*Proof.* The problem of determining a lowest-fee path in  $G$  from vertex  $s$  to vertex  $t$  where the amount transferred to  $t$  is  $a$  is equivalent to the problem of determining a lowest-fee path in  $G^T$  from vertex  $t$  to vertex  $s$  where the amount transferred to  $t$  is  $a$  (see Theorem 1). By equating  $G^T$  with  $G'$ ,  $s$  with  $s'$ ,  $t$  with  $t'$ ,  $a$  with  $a'$  and  $f_{br}^\infty$  with  $f_d$ , the path planning problems in the graphs  $G'$  and  $G^T$  are equivalent. In turn, the path planning problems in the graphs  $G$  and  $G'$  are equivalent.  $\square$

The equivalence established in Lemma 4 means that several results for path planning in TDNs generalise to path planning in an LN. This is a useful property, because it allows us to leverage existing research on the topic of TDNs that has been developed over several decades [29, 31]. Toward this goal, we define an algorithm denoted by Algorithm 1'. This algorithm modifies Algorithm 1 by replacing the fee map  $f_{br}$  with  $f_{br}^\infty$  and modifying the condition used to determine if a given arc is explored in the search for a lowest fee path. We subsequently prove the correctness and computational complexity of Algorithm 1' before generalising these results to Algorithm 1.

**Definition 2.** *Replace in Algorithm 1 the fee map  $f_{br}$  with  $f_{br}^\infty$  on line 9 and replace the statement  $c_{v'} < c(v') \wedge c(v) + a \leq b(v, v')$  with  $c_{v'} < c(v')$  on line 10. We denote this modified algorithm as Algorithm 1'.*

**Lemma 5.** *Algorithm 1' for computing a lowest-fee path in the LN graph  $G = (V, E)$  with fee map  $f_{br}^\infty$  is correct.*

*Proof.* Algorithm 1' is a variant of Dijkstra's algorithm. Kaufman et al. [29] prove that using this type of algorithm variant to determine a minimum time path in a TDN with a consistent travel time map is correct (see Theorem 4 in the original paper by Kaufman et al.). Lemma 3 proves that the fee map  $f_{br}^\infty$  is a consistent fee map. Lemma 4 proves that the problem of determining a lowest-fee path in an LN with fee map  $f_{br}^\infty$  is equivalent to the problem of determining a minimum time path in a TDN with travel time map  $f_d$ . Combining the above three results, implies that Algorithm 1' is correct.  $\square$

**Lemma 6.** *The computational complexity of applying the path planning algorithm Algorithm 1' to the LN graph  $G = (V, E)$  with fee map  $f_{br}^\infty$  is  $O(|E| + |V| \log |V|)$ .*

*Proof.* It was proved by Kaufman et al. (see Theorem 4 in [29]) that applying Dijkstra's algorithm to a TDN has equal computational complexity to applying Dijkstra's algorithm to a TDN where each arc has a constant and non-negative travel time. Hence, computational complexity results for applying Dijkstra's algorithm to graphs with constant and non-negative weights generalise to Algorithm 1'. The algorithm is applied to the transpose of the graph  $G$  which contains an equal number of vertices and arcs. Using a Fibonacci heap data structure for the priority queue  $Q$ , an application of Dijkstra's algorithm to this graph has computational complexity  $O(|E| + |V| \log |V|)$ .  $\square$

The above results prove the correctness and computational complexity respectively of Algorithm 1'. In the following, we prove that these results generalise to Algorithm 1. These proofs are based on the insight that Algorithm 1 prunes the search space of Algorithm 1' to only consider feasible payment paths.

**Theorem 7.** *Algorithm 1 for computing a lowest-fee path in the LN graph  $G = (V, E)$  with fee map  $f_{br}$  is correct.*

*Proof.* In the following, we prove that Algorithms 1 and 1' compute equivalent results. The search space of Algorithm 1 is a pruned search space relative to Algorithm 1'. Specifically, the search space of Algorithm 1 is the space of feasible payment paths. The search space of Algorithm 1' is the space of both feasible and infeasible payment paths. We assume that the range of  $f_{br}$  is bounded and therefore all feasible payment paths have a finite fee. On the other hand, all infeasible payment paths will contain an arc where  $f_{br}^\infty$  evaluates to  $\infty$  and therefore have an infinite fee. If a feasible lowest fee path exists, since Algorithm 1' is correct from Lemma 5, Algorithm 1' will return it. On the other hand, since it is feasible, this payment path will exist in the pruned search space and will also be returned by Algorithm 1. If no feasible payment path exists, Algorithm 1' will return an infeasible payment path with an infinite fee. On the other hand, Algorithm 1 will not return a payment path. If we equate returning an infeasible payment path with not returning a payment path, then Algorithms 1 and 1' compute equivalent results. Therefore, the correctness of Algorithm 1' implies the correctness of Algorithm 1.  $\square$

