

# Settling Payments Fast and Private: Efficient Decentralized Routing for Path-Based Transactions

Stefanie Roos  
University of Waterloo  
sroos@uwaterloo.ca

Pedro Moreno-Sanchez  
Purdue University  
pmorenos@purdue.edu

Aniket Kate  
Purdue University  
aniket@purdue.edu

Ian Goldberg  
University of Waterloo  
iang@cs.uwaterloo.ca

**Abstract**—Path-based transaction (PBT) networks, which settle payments from one user to another via a path of intermediaries, are a growing area of research. They overcome the scalability and privacy issues in cryptocurrencies like Bitcoin and Ethereum by replacing expensive and slow on-chain blockchain operations with inexpensive and fast off-chain transfers. In the form of credit networks such as Ripple and Stellar, they also enable low-price real-time gross settlements across different currencies. For example, SilentWhispers is a recently proposed fully distributed credit network relying on path-based transactions for secure and in particular private payments without a public ledger.

At the core of a decentralized PBT network is a routing algorithm that discovers transaction paths between payer and payee. During the last year, a number of routing algorithms have been proposed, including SilentWhispers' landmark routing and the link state protocol Flare for Bitcoin's Lightning Network. However, the existing ad hoc efforts lack either efficiency or privacy, and a comprehensive analysis is needed to ensure the success of PBT networks in practice.

In this work, we first identify several efficiency concerns in SilentWhispers. Armed with this knowledge, we design and evaluate SpeedyMurmurs, a novel routing algorithm for decentralized PBT networks using efficient and flexible embedding-based path discovery and on-demand efficient stabilization to handle the dynamics of a PBT network. Our simulation study, based on real-world data from the currently deployed Ripple credit network, indicates that SpeedyMurmurs reduces the overhead of stabilization by up to two orders of magnitude and the overhead of routing a transaction by more than a factor of two. Furthermore, using SpeedyMurmurs maintains at least the same success ratio as decentralized landmark routing, while providing lower delays. Finally, SpeedyMurmurs achieves key privacy goals for routing in PBT networks.

## I. INTRODUCTION

Since the advent of Bitcoin [23], many other blockchain-based payment systems have been proposed and deployed in practice to serve a multitude of purposes. For instance, credit networks [7], [11] such as Ripple [3], [5] or Stellar [4] leverage blockchain technology to enable real-time gross settlement [2] between two end users across different currencies and assets significantly cheaper than the current central banking system. Ethereum [10] builds on top of a blockchain to construct a platform to run fully expressive smart contracts.

However, the growing base of users and transactions is resulting in blockchain scalability issues [6], [25]. Moreover, the public nature of the blockchain leads to demonstrable privacy breaches of sensitive data such as the identities of the transaction partners and the transaction value [13], [17], [18],

[22], [28]. Academic and industry efforts are leading towards peer-to-peer (P2P) path-based transaction (PBT) networks such as the Lightning Network [25] for Bitcoin, the Raiden Network [1] for Ethereum, SilentWhispers [15] for credit networks, or InterLedger [33] and Atomic-swap [14] for inter-blockchain transactions; these PBT networks are promising for addressing scalability, efficiency, and interoperability concerns with blockchains and cryptocurrencies through off-chain transactions requiring no expensive mining efforts. In fact, at a recent blockchain event, the InterLedger team demonstrated a transaction through seven ledgers including Bitcoin, Ethereum, and Ripple [36].

In a PBT network, two users  $u$  and  $v$  locally maintain a weighted link between them (also called a payment channel, state channel, or credit link, depending on the application). The link's weight characterizes the amount of funds (or assets) that one user can transfer to the other, the exact nature of the link depending on the application. For instance, in a credit network, the weight defines the difference between the amount of credit  $u$  is willing to grant  $v$  and the amount  $v$  already owes  $u$ . However, at their core, all PBT networks realize the transaction of funds between a payer and a payee by transferring funds along one or multiple paths between them.

A PBT network builds on top of three key algorithms: *routing*, *payment* and *accountability*. The routing algorithm is in charge of finding paths with enough funds from payer to payee. The payment algorithm settles the funds between payer and payee along the paths connecting them. Finally, the accountability algorithm allows the resolution of disputes in the presence of misbehaving users.

While frequently omitted or disregarded as an orthogonal problem, the design of the routing algorithm is key to the PBT network's effectiveness, characterized by the fraction of successfully resolved transactions; efficiency, characterized by the delays experienced during a transaction as well as the overhead created by transactions; and scalability, characterized by the ability of a PBT network to maintain effectiveness and efficiency for a growing base of users and transactions. PBT networks relying on the max-flow algorithm are unlikely to improve scalability beyond blockchains whereas other approaches might fail to discover a large fraction of paths. Efficiency, effectiveness, and scalability are supposedly the main concerns of companies interested in deploying or investing in PBT networks. As a consequence, the importance of privacy might be overlooked, entailing that the routing algorithm might reveal sensitive information such as the transaction value, the identity of payer and payee, and the debt of one user to another. In this

paper, we stress that all of effectiveness, efficiency, scalability, and privacy are important to the design of a routing algorithm. A privacy-preserving but inefficient algorithm is unlikely to be deployed.