**Theorem 8.** *The computational complexity of applying the path planning algorithm presented in Algorithm 1 to the LN graph  $G = (V, E)$  with fee map  $f_{br}$  is  $O(|E| + |V| \log |V|)$ .*

*Proof.* Since Algorithm 1 searches a pruned search space relative to Algorithm 1', it will have computational complexity less than or equal to Algorithm 1'. We assume that payment amounts and the range of  $f_{br}$  are both bounded. Therefore, the amount that one may wish to forward along an arc is bounded. In a worst case where each arc has a balance greater than or equal to this bound, the search spaces of Algorithm 1 and Algorithm 1' are equal. Hence Algorithm 1 and Algorithm 1' have equal computational complexity. Given this, the proof follows from Lemma 6.  $\square$

The above proofs of correctness and computational complexity for Algorithm 1 assume that the fee map in question is consistent. If this were not the case, the algorithm would not be correct. To demonstrate this, consider the LN graph  $G$  and corresponding transpose graph  $G^T$  displayed in Figures 4(a) and 4(b) respectively. Furthermore, let the fee map  $f_{br}$  corresponding to the transpose graph  $G^T$  have the properties  $f_{br}((t, i), 100) = 10$ ,  $f_{br}((i, j), 110) = 10$ ,  $f_{br}((j, s), 120) = 5$ ,  $f_{br}((t, j), 100) = 10$  and  $f_{br}((j, s), 110) = 20$ . This fee map is not consistent because the following inequality is not satisfied  $110 + f_{br}((j, s), 110) = 110 + 20 = 130 \leq 120 + f_{br}((j, s), 120) = 120 + 5 = 125$ .

Consider the problem of making a payment from  $s$  to  $t$  in  $G$  where the amount transferred to  $t$  is 100. Applying Algorithm 1 to  $G^T$  to determine a path from  $t$  to  $s$  returns the path  $(t, j), (j, s)$  which has a fee of  $f_{br}((t, j), 100) + f_{br}((j, s), 110) = 10 + 20 = 30$ . However,




Figure 4: An LN graph  $G$  and corresponding transpose graph  $G^T$  displayed in (a) and (b) respectively. The fee map  $f_{br}$  corresponding to  $G^T$  has the properties  $f_{br}((t, i), 100) = 10$ ,  $f_{br}((i, j), 110) = 10$ ,  $f_{br}((j, s), 120) = 5$ ,  $f_{br}((t, j), 100) = 10$  and  $f_{br}((j, s), 110) = 20$ .

the lowest fee path from  $t$  to  $s$  is in fact  $(t, i), (i, j), (j, s)$ , which has a fee of  $f_{br}((t, i), 100) + f_{br}((i, j), 110) + f_{br}((j, s), 120) = 10 + 10 + 5 = 25$ . This is because when the algorithm processes the arc  $(t, j)$  it will insert the element  $(110, j)$  into the queue  $Q$  and consider the path  $(t, j)$  as the lowest fee path from  $t$  to  $j$ . When the algorithm later processes the arc  $(i, j)$  it will observe that  $120 \not\leq 110$  and, consequently will not insert the element  $(120, i)$  into the queue  $Q$  nor consider the path  $(t, i), (i, j)$  as the lowest fee path from  $t$  to  $j$ . Hence, Algorithm 1 does not return the lowest fee path and is therefore incorrect in this case.

If the fee map for a LN graph is an arbitrary map that is not consistent, the corresponding path planning can be NP-hard. We present a proof of this fact in the following theorem.

**Theorem 9.** *Consider a LN graph  $G = (V, E)$  with fee map  $f : E \times \mathbb{Z}^{\geq} \rightarrow \mathbb{Z}^{\geq}$ . If  $f$  is an arbitrary map that is not consistent, the problem of determining a lowest-fee path in  $G$  from  $s$  to  $t$ , where the amount transferred to  $t$  is  $a$ , can be NP-hard.*

*Proof.* In Lemma 4 we proved an equivalence between the path planning problem in question and a specific type of path planning problem in a TDN. The latter problem is denoted *forbidden waiting* by Orda et al. [32], and Zeitz [33] proved that this problem is NP-hard. This implies the path planning problem in question is also NP-hard.  $\square$



## 6 Bidirectional Variant of Dijkstra's algorithm

The path planning algorithm described in the previous section can be considered *unidirectional* in that it searches for a path in a single direction from the destination vertex to the source vertex. On the other hand, bidirectional path planning algorithms search for a path in two directions simultaneously. The first search takes place from the destination vertex towards the source vertex, the second search takes place from the source vertex towards the destination vertex. A path is then established when both searches intersect [34]. If applicable to the LN, bidirectional path planning has the potential to reduce the time and space requirements relative to unidirectional path planning. This has been demonstrated empirically in many types of real-world networks including street networks [35].

To illustrate how bidirectional path planning can reduce the computational effort needed for path planning in the LN, consider the example LN displayed in Figure 5 which has a hub-and-spoke topology. Although this is not an accurate representation of the real LN, the LN is known to exhibit topologies with hub-and-spoke structures [20]. In this context, vertices on the spokes likely correspond to customers and merchants, while vertices in the hub correspond to routing vertices. In this example, consider the case where we wish to transfer an amount  $a$  from vertex  $s$  to vertex  $t$  and let us assume that all arcs have sufficient balance to transfer this amount plus any necessary fees. There is a single path  $(s, r), (r, t)$  from  $s$  to  $t$  which contains a single intermediate vertex  $r$ . If we apply the unidirectional path planning algorithm described in the previous section, this algorithm will explore every vertex adjacent to  $r$  before finding the lowest-fee path. On the other hand, if we could apply a bidirectional path planning algorithm, this algorithm would only need to explore the three vertices  $s, r$  and  $t$  to find the lowest-fee path. This is because both searches intersect when they explore the vertex  $r$  and therefore do not need to explore any other vertices on the spokes.

The asymptotic computational complexity of both the unidirectional and bidirectional path planning algorithms are both  $O(|E| + |V| \log |V|)$ . The former corresponds to applying Dijkstra's algorithm once, while the latter corresponds to applying Dijkstra's algorithm twice. However, the empirical benefits of the bidirectional algorithm can be explained by considering the branching factor  $b$  of the searches and the number of arcs  $d$  in the lowest-fee path. In graph searching algorithms such as Dijkstra's algorithm, the branching factor equals the average number of branching paths the search can explore from a given vertex [34, Chapter 3]. This approximately equals the average vertex degree. If a path planning algorithm is applied, the number of vertices explored will be approximately  $O(b^d)$ . This follows from the fact that the search will explore a search tree to a depth  $d$  before the lowest-fee path is found and at each step, the search tree will expand by a factor of  $b$ . On the other hand, if a bidirectional path planning algorithm is applied, the number of vertices explored will be approximately  $2 \times O(b^{d/2})$  which is less than  $O(b^d)$  [34, Chapter 3]. Here, the multiple 2 represents the fact that two searches are applied. The  $d/2$  exponent represents the fact that each of these




Figure 5: An example LN with a hub-and-spoke topology is displayed.

searches will explore a search tree to a depth of  $d/2$  before the lowest-fee path is found. The above analysis implies that, if bidirectional path planning could be successfully applied to the LN, this would reduce the number of vertices explored. As discussed in the previous section, applying a path planning algorithm from the destination vertex  $t$  to the source vertex  $s$  is possible using Algorithm 1. However, applying a path planning algorithm in the opposite direction is challenging.

To partially overcome this challenge, we can use the fact that fees charged by outgoing arcs adjacent to the vertex  $s$  are paid to  $s$ . That is, any fees paid by the vertex  $s$  will automatically be refunded and therefore can be considered to be zero. This fact can be modelled as a transformation of the LN such that  $f_b(s, v) = 0$  and  $f_r(s, v) = 0.0$  for all  $v \in \{u : (s, u) \in E\}$ . For example, applying this transformation to the LN in Figure 3(a) gives the LN displayed in Figure 6. Note that the transformation does not alter arc balances. As a consequence of the fact that the LN does not support negative fees, the value 0.0 is a tight lower bound for the fees charged for transferring any amount. Therefore, when a vertex  $v$  in the set  $\{u : (s, u) \in E\}$  has been explored in a unidirectional search from  $t$  and the arc  $(s, v)$  has sufficient balance, a lowest cost path from  $t$  to  $s$  can be inferred by appending the arc  $(s, v)$  to the lowest-fee path from  $v$  to  $t$ . In this case, the search can be terminated before exploring the vertex  $s$ . For example, consider again the LN in Figure 3(a). When the vertex  $i$  has been explored in a unidirectional search from  $t$ , a lowest cost path from  $t$  to  $s$  can be inferred by appending the arc  $(s, i)$  to the lowest cost path from  $i$  to  $t$ .

Using this insight, we can extend the termination condition used in the unidirectional path planning method shown in Algorithm 1. Specifically, instead of just terminating when the




Figure 6: An example LN is displayed.

next vertex  $v$  to be explored is the source  $s$  (line 6), we can also terminate when  $v$  belongs to the set  $\{u : (u, s) \in E\}$  and the arc in question has a sufficient balance (line 8). Algorithm 2 provides a full description of this new algorithm. This is the same as Algorithm 1 except for the additional termination condition at line 8. We now prove the correctness of this algorithm.