The few routing algorithms proposed so far for PBT networks fail to achieve either privacy, efficiency, or scalability. For instance, Canal [35] relies on a single server to store the complete PBT network, find paths, and settle payments between users. Therefore, the server is trivially aware of all links between users and their transactions. PrivPay [21] leverages trusted hardware to encrypt the payment network data at the server and uses oblivious algorithms to hide the access patterns, thereby increasing the privacy for the links between users and their payments. Nevertheless, PrivPay still suffers from a single point of failure and low scalability. Flare [27], a routing algorithm for the Bitcoin Lightning Network, requires every user in the path from payer to payee to send the current fund amounts for their state channels to the payer, thereby leaking sensitive information [27, Section 3.6]. The most promising approach with regard to privacy is SilentWhispers [15], a fully distributed PBT network without a public ledger. The routing algorithm in SilentWhispers provides formal privacy guarantees but lacks efficiency, as we show in this paper.

In this work, we present SpeedyMurmurs, a routing algorithm for PBT networks that provides formal privacy guarantees in a fully distributed setting and outperforms the state-of-the-art routing algorithms in terms of effectiveness and efficiency. SpeedyMurmurs extends VOUTE [30], a privacy-preserving embedding-based [24] routing algorithm for message delivery in route-restricted P2P networks. Targeting message transmission in undirected and unweighted networks rather than payments, VOUTE is unequipped for dealing with weighted links and specifically changes of these weights as a result of previous transfers. SpeedyMurmurs combines the underlying ideas of VOUTE with the specifics of credit networks. In particular:

- SpeedyMurmurs considers both the available funds and the closeness to the destination of a neighbor when routing a payment, resulting in an efficient algorithm with flexible path selection.
- SpeedyMurmurs employs an on-demand efficient stabilization protocol that reacts to changes of links if necessary but keeps the overhead corresponding to these changes low.
- SpeedyMurmurs provides an improved handling of concurrent transactions by allowing nodes to proactively allocate exactly the amount of funds required for a transaction rather than barring concurrent transactions from using a link altogether or risking failures during the subsequent payment phase.
- In our simulation study, which models a credit network and transactions based on a real-world dataset of Ripple ranging from 2013 to 2016, SpeedyMurmurs performs transactions at about twice the speed of SilentWhispers and reduces the communication overhead of transactions by at least a factor of 2 while maintaining a similar or higher effectiveness.
- SpeedyMurmurs reduces the overhead of managing link changes by 2–3 orders of magnitude except for rare

phases (approximately one per year) in the Ripple dataset corresponding to sudden rapid growth.

- SpeedyMurmurs achieves a number of key privacy goals, in particular hiding the value of a transaction as well as the identities of payer and payee.

In summary, SpeedyMurmurs offers an efficient solution for privacy-preserving routing in credit networks, thus being a promising candidate for upcoming deployment of such networks.

## II. STATE OF THE ART AND LIMITATIONS

In this section, we describe SilentWhispers [15], the most promising approach for a distributed privacy-preserving credit network, and the concept of embedding-based routing in VOUTE [30], which we adapt to PBT networks in Section IV.

### A. Landmark Routing

The landmark routing technique [34] enables the computation of a subset of paths between a payer and a payee in a network graph without relying on the cost-intensive max-flow approach. The key idea of landmark routing is to determine a path from payer to payee through an intermediate node, called a *landmark*, usually a well-known node of high connectivity. While landmark routing does not discover all possible paths and hence might lead to a lower probability for a successful payment, past work indicates that the decrease of success is small in comparison to the gain in performance [21], [35].

Initially, each landmark starts two instances of the Breadth-First Search (BFS) algorithm. In the first instance, only forward edges are considered and shortest paths from the landmark to each node are calculated. The second instance considers only reverse edges and results in shortest paths between each node and the landmark. As PBT networks change over time, landmarks repeat the initialization process periodically.

The path discovery between a payer and payee then concatenates the path from the payer to the landmark (using reverse edges) and the path from the landmark to the payee (using forward edges). The payer can send funds along the path as long as the amount of funds is at most as high as the available credit on each link on the path.

There are two versions of landmark routing. The first version always concatenates a path from the source to a landmark and from the landmark to the destination. The second version, which we call *tree-only routing*, discovers the shortest path in the BFS tree, which does not necessarily contain a landmark.

1) *Landmark Routing in SilentWhispers*: SilentWhispers utilizes landmark routing to discover multiple paths and then performs multi-party computation to determine the amount of funds to send along each path. The initialization of the landmark routing follows the above description, using multiple landmarks that perform periodic BFSs.

The actual payment relies on two operations: a probe operation and the actual payment operation. We here describe the probe operation as it performs the routing and decides on the credit to be transferred along each path. The payment




Fig. 1: Examples of different tree routing schemes for landmark  $lm$ , sender  $s$ , receiver  $r$ .

operation then merely executes the transfers suggested by the probe operation in a secure manner.

At the core of the probe operation is a secret-sharing-based multiparty computation that computes the credit available in a path. After discovering paths between payer and payee using landmark routing, each pair of adjacent users in the path sends a share of their link’s value to each of the landmarks. The payer and payee must construct additional shares that act as padding in order to hide the actual length of the path, and effectively preserve the identities of the actual payer and payee. With the help of cryptographic signatures, relying on fresh keys to hide the identities of the nodes on the path, and using multiparty computation, the landmarks determine shares that the payer can combine to obtain the minimal available credit  $z_i$  of the  $i^{\text{th}}$  path. If the sum of the  $z_i$  values is at least equal to the total payment amount, the payer assigns values  $c_1, \dots, c_{|L|}$  to the paths such that  $c_i \leq z_i$ . The result of the probe operation are these values  $c_i$  and handles to the paths, which the payment operation leverages to perform the actual transfer.

2) *Weaknesses of SilentWhispers:* Based on the above description, we identify four issues related to the implementation of the routing algorithm in SilentWhispers. First, the periodic tree creation (execution of BFS) fails to take into account changes in the network immediately, which can lead to significant failure rates due to outdated information. Moreover, periodic tree creation induces unnecessary overhead due to reconstructing parts of the spanning tree that might not have changed.

Second, all paths include the landmarks even if i) the sender and receiver of a payment are in the same branch, or ii) there is a short path between sender and receiver but the links are not part of the spanning tree. Thus, the overall path used for the payment can be unnecessarily long, leading to longer delays and a lower success ratio due to the increased chance of encountering at least one link without enough funds.

Third, the probe operation requires that the nodes included in a transaction path send shares to all landmarks. This means that the transaction overhead scales quadratically in the number of landmarks.

Fourth, SilentWhispers does not provide a suitable solution for concurrency. Assume that one or more probe operations aim to use the same link. The probe operation can either provide the same amount of available credit for both links or block use of the link for some time after the first time a probe operation finds a path containing the link. The former can lead to failures later on as the available credit, while sufficient for

each transfer on its own, might not cover multiple transfers. While a block prevents such complications, it increases the likelihood of failures because probe operations cannot use certain links, which might have enough credit to execute multiple transactions. Hence, both approaches to concurrency have severe drawbacks.

In summary, landmark routing and SilentWhispers have various weaknesses that we aim to overcome in this work.

### B. Embedding-based Routing

Routing is essential for any network that forwards packets from a source to a destination. However, in contrast to PBT networks, the capacity of traditional network links does not change as a result of transferring a message. In the following, we summarize the concept of greedy embeddings [24], an efficient routing protocol requiring little state information. The fact that greedy embeddings maintain little state information and hence reveal little information about the network topology encourages their use for privacy-preserving communication systems such as Friend-to-Friend overlays [12], [30].

Embeddings rely on assigning coordinates to nodes in a network and having nodes forward packets based on the distances between coordinates known to that node and a destination coordinate. Greedy embeddings are similar to landmark routing in that they assign coordinates based on a node’s position in a spanning tree. However, greedy embeddings disregard the spanning tree after assigning the coordinates and in particular discover shorter paths using links that are not in the spanning tree. We refer to links that are not contained in the tree but are used during routing as *shortcuts*. When a node  $v$  forwards a message addressed to a destination coordinate,  $v$  chooses the neighbor with the coordinate closest to the destination coordinate to forward the message to. Hence,  $v$  might either use a link in the spanning tree (forwarding to a child or parent), or a shortcut.

Despite the fact that routes can contain shortcuts, there is no guarantee that routes with shortcuts exist. Hence, the links in the trees provide the guarantee that the routing works and removing any such links likely leads to failures. As a consequence, it is important to adapt the tree when the nodes or links change.

Prefix Embedding [12] is a greedy embedding that enables routing of messages in F2F overlays. As illustrated in Fig. 1, Prefix Embedding assigns coordinates in the form of vectors, starting with an empty vector at the landmark/root. Each internal node of the spanning tree enumerates its children and

appends the enumeration index of a child to its coordinate to obtain the child coordinate. The distance between two such coordinates corresponds to the length of the shortest path in the spanning tree between them; i.e., the distance of two coordinates  $id(u)$  and  $id(v)$  with  $|id(v)|$  denoting the coordinate length and  $cpl(id(u), id(v))$  denoting the common prefix is

$$d(id(u), id(v)) = |id(u)| + |id(v)| - 2cpl(id(u), id(v)). \quad (1)$$

Based on Eq. 1, nodes determine which neighbor is closest to the receiver in terms of their coordinates' distance and forwards a message accordingly. Figure 1 displays an example to illustrate the difference between various tree-based routing schemes and illustrates the coordinate assignment in Prefix Embedding.

1) *Prefix Embeddings in VOUTE*: VOUTE [30] is a routing algorithm building upon Prefix Embedding with the goal of anonymous and efficient message delivery for a dynamic route-restricted overlay; i.e., an overlay that does not allow the establishment of links between arbitrary nodes. We quickly describe how VOUTE addresses the issues of anonymity and dynamics.

Prefix Embedding reveals the unique receiver address of the destination. In contrast, VOUTE allows nodes to provide anonymous return addresses instead of their original coordinates. A receiver generates a return address by padding its coordinates to a fixed length and generating keyed hashes of the coordinate elements. The anonymous return address then corresponds to the keyed hashes and the key, allowing forwarding nodes to determine the common prefix length required in Eq. 1. Based on the common prefix length of the receiver coordinate  $id(r)$  and a neighbor's coordinate  $id(u)$ , forwarding nodes can compute  $d(id(u), id(r)) + \Delta$  with  $\Delta$  corresponding to the constant length of the padding. Hence, they can forward the message along the same path as when using clear-text coordinates, while maintaining the privacy of the receiver's true coordinates.

The original Prefix Embedding coordinates reflect an enumeration and hence have little entropy. As a consequence, VOUTE replaces the enumeration index with random  $b$ -bit numbers; e.g., for  $b = 128$ . In this manner, guessing the coordinate of an unknown node becomes computationally unfeasible for an adversary.

Rather than periodically reconstructing the spanning tree, VOUTE addresses dynamics with an on-demand stabilization protocol. When constructing the tree, nodes send invitations to all neighbors stating their coordinate and offering to become a parent. Each node accepts one such invitation but keeps the most recent invitation of all neighbors to quickly react to network dynamics. If nodes establish a new link, nodes already contained in the spanning tree offer invitations to their new neighbors. If a node is not yet part of the tree, it accepts the invitation. Otherwise, it stores it for future consideration. On the other hand, if a link in the spanning tree ceases to exist, the child node and all its descendants choose a new parent based on their remaining invitations. They then disseminate their new coordinate to all neighbors. In this manner, spanning trees and embeddings have an on-demand repair mechanism rather than periodic re-computation.