**Theorem 10.** *Algorithm 2 for computing a lowest-fee path in the LN is correct.*

*Proof.* It follows from Theorem 5 that the algorithm computes a lowest-fee path from the vertex  $t$  to the vertex  $s$  or a vertex in the set  $v \in \{u : (u, s) \in E\}$ . When either of these events occurs, the algorithm, in turn, computes a lowest-fee path from  $t$  to  $s$  in  $G^T$ . Hence, Algorithm 2 is correct.  $\square$

Algorithm 2 is a variant on bidirectional path planning in which a full search is initiated from  $t$  and only a partial search, consisting of a single step, is initiated from  $s$ . Hence, we refer to this path planning algorithm as a *partial bidirectional path planning* algorithm. As discussed earlier, a bidirectional path planning algorithm explores fewer vertices than a unidirectional path planning algorithm. Since the partial bidirectional path planning terminates the search one step earlier, it will explore approximately  $O(b^{d-1})$  where  $b$  is the branching factor and  $d$  is the number of arcs in the lowest-fee path. This is less than that of the unidirectional path planning algorithm equals  $O(b^d)$  but greater than that of the bidirectional path planning algorithm, which equals  $2O(b^{d/2})$ . The number of vertices explored by the partial bidirectional algorithm is not significantly less than that for unidirectional path planning. However, as discussed in the results section below, we find in practice that partial bidirectional path planning exhibits better performance. This can be attributed to the hub and spoke topology of the LN described earlier.

---

**Algorithm 2:** Bidirectional Variant of Dijkstra's algorithm
 

---

**Input:** A graph  $G^T = (V, E)$  that equals the transpose of the LN graph  $G$ ; a payment source  $s \in V$ ; a payment destination  $t \in V$ ; and a payment amount  $a \in \mathbb{Z}^>$  to be transferred to  $t$ .

**Output:** A map  $c : V \rightarrow \mathbb{R}$  that, for each vertex  $v$ , returns the fee for a lowest-fee path from  $t$  to  $v$  in  $G^T$  where the amount transferred to  $t$  is  $a$ ; a map  $p : V \rightarrow V$  that, for each vertex  $v$ , returns the previous vertex in a lowest-fee path from  $t$  to  $v$  in  $G^T$ .

```

1 For all  $v \in V$ , set  $c(v) = \infty$ 
2 Set  $c(t) = 0$  and insert the ordered pair  $(c(t), t)$  into a priority queue  $Q$ 
3 while  $|Q| > 0$  do
4   | Let  $(c(v), v)$  be the element in  $Q$  with the minimum value for  $c(v)$ 
5   | Remove the element  $(c(v), v)$  from  $Q$ 
6   | if  $v = s$  then
7   |   | break
8   | if  $v \in \{u : (u, s) \in E\} \wedge c(v) + a \leq b(v, s)$  then
9   |   | Set  $c(s) = c(v)$ , set  $p(s) = v$ , and break
10  | forall  $v' \in \{u : (v, u) \in E\}$  do
11  |   |  $c_{v'} = c(v) + f_{br}((v, v'), a)$ 
12  |   | if  $c_{v'} < c(v') \wedge c(v) + a \leq b(v, v')$  then
13  |   |   | Add the element  $(c_{v'}, v')$  to  $Q$  and, if present, remove element  $(c(v'), v')$ 
14  |   |   | from  $Q$ 
14  |   |   | Set  $c(v') = c_{v'}$  and set  $p(v') = v$ 
    
```

---

## 7 Experimental Results

In this section, we present an empirical evaluation of the unidirectional and partial bidirectional path planning algorithms described in Sections 4 and 6 respectively. As discussed in the Related Works section, existing path planning algorithms can broadly be categorised as centralised algorithms and decentralised algorithms. The unidirectional algorithm is the most commonly used centralised algorithm and therefore is a suitable benchmark to compare the partial bidirectional algorithm against. On the other hand, decentralised algorithms are not commonly used due to the lack of payment privacy they provide. In turn, implementations of these methods are difficult to obtain. Therefore, we do not consider any decentralised algorithms in our comparison.

In our trials, we used a real snapshot of the LN captured on 22 February 2024 using

the Lightning Network Daemon (Ind) implementation of an LN node.<sup>1</sup> On this same date, one Satoshi (the unit of currency used by the LN) had a value of approximately 0.0004 GBP. The graph representation of the above snapshot contains 9,420 vertices and 32,202 arcs; however, many of these arcs had no fee information making them unusable for routing and were therefore removed. Furthermore, many vertices were isolated with no adjacent arcs and these were also removed. Following these steps, the resulting graph contained 2,453 vertices and 26,000 arcs. The mean and median LN snapshot channel capacities were 10,490,032 Satoshis and 4,000,000 Satoshis respectively. The 10th and 90th percentiles of LN snapshot channel capacities were 500,000 Satoshis and 17,137,075 Satoshis respectively. To help maintain privacy and reduce communication bandwidth, individual arc balances are not shared. We therefore assumed that all channels were balanced. That is, to each of the two arcs corresponding to a given channel, we assigned a balance equal to half of the corresponding channel capacity. Currently, there is no historical transaction data for the LN and such data is difficult to infer by design. Therefore, to test the two path planning algorithms, we constructed three sets of simulated payments on the LN snapshot described above. The first set of payments contained 10,000 elements and was constructed using the following rejection sampling-based approach. We first sampled without replacement the payment source and destination vertices from the set of graph vertices  $V$ . This ensures that the pair of vertices is distinct. Next, we sampled the payment amount uniformly at random from the set  $\{1, 2, \dots, 1000000\}$ . A recent report by the cryptocurrency company River indicates that most LN payment amounts lie in this interval [36]. Given the above payment source, destination, and amount, we then determined if the payment in question was feasible; that is, whether there existed a path from the source to the destination that could transfer the amount. If the payment was feasible, it was added to the set of payments. Otherwise, it was discarded. This process was repeated until the set of payments contained 10,000 elements.

For each simulated payment, we computed the number of vertices explored by the unidirectional and partial bidirectional algorithms by counting the number of corresponding for-loop iterations. (The loops in question start at lines 8 and 10 of Algorithms 1 and 2 respectively). For the unidirectional algorithm, the mean and standard deviation of the number of vertices explored was  $7,452 \pm 3,190$ . For the partial bidirectional algorithm, the corresponding figures were  $4,069 \pm 3,594$ , equating to a 45% reduction in the mean. Next, for each payment, we computed the percentage reduction in vertices explored by the partial bidirectional algorithm. That is, the number explored by unidirectional minus the number explored bidirectional divided by the number explored by unidirectional. This percentage represents the reduction in computation provided by the proposed bidirectional algorithm for the path planning task in question. Figure 7(a) displays a histogram of these percentage values. The mean and standard deviation of these percentages are 47% and 35% respectively.

---

<sup>1</sup><https://docs.lightning.engineering/>




Figure 7: For two sets of simulated payments, the corresponding histograms of the percentage reduction in the number of vertices explored by the partial bidirectional algorithm relative to the unidirectional algorithm are displayed. The vertical axis is the number of payments, while the horizontal axis gives the percentage reduction.

Our second set of simulated payments also contained 10,000 elements. This set was constructed using an approach similar to the previous though, in this case, we sampled the source and destination vertices from the set of all graph vertices with out-degrees of less than four. The motivation for considering these low-degree vertices is that they tend to correspond to merchants and customers instead of routing peers. Similarly to before, we then computed the percentage reduction in the number of vertices explored by the partial bidirectional algorithm. In this case, the mean number of vertices explored by the unidirectional algorithm was  $8,102 \pm 2,882$  whereas, for the partial bidirectional algorithm, these figures were  $5,491 \pm 3,574$ , giving a 32% reduction in the mean. For each payment, we then computed the percentage reduction in vertices explored by the partial bidirectional algorithm. Figure 7(b) displays a histogram of these values. The mean and standard deviation of these percentage values are 33% and 32% respectively.

Our final set of simulated payments contained 100,000 elements. This set was constructed using the same approach as the first set. We then used this set of payments to measure the wall-clock running time of both path planning algorithms. Both algorithms were implemented as single-thread computer programs using the Rust programming language and executed on a PC with an Inter i9-14900K CPU and 32GB of RAM. We computed the set of payments using each algorithm and measured the corresponding wall-clock running times. We then repeated this step ten times and computed the corresponding mean wall-clock running times. The mean wall-clock running times for the unidirectional and partial bidirectional path planning algorithms were 66 and 48 seconds respectively. This equates to a mean wall-clock running

time for an individual payment of 0.00066 and 0.00048 seconds respectively. The latter algorithm therefore provides a 27% reduction in mean wall-clock running time.

In summary, for all three sets of simulated payments considered, the partial bidirectional path planning algorithm explores significantly fewer vertices than the unidirectional path planning algorithm and, in turn, provides improved computational performance.