2) *Limitations of VOUTE*: VOUTE has not been defined in the context of PBT networks and therefore presents several limitations that must be overcome before considering it as a routing algorithm in PBT networks. In particular, VOUTE has incompatible assumptions with regard to the nature of links and network dynamics.

First, VOUTE considers undirected and *unweighted* links between pairs of users. In a PBT network instead, links are weighted and directed, as are payments. While all links allow message transfer in VOUTE, a link in a PBT network might not carry enough funds to perform a transfer. The directed nature of the links indicates that VOUTE's construction protocol is insufficient as it is unclear how to deal with unidirectional links. If unidirectional links are part of the spanning tree, a node (and its descendants) might only be able to send or receive funds but not both. The weighted nature of links and the impossibility to use links for all payments contradicts one of the key assumptions of VOUTE's protocol, namely that in the absence of link failures, all links can transfer messages. Applying VOUTE in the context of credit networks, it is necessary to design protocols that deal with weighted links and transfers.

Second, VOUTE considers dynamics in the form of nodes joining and leaving the network. However, in PBT networks, the weights of the links are the main source of change. In particular, each successful transaction might change several links. A variant of VOUTE for PBT networks would likely be inefficient if it reacts to all of these changes. Deciding on when and how to adapt to changes of links is important for the design of such a variant.

Finally, VOUTE does not have to deal with concurrency issues. While concurrent message transfers might increase delay and congestion, they do not change the capacity of links and transmitting a message does not affect the ability of the link to transmit future messages. However, separated probe and payment operations as in SilentWhispers, creates concurrency issues. SilentWhispers provides insufficient solutions here, so we require a new concurrency protocol.

In summary, although VOUTE presents an interesting alternative to landmark routing as implemented in SilentWhispers for the routing operation, its application in PBT network scenarios is not straightforward.

### III. SYSTEM MODEL AND GOALS

We start with a generic system model for distributed routing algorithms, followed by our privacy goals expressed in cryptographic games and our performance metrics.

#### A. Our Model

We model a distributed PBT network  $(G, w)$  as a directed graph  $G = (V, E)$  and a weight function  $w$  on the set of edges. The set of nodes  $V$  corresponds to the participants of the PBT network. A link (edge) from node  $u$  to  $v$  exists if  $u$  can transfer funds to  $v$ . We define the set of outgoing neighbors of a node  $v$  as  $N_{out}(v) = \{u \in V : (v, u) \in E\}$ . Correspondingly, we define the set of incoming neighbors of a node  $v$  as  $N_{in}(v) = \{u \in V : (u, v) \in E\}$ . Furthermore, a path  $p$  is a sequence of links  $e_1 \dots e_\ell$  with  $e_i = (v_{i1}, v_{i2})$ , and  $v_{i2} = v_{(i+1)1}$  for

$1 \leq i \leq \ell - 1$ . Moreover we denote by  $L = \{l_1, \dots, l_{|L|}\}$  a set of highly connected nodes, called landmarks, that are well known to other users in the PBT network. We denote by  $|L|$  the size of the set  $L$ .

The function  $w$  describes the amount of funds that can be transferred between two nodes sharing an edge. We thereby abstract from the specific implementation of the function  $w$ . For instance, in the Bitcoin Lightning Network, the function  $w : E \rightarrow \mathbb{R}$  is defined as  $w(u, v) = \text{pot}(u, v)$ , where  $\text{pot}$  defines how many bitcoins  $u$  can transfer to  $v$  in a payment channel opened between them. Moreover, each edge has associated an upper bound  $\text{cred}(u, v)$  and a lower bound set to 0. In general  $0 \leq w(u, v) \leq \text{cred}(u, v)$ .

We define the funds available in a path  $e_1, \dots, e_\ell$  as the *minimum pot*  $\text{pot}(e_i)$ . Moreover, we define the net balance of a node  $v$  as  $\text{cnode}(v) = \sum_{u \in N_{in}(v)} w(u, v) - \sum_{u \in N_{out}(v)} w(v, u)$ .

1) *Operations*: A PBT network consists of a tuple of algorithms ( $\text{setRoutes}$ ,  $\text{setCred}$ ,  $\text{routePay}$ ) defined as follows:

$\text{setRoutes}(L)$ : Given the set of landmarks  $l_1, \dots, l_{|L|}$ ,  $\text{setRoutes}$  initializes the routing information required by each node in the PBT network.

$\text{setCred}(c, u, v)$ : Given the value  $c$  and the nodes  $u$  and  $v$ ,  $\text{setCred}$  sets  $w(u, v) = c$ . In addition,  $\text{setCred}$  might alter the routing information initially generated by  $\text{setRoutes}$ .

$((\text{path}_1, c_1), \dots, (\text{path}_{|L|}, c_{|L|})) \leftarrow \text{routePay}(c, u, v)$ . Given a value  $c$ , a sender  $u$  and a receiver  $v$ ,  $\text{routePay}$  returns a set of tuples  $(\text{path}_i, c_i)$ , denoting that  $c_i$  funds are routed through the path described by  $\text{path}_i := e_{i1}, \dots, e_{i\ell_i}$ .

a) *Correctness*: A key property of a payment system is correctness. Intuitively, correctness indicates that the routing algorithm i) suggests to spend the desired funds  $c$  rather than a higher value and ii) suggests paths that indeed have sufficient funds. Let  $(\text{setRoutes}, \text{setCred}, \text{routePay})$  be the routing operations of a PBT network. We say that the PBT network is *correct* if for all results  $((\text{path}_1, c_1), \dots, (\text{path}_{|L|}, c_{|L|}))$  of  $\text{routePay}(c, u, v)$ , the following two conditions hold:

- $\sum_i c_i \leq c$
- For each  $\text{path}_i := e_{i1}, \dots, e_{i\ell_i}$  and each  $e_{ij}$ ,  $c_i \leq w(e_{ij})$ .