## 8 Conclusions

In this article, we have considered the problem of path planning in PCNs. We have provided correctness and computational complexity results for a variant of Dijkstra's algorithm, which is the most cited algorithm for solving this problem. These results show that if the PCN fee function has a consistency property, the path planning algorithm is correct and has polynomial computational complexity. This finding therefore has the potential to inform the development of new or existing PCN fee functions. In this article, we have also proposed a slight modification to the above algorithm based on the concept of bidirectional search. We demonstrate empirically that this proposed modification provides improved computational performance.

These findings notwithstanding, solving the path planning problem is only one of many challenges that must be overcome before PCNs can fully deliver on their potential to solve the scaling challenges of cryptocurrencies [37]. Other challenges include mitigating security attacks such as denial-of-service [38] and griefing [39] attacks, ensuring the privacy of users [40], ensuring network topologies that allow efficient payments [41], and allowing users with edge devices such as mobile phones to make and receive payments. In future work, we hope to help address some of these other challenges.

## Acknowledgements

This work was supported by the Security, Crime and Intelligence Innovation Institute at Cardiff University through their Kickstarter Funding scheme.

## References

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, page 21260, 2008.
- [2] Yang Xiao, Ning Zhang, Wenjing Lou, and Y Thomas Hou. A survey of distributed consensus protocols for blockchain networks. *IEEE Communications Surveys & Tutorials*, 22(2):1432–1465, 2020.

- [3] Abdelatif Hafid, Abdelhakim Senhaji Hafid, and Mustapha Samih. Scaling blockchains: A comprehensive survey. *IEEE Access*, 8:125244–125262, 2020.
- [4] Andreas M Antonopoulos, Olaoluwa Osuntokun, and René Pickhardt. *Mastering the Lightning Network*. O’Reilly Media, Inc., 2021.
- [5] Christian Decker, Rusty Russell, and Olaoluwa Osuntokun. Eltoo: A simple layer2 protocol for bitcoin. *White paper: <https://blockstream.com/eltoo.pdf>*, 2018.
- [6] Gabriel Antonio F Rebello, Gustavo F Camilo, Lucas Airam C de Souza, Maria Potop-Butucaru, Marcelo Dias de Amorim, Miguel Elias M Campista, and Luís Henrique MK Costa. A survey on blockchain scalability: From hardware to layer-two protocols. *IEEE Communications Surveys & Tutorials*, 2024.
- [7] Vibhaalakshmi Sivaraman, Shaileshh Bojja Venkatakrishnan, Kathleen Ruan, Parimarjan Negi, Lei Yang, Radhika Mittal, Giulia Fanti, and Mohammad Alizadeh. High throughput cryptocurrency routing in payment channel networks. In *USENIX Symposium on Networked Systems Design and Implementation*, 2020.
- [8] Vivek Bagaria, Joachim Neu, and David Tse. Boomerang: Redundancy improves latency and throughput in payment-channel networks. In *Financial Cryptography and Data Security, Kota Kinabalu, Malaysia*, pages 304–324. Springer, 2020.
- [9] Rene Pickhardt and Stefan Richter. Optimally reliable & cheap payment flows on the lightning network. *arXiv preprint arXiv:2107.05322*, 2021.
- [10] Stefanie Roos, Pedro Moreno-Sanchez, Aniket Kate, and Ian Goldberg. Settling payments fast and private: Efficient decentralized routing for path-based transactions. *arXiv preprint arXiv:1709.05748*, 2017.
- [11] Pavel Prihodko, Slava Zhigulin, Mykola Sahno, Aleksei Ostrovskiy, and Olaoluwa Osuntokun. Flare: An approach to routing in lightning network. *White Paper*, 144, 2016.
- [12] Changting Lin, Ning Ma, Xun Wang, and Jianhai Chen. Rapido: Scaling blockchain with multi-path payment channels. *Neurocomputing*, 406:322–332, 2020.
- [13] Rene Pickhardt and Mariusz Nowostawski. Imbalance measure and proactive channel rebalancing algorithm for the lightning network. In *IEEE International Conference on Blockchain and Cryptocurrency*, pages 1–5, 2020.
- [14] Weizhao Tang, Weina Wang, Giulia Fanti, and Sewoong Oh. Privacy-utility tradeoffs in routing cryptocurrency over payment channel networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(2):1–39, 2020.
- [15] Maya Dotan, Saar Tochner, Aviv Zohar, and Yossi Gilad. Twilight: A differentially private payment channel network. In *USENIX Security Symposium*, pages 555–570, 2022.
- [16] Giovanni Di Stasi, Stefano Avallone, Roberto Canonico, and Giorgio Ventre. Routing payments on the lightning network. In *IEEE international conference on internet of things and IEEE green computing and communications and IEEE cyber, physical and social computing and IEEE smart data*, pages 1161–1170, 2018.
- [17] Oğuzhan Ersoy, Stefanie Roos, and Zekeriya Erkin. How to profit from payments channels. In *Financial Cryptography and Data Security*, pages 284–303. Springer, 2020.
- [18] Yuup Van Engelshoven and Stefanie Roos. The merchant: Avoiding payment channel depletion