We note that the  $\text{routePay}$  operation could return paths that contribute  $\sum_i c_i < c$ , and it is still considered correct. This accounts for the cases where the PBT network does not provide enough liquidity between the sender and receiver to perform a transaction.

## B. Attacker Model

We consider a fully distributed network. Our primary attack scenario is companies and individuals interested in a user's financial situation rather than governmental security agencies. The adversary controls a subset of the nodes in the network either by inserting its own nodes or corrupting existing nodes. We assume that the adversary cannot choose the set of users at will, as some users will be harder to corrupt by social engineering or malware attacks. In general, we assume that the attacker does not know *all* links and nodes in the network

and in particular cannot access routing information locally stored at non-compromised nodes. The assumption that the attacker does not know the complete topology of a large-scale distributed system with participants from a multitude of regions and countries seems realistic for our attack scenario.

Our adversary aims to undermine the privacy rather than perform a large-scale denial-of-service attack. We argue that the primary defense against denial-of-service attacks is detection and expulsion of malicious nodes. While related to routing, different operations are required for realizing detection and expulsion, and they are out of scope for this paper.

While our overall adversary model limits the adversary's capacities, we nonetheless define our value privacy goal for an attacker that has a global view of the topology, indicating that we can still achieve some privacy against a stronger adversary.

## C. Privacy Goals

The hope that cryptography and decentralization might ensure robust privacy was among the strongest drivers of Bitcoin's and blockchains' early success. We expect businesses and customers employing the PBT networks to be interested in hiding their transactions from competitors and even service providers. Therefore, ensuring privacy for path-based transactions is important.

Like PrivPay [21] and SilentWhispers [15], we aim to hide values (*value privacy*), and the identities of payer and payee (*payer/payee privacy*) of path-based transactions. We use the term *transaction privacy* to refer to meeting all three of these notions. Next, we informally describe these privacy properties for PBT networks, and refer the readers to the PrivPay paper [21] for the formalized versions defined in the context of credit networks.

*Value Privacy*: A PBT network achieves value privacy if it is not possible for any adversary to determine the total value of a transaction between non-compromised users as long as none of the employed intermediate nodes is compromised.

Let  $s$  and  $r$  be two non-compromised users, and let  $(\text{path}_1, c_1), \dots, (\text{path}_{|L|}, c_{|L|})$  be the result of a  $\text{routePay}(c, s, r)$  operation. If for every path  $\text{path}_i$ , all nodes on that path are non-compromised, the adversary (even a global passive adversary) obtains no information about the transaction value  $c$ .

Notice that, as elaborated in Section IV-G, we can provide a weaker form of value privacy even when the adversary compromises some intermediate nodes as long as all nodes on at least one of the employed paths remain non-compromised.

*Payer Privacy*: A PBT network achieves payer privacy if it is not possible for any adversary to determine the payer in a path-based transaction between non-compromised users.

In particular, for two non-compromised users  $s$  and  $r$ , the attacker should not be able to determine the payer  $s$  of any routing operation  $\text{routePay}(z, s, r)$ , unless she has complete knowledge of  $s$ 's incoming links i.e., she knows the set  $N_{in}(s)$  though not necessarily the funds of the links  $e \in N_{in}(s)$ .

Note that although the local attacker without a global view of the network might know (and even control) all nodes in

$N_{in}(s)$ , she might not be aware that she does control *all* such nodes. As a consequence, similar to P2P anonymity systems [19], [20], [29], controlling all neighbors does not automatically mean she can be sure that  $s$  did initiate the routing. Therefore, we expect the payer privacy to hold even when the attacker controls all nodes in  $N_{in}(s)$  for the payer  $s$  but does not know that she does control the whole set.

*Payee Privacy* is defined analogously to payer privacy, and the adversarial assumptions also remain the same except that instead of neighboring nodes  $N_{in}(s)$  of the payer  $s$ , now for payee privacy, we consider  $N_{out}(r)$  of the payee  $r$ .

#### D. Performance Metrics

In this section, we describe the goals to be achieved by a routing protocol, which we denote generically by  $\mathcal{R}$ . In the following, we denote by  $(G_t, w_t)$  the snapshot of a PBT network at time  $t$ . Note that although we abstract away the payment and accountability protocols in this work, a PBT network must implement them and therefore a PBT network is dynamic. Let  $\{(t_i, c_i, s_i, r_i)\}$  be a set of payment requests from  $s_i$  to  $r_i$  for an amount  $c_i$  at time  $t_i$ .

The performance of a routing protocol  $\mathcal{R}$  is characterized by the following four metrics:

- **Success ratio:** Let  $((path_1, c_1), \dots, (path_{|L|}, c_{|L|}))$  be the set of paths returned by  $routePay(c, s, r)$  as implemented in  $\mathcal{R}$ . We consider the transaction successful only if  $\sum_i c_i = c$ . The success ratio describes the fraction of transactions that are successful.<sup>1</sup>
- **(Hop) Delay:** The delay of  $\mathcal{R}$  with regard to a transaction  $(t, c, s, r)$  is the difference between the time of termination and the initiation time  $t$ . In the absence of a concrete implementation including realistic computation and communication latencies, we provide an abstract measurement of the delay as follows. Let  $m_1$  and  $m_2$  be messages sent by  $\mathcal{R}$ . We say  $m_2$  is *subsequent* to  $m_1$  if a node sends  $m_2$  as a result of receiving  $m_1$ . The hop delay is the length of the longest chain of subsequent messages sent by  $\mathcal{R}$ .
- **Transaction Overhead:** Nodes exchange messages to execute a transaction  $(t, c, s, r)$ . The transaction overhead denotes the number of exchanged bits. As before, in the absence of a concrete implementation, we abstract it by assuming equal-sized messages for each implementation and stating the number of messages as the overhead.
- **Stabilization Overhead:** Analogously to the transaction overhead, the stabilization overhead corresponds to the number of bits sent within a certain time interval to maintain necessary state information, as required by the implementation of the operation  $setRoutes()$  in  $\mathcal{R}$ . Again, we abstract from the concrete implementation by stating the number of messages instead of the number of bits.

The first two metrics heavily impact the perceived quality of service while the latter two directly relate to network congestion and hence impact the delay. Furthermore, the overhead determines the load on the user devices.

<sup>1</sup>This inherently assumes a payment protocol that always succeeds after a route with enough credit has been found. We thereby abstract away the details of the payment protocol.

## IV. OUR CONSTRUCTION

We first describe the key ideas of our design and then discuss each of the three operations in detail. We present pseudocode for centralized versions of the algorithms, which allows presenting the algorithms in a linear compact manner. We then describe how the distributed versions used within a PBT network differ from the centralized version.

### A. Assumptions

Each user in the PBT network locally maintains the information of the links with her neighbors. We further assume that users sharing a link can send messages to each other through an authenticated and confidential communication channel. Moreover, we assume that there exist a set of highly connected nodes, called landmarks, that are well known to other users in the PBT network. We note that these assumptions are in tune with other distributed PBT networks such as SilentWhispers.

Throughout this section, we refer to links that have non-zero funds in both directions, i.e., links  $(u, v)$  and  $(v, u)$  with  $pot(u, v) > 0$  and  $pot(v, u) > 0$  as bidirectional.  $u$  and  $v$  have a unidirectional link if one of the two links does not exist or has zero funds.

### B. Overview and Key Ideas

We here describe the key ideas of SpeedyMurmurs with regard to the operations  $setRoutes$ ,  $setCred$ , and  $routePay$ . In particular, we focus on the predominant differences to existing protocols.

*setRoutes:* In this algorithm, we construct multiple embeddings, one for each landmark. As mentioned in Section II-B, VOUTE offers a protocol for BFS-based coordinate assignment that assumes unweighted and undirected links. We modify this protocol by dividing it into two phases. First, we only add bidirectional links. In the second phase of the protocol, nodes that are not yet part of the spanning tree join by adding unidirectional links to the spanning tree.

*setCred:* We leverage SilentWhispers' protocol for setting the funds to  $c$ . The main challenge lies in designing suitable changes to the embedding. VOUTE does not provide any guidance on how to react to changes of weights. In agreement with *setRoutes*, we decided to only initiate changes if the two nodes i) established a new link of non-zero weight (or set the value of link from 0 to a non-zero value), and ii) removed a non-zero link (or set its value to 0). If they established a new link, one of the nodes can choose the other as a parent if it does not have a parent or the link to its current parent only has credit in one direction. In contrast, if a link is removed, one of the nodes has to select a new parent (and coordinate) if the other node previously was its parent. Furthermore, any descendants of the affected node have to change coordinates.

*routePay:* The routing consists of three steps: i) the payee generates anonymous return addresses and sends them to the payer, ii) the payer randomly splits the transaction value on  $|L|$  paths, one for each landmark, and iii) VOUTE's routing algorithm finds a path between payer and payee, restricted to links that have sufficient funds. Our algorithm allows a flexible choice of routes, preferring paths with high funds. Determining

the funds transferred along each path beforehand enables nodes to block a certain amount of credit during the probe operation and ensure that the subsequent payment succeeds without blocking all of the link's funds.

### C. *setRoutes*

In the initialization phase, described in Algorithm 1, we derive the embeddings. Iterating over all landmarks, Algorithm 1 assigns the landmark coordinate to be the empty vector (Line 3) and appends the landmark to a queue (Line 5). The main loop of the algorithm then processes the queue. In each step, the algorithm removes a node from the queue (Line 8) and considers all its neighbors. If a neighbor  $n$  does not have a coordinate yet and is eligible to have one, the algorithm integrates  $n$  into the spanning tree, assigns a coordinate by concatenating the parent coordinate and a random  $b$ -bit number, and appends it to the queue (Lines 14–17). The criterion determining the eligibility to be part of the spanning depends on the phase of the algorithm: Initially ( $bi = true$ , Line 6), a node is eligible if the available credit on the link to and from its potential parent is non-zero. In the second phase of the algorithm ( $bi = false$ ), triggered by an empty queue (Lines 18–21), all nodes can join the spanning tree. Note that Algorithm 1 does not prevent a child from choosing a parent such that they do not have funds in either direction. As such links do not serve any purpose in a PBT network, we assume that the network does not allow them. Alternatively, Algorithm 1 can check if the funds on a unidirectional link is non-zero before allowing a parent-child relation. The algorithm terminates once the queue is empty, indicating that all nodes in a connected graph have coordinates.