- through incentives. In *IEEE International Conference on Decentralized Applications and Infrastructures*, pages 59–68, 2021.
- [19] Saar Tochner, Stefan Schmid, and Aviv Zohar. Hijacking routes in payment channel networks: A predictability tradeoff. *arXiv preprint arXiv:1909.06890*, 2019.
- [20] Stefano Martinazzi and Andrea Flori. The evolving topology of the lightning network: Centralization, efficiency, robustness, synchronization, and anonymity. *PIOS ONE*, 15(1):e0225966, 2020.
- [21] István András Seres, László Gulyás, Dániel A Nagy, and Péter Burcsi. Topological analysis of bitcoin’s lightning network. In *International Conference on Mathematical Research for Blockchain Economy, Santorini, Greece*, pages 1–12, 2020.
- [22] Elias Rohrer, Julian Malliaris, and Florian Tschorsch. Discharged payment channels: Quantifying the lightning network’s resilience to topology-based attacks. In *IEEE European symposium on security and privacy*, pages 347–356, 2019.
- [23] Philipp Zabka, Klaus-T Förster, Christian Decker, and Stefan Schmid. A centrality analysis of the lightning network. *Telecommunications Policy*, 48(2):102696, 2024.
- [24] Arad Kotzer and Ori Rottenstreich. Braess paradox in layer-2 blockchain payment networks. In *IEEE International Conference on Blockchain and Cryptocurrency, Dubai, United Arab Emirates*, 2023.
- [25] Charalambos A Charalambides. *Enumerative combinatorics*. CRC Press, 2002.
- [26] Satwik Prabhu Kumble, Dick Epema, and Stefanie Roos. How lightning’s routing diminishes its anonymity. In *International Conference on Availability, Reliability and Security*, pages 1–10, 2021.
- [27] Dan Andreescu. Comparing lightning routing protocols to routing protocols with splitting. *bachelor thesis*, 2021.
- [28] Peter Sanders, Kurt Mehlhorn, Martin Dietzfelbinger, and Roman Dementiev. *Sequential and Parallel Algorithms and Data Structures*. Springer, 2019.
- [29] David E Kaufman and Robert L Smith. Fastest paths in time-dependent networks for intelligent vehicle-highway systems application. *Journal of Intelligent Transportation Systems*, 1(1):1–11, 1993.
- [30] Jorge Nocedal and Stephen J Wright. Nonlinear equations. *Numerical Optimization*, pages 270–302, 2006.
- [31] Michel Gendreau, Gianpaolo Ghiani, and Emanuela Guerriero. Time-dependent routing problems: A review. *Computers & operations research*, 64:189–197, 2015.
- [32] Ariel Orda and Raphael Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *Journal of the ACM*, 37(3):607–625, 1990.
- [33] Tim Zeitz. NP-hardness of shortest path problems in networks with non-fifo time-dependent travel times. *Information Processing Letters*, 179:106287, 2023.
- [34] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: a modern approach*. Pearson, 4 edition, 2021.
- [35] Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F Werneck. Route planning in transportation

- networks. *Algorithm engineering: Selected results and surveys*, pages 19–80, 2016.
- [36] River. Research report: The lightning network grew by 1212% in 2 years, why it’s important to pay attention. Technical report, River, 2023.
- [37] Maya Dotan, Yvonne-Anne Pignolet, Stefan Schmid, Saar Tochner, and Aviv Zohar. Survey on blockchain networking: Context, state-of-the-art, challenges. *ACM Computing Surveys*, 54(5):1–34, 2021.
- [38] Clara Shikhelman and Sergei Tikhomirov. Unjamming lightning: A systematic approach. *Cryptography ePrint Archive*, 2022.
- [39] Subhra Mazumdar, Prabal Banerjee, Abhinandan Sinha, Sushmita Ruj, and Bimal Kumar Roy. Strategic analysis of griefing attack in lightning network. *IEEE Transactions on Network and Service Management*, 2022.
- [40] George Kappos, Haaron Yousaf, Ania Piotrowska, Sanket Kanjalkar, Sergi Delgado-Segura, Andrew Miller, and Sarah Meiklejohn. An empirical analysis of privacy in the lightning network. In *Financial Cryptography and Data Security*, pages 167–186. Springer, 2021.
- [41] Paolo Guasoni, Gur Huberman, and Clara Shikhelman. Lightning network economics: Topology. Available at SSRN 4439190, 2023.