---

#### Algorithm 1 *setRoutes*

---

```

1: # Input: Graph  $G$ , landmarks  $l_1, \dots, l_{|L|}$ 
2: for  $i = 1$  to  $|L|$  do
3:    $id_i(l_i) = ()$ 
4:    $q =$  empty queue
5:   add  $l_i$  to  $q$ 
6:    $bi = true$  # first consider nodes with links in both
   directions
7:   while  $q$  is not empty do
8:      $node =$  remove head of  $q$ 
9:     for all  $n$ : neighbors of node do
10:       $n$  stores  $id_i(node)$ 
11:      # Assign coordinate if  $n$  does not have one
12:      if  $id_i(n)$  is not set then
13:        if ( $pot(node, n) > 0$  and  $pot(n, node) > 0$ ) or
         $!bi$  then
14:           $parent_i(n) = node$ 
15:           $n$  chooses random  $b$ -bit number  $r$ 
16:           $id_i(n) = concatenate(id_i(node), r)$ 
17:          add  $n$  to  $q$ 
18:      # Add nodes with unidirectional links
19:      if  $q$  is empty and  $bi$  then
20:         $bi = false$ 
21:        add all nodes  $n$  with  $id_i(n)$  set to  $q$ 

```

---

In a distributed scenario, there are no central queues. Rather, nodes send messages to their neighbors when they join a spanning tree. Each message includes the index  $i$  of

the tree and the coordinate  $id_i(node)$  of the potential parent. Starting the second phase is tricky in a distributed scenario and will not be simultaneous for all nodes. Rather, we choose a time limit  $\tau$  that represents an upper limit on the time the initialization should take. If a node  $n$  receives a message of the form  $(i, id_i(node))$  from a neighbor  $node$  with only one link of non-zero weight,  $n$  waits for time  $\tau$ . If none of  $n$ 's neighbors with bidirectional links to  $n$  indicate that they are potential parents until the waiting period ends,  $n$  selects  $node$  as its parent.

### D. *setCred*

*setCred* reacts to a pair of nodes  $(u, v)$  that want to change the value of their shared link to  $c$ . In Algorithm 2,  $u$  and  $v$  first execute SilentWhispers' secure Link Setup protocol [15, Protocol 1] to change the available credit of the link  $(u, v)$ . Afterwards, the algorithm determines if the value change should lead to coordinate changes. In total there are three situations that indicate a need for a coordinate change:

- 1) New non-zero link: One of the nodes is not yet part of the tree and should hence choose the other as their parent (Lines 9–13) to be able to participate in the routing.
- 2) New non-zero link:  $u$  and  $v$  share a bidirectional non-zero link and (without loss of generality)  $u$  has only a one-directional link to its current parent. Then  $u$  should change its parent to  $v$  if  $v$  has a bidirectional link to its parent (Lines 14–22). In this manner, a bidirectional connection replaces a unidirectional link in the spanning tree and increases the likelihood of successfully transferring funds.
- 3) Removed link:  $u$  is a child of  $v$  or  $v$  is a child of  $u$  (Lines 25–28). The child node should select a new parent to increase the number of non-zero links in the spanning tree and the likelihood of transferring funds.

If one of  $u$  or  $v$  changes its parent, all descendants remove their coordinates and inform their neighbors of the removal. Afterwards, they all choose a new parent and corresponding coordinate. In agreement with the initialization *setRoutes*, nodes first consider only neighbors to whom they have non-zero links in both directions. However, if a node does not have such links to any neighbor, it considers links in one direction. If they have several suitable parents, they choose their parent randomly from those candidates with the shortest coordinates, as having short routes to the landmark reduces the lengths of paths [30]. After choosing a new coordinate, nodes forward the new coordinate and the tree index to all their neighbors. We do not present the pseudocode, as it is very similar to Algorithm 1.

The distributed variant of Algorithm 2 follows the same principles but requires the exchange of messages for nodes to communicate information.  $u$  and  $v$  exchange information about the link to their parents. Each of them then individually decides if they want to add or remove the other as a parent. Starting from the node *reset* that aims to reset its coordinate, all descendants inform their neighbors first that they remove their old coordinate for the tree  $i$ . Children of a node in turn remove their own coordinate and send the respective messages. In the second phase, nodes select their new coordinates and inform their neighbors. As the two phases are likely to run

---

**Algorithm 2** *setCred*

---

```
1: # Input: Graph  $G$ ,  $u, v \in V$ , new value  $c$ 
2:  $old = pot(u, v)$  # Previous value of  $pot(u, v)$ 
3:  $u$  and  $v$  execute SilentWhisper's Setup Protocol
4: # check if coordinate change necessary
5: for  $i = 1$  to  $|L|$  do
6:    $reset = null$  # node whose coordinate should change
7:   # case:add link
8:   if  $old == 0$  and  $c > 0$  then
9:     # If one node does not have a coordinate
10:    if  $id_i(v)$  is not set and  $id_i(u)$  is set then
11:       $reset = v$ 
12:    if  $id_i(u)$  is not set and  $id_i(v)$  is set then
13:       $reset = u$ 
14:    # One node has unidirectional link to parent
15:    if  $reset = null$  then
16:      if  $pot(u, v) > 0$  and  $pot(v, u) > 0$  then
17:         $a_1 = (pot(u, parent_i(u)) == 0$  or
18:               $pot(parent_i(u), u) == 0)$ 
19:         $a_2 = (pot(v, parent_i(v)) == 0$  or
20:               $pot(parent_i(v), v) == 0)$ 
21:        if  $a_1$  and  $!a_2$  then
22:           $reset = v$ 
23:        if  $a_2$  and  $!a_1$  then
24:           $reset = u$ 
25:    # case:remove link
26:    if  $old > 0$  and  $c == 0$  then
27:      if  $parent_i(u) == v$  then
28:         $reset = u$ 
29:      if  $parent_i(v) == u$  then
30:         $reset = v$ 
31:    # change coordinates
32:    if  $reset != null$  then
33:      delete coordinates of  $reset$  and descendants
34:      have nodes choose new parent
```

---

in parallel in the distributed setting, nodes have to ensure that they do not choose a previous descendant as a parent before the descendant chooses a new coordinate. However, the nature of the coordinates makes it easy to prevent such cycles in the tree by disallowing a node  $v$  from choosing a parent whose coordinate contains  $v$ 's previous coordinate as a prefix.