## A Recurrence Relation

In the following theorem, we solve the recurrence relation defined in Equation (2) to give a closed-form expression for each  $a_i$  value.

**Theorem 11.** *The solution of the recurrence relation in Equation (2) with the condition  $a_{n+1} = a$  is*

$$a_i = \left( \prod_{k=i+1}^{n+1} (1 + f_r(e_k)) \right) \left( a_{n+1} + \sum_{m=i+1}^{n+1} \frac{f_b(e_m)}{\prod_{k=m}^{n+1} (1 + f_r(e_k))} \right). \quad (7)$$

*Proof.* Let  $e'_1, e'_2, \dots, e'_n$  be the sequence of arcs where  $e'_i = e_{n+1-i}$ , and let  $a'_1, a'_2, \dots, a'_{n+1}$  be the sequence of values where  $a'_i = a_{n+2-i}$ . That is, we reverse the order of both sequences. By performing this change of variables, Equation (2) can be rewritten as the following recurrence relation with the initial condition  $a'_1 = a$ :

$$a'_{i+1} = a'_i + f_b(e'_i) + f_r(e'_i)a'_i. \quad (8)$$

The recurrence relation in Equation (8) can be refactored as

$$a'_{i+1} = a'_i + f_b(e'_i) + f_r(e'_i)a'_i \quad (9a)$$

$$= (1 + f_r(e'_i))a'_i + f_b(e'_i). \quad (9b)$$

The recurrence relation in Equation (9b) has the properties of being linear, first-order, non-homogeneous and having variable coefficients [25]. First-order refers to the fact that the term  $a'_{i+1}$  is a function of only the previous term  $a'_i$  in the sequence. Non-homogeneous means that the term  $f_b(e'_i)$  is not a multiple of  $a'_{i-1}$  or  $a'_i$ . Finally, having variable coefficients refers to the fact that the coefficients  $f_b(e'_i)$  and  $(1 + f_r(e'_i))$  are a function of  $i$ . A direct application of Theorem 7.1 in [25] solves this recurrence relation, where the solution is defined as

$$a'_i = \left( \prod_{k=1}^{i-1} (1 + f_r(e'_k)) \right) \left( a'_1 + \sum_{m=1}^{i-1} \frac{f_b(e'_m)}{\prod_{k=1}^m (1 + f_r(e'_k))} \right). \quad (10)$$

Performing an inverse of the change of variables previously performed gives the following solution with the initial condition  $a_{n+1} = a$ :

$$a_{n+2-i} = \left( \prod_{k=n+1}^{n+1-i} (1 + f_r(e_k)) \right) \left( a_{n+1} + \sum_{m=n+1}^{n+1-i} \frac{f_b(e_m)}{\prod_{k=n+1}^m (1 + f_r(e_k))} \right) \quad (11)$$

implying that

$$a_i = \left( \prod_{k=n+1}^{i+1} (1 + f_r(e_k)) \right) \left( a_{n+1} + \sum_{m=n+1}^{i+1} \frac{f_b(e_m)}{\prod_{k=n+1}^m (1 + f_r(e_k))} \right) \quad (12)$$

$$= \left( \prod_{k=i+1}^{n+1} (1 + f_r(e_k)) \right) \left( a_{n+1} + \sum_{m=i+1}^{n+1} \frac{f_b(e_m)}{\prod_{k=m}^{n+1} (1 + f_r(e_k))} \right) \quad (13)$$

as required.  $\square$

The following Python code listing implements the above method for directly solving the recurrence relation in question. By comparing solutions computed using the recursive and direct methods defined in Equations (2) and (7) respectively, this code can be used to empirically verify the correctness of the above theorem.

```

1 import numpy
2
3 a = 102 #Amount a
4 b = [10, 5, 3.4, 11, 7] # path arc base fees
5 r = [0.1, 0.211, 0.15, 0.12, 0.11] # path arc fee rates
6 n = len(b)-1
7
8 #Compute the sequence of values recursively using Equation (2)
9 ai = a
10 print("a[" , n+1, "]: " , ai)
11 for i in range(len(b)-1, -1, -1):

```

```
12     ai += b[i] + ai*r[i]
13     print("a[" , i, "]: " , ai)
14 print()
15
16 #Compute the sequence of values directly using Equation (6)
17 for i in range(n+1, -1, -1):
18     ai = numpy.prod([(1+r[k]) for k in range(i,n+1)]) * (a + numpy.sum
19         [(b[m]/numpy.prod([(1+r[k]) for k in range(m,n+1)]) for m in
20             range(i,n+1)]))
21     print("a[" , i, "]: " , ai)
```

Listing 1: Solving the recurrence relation recursively and directly.