### E. routePay

*routePay* discovers a set of paths from the payer to the payee. It corresponds to the probe operation in SilentWhispers. Algorithm 3 divides the process into three steps: i) generation of receiver addresses (Lines 2–5), ii) splitting the total transaction value  $c$  randomly on  $|L|$  paths, and iii) finding paths for all embeddings that can transmit the required value.

First, the payee generates anonymous return addresses for all embeddings and sends them to the payer (Lines 2–5). Second, the payer splits the transaction value randomly between all paths (Line 7). By defining a per-path value before routing, we i) avoid the costly multiparty computation of SilentWhispers and ii) allow the algorithm to choose between several possible routes. Avoiding the multiparty computation of the minimum also removes a privacy leakage, as knowing

the minimum value of funds available on each link naturally reveals information about the links.

Third, the route discovery starts at  $v$  and each node selects a neighbor to be the next node on the route. In VOUTE, each node would select the neighbor with the coordinate closest to the destination, using the function  $\tilde{d}$  that compares a coordinate with an anonymous return address. However, such a choice might not be suitable for routing funds as the link might have insufficient available credit. As a consequence, the routing only considers links  $(v, u)$  with *guaranteed* available credit  $pot_A(v, u)$  of at least  $c_i$  (Line 14). We differentiate between available credit  $pot(v, u)$  and *guaranteed* available credit to deal with concurrency.  $pot_A(v, u)$  is a lower bound on the available credit if ongoing probe operations succeed. Initially,  $pot_A$  equals the actual available credit  $pot$ . We do not include the initialization in Algorithm 3 as multiple concurrent executions of *routePay* can impact  $pot_A(v, u)$  and the algorithm might start with  $pot_A(v, u) < pot(v, u)$ . If a probe operation indicates that a payment will transmit funds  $c_i$  along a link  $(v, u)$ , we proactively decrease the *guaranteed* available credit by  $c_i$  (Line 17) to keep future routings from using the link unless they require at most the guaranteed available credit. If the routing fails, we add  $c_i$  to  $pot_A(v, u)$  again (Lines 21–23). The routing fails if a node  $v$  cannot find a neighbor with a coordinate closer to the destination than  $v$ 's coordinate and a link of sufficient guaranteed available credit.

---

**Algorithm 3** *routePay*

---

```
1: # Input: Graph  $G$ , payer  $src$ , payee  $dst$ , value  $c$ 
2: # get addresses
3: for  $i = 1$  to  $|L|$  do
4:   use VOUTE's algorithm to generate return address
5:    $add_i(dst)$ 
6:    $dst$  sends  $add_i(dst)$  to  $src$ 
7: # value shares for each path
8:  $src$  splits  $c$  into shares  $c_1, \dots, c_{|L|}$ 
9: # routing
10:  $path_i =$  empty list of links
11: for  $i = 1$  to  $|L|$  do
12:    $v = src$ 
13:    $fail = false$ 
14:   while  $!fail$  and  $v != dst$  do
15:      $C = \{u \in N(v) : \tilde{d}(id_i(u), add_i(dst)) <$ 
16:            $\tilde{d}(id_i(v), add_i(dst)), pot_A(v, u) \geq c_i\}$ 
17:     if  $C$  not empty then
18:        $next = u$  in  $C$  with minimal  $\tilde{d}(id_i(u), add_i(dst))$ 
19:        $pot_A(v, u) = pot_A(v, u) - c_i$ 
20:        $v = next$ 
21:     else
22:        $fail = true$  # Routing failed
23:   if routing failed then
24:     for all  $i = 1 \dots |L|$ ,  $e \in path_i$  do
25:        $pot_A(e) = pot_A(e) + c_i$ 
```

---

Algorithm 3 achieves correctness, as defined in Section III, because i)  $\sum_{i=1}^{|L|} c_i = c$  and ii) nodes always select links  $e$  with  $pot(e) \geq pot_A(e) \geq c_i$  on the  $i^{\text{th}}$  path.

In the distributed variant of Algorithm 3, nodes send messages to the next node on the path, which contain the



address  $add_i$  and the partial value  $c_i$ . Nodes report failures and successes to the payee by sending messages along the reverse path. To account for messages getting lost, nodes also reset  $pot_A$  if a payment operation does not follow a probe operation within a certain time.

#### F. Parameters

Several parameters govern the performance of the above routing algorithm. First, the number  $|L|$  of landmarks determines the number of returned paths. The transaction and stabilization overhead increases roughly linearly with  $|L|$  as routing and stabilization is required for each embedding. Similarly, the delay corresponds to the longest route in any embedding and hence is likely to increase with  $|L|$ . The impact of  $|L|$  on the success ratio highly depends on the scenario. The second parameter is  $a$ , the number of transaction attempts. A sender  $s$  can attempt to perform a transaction up to  $a$  times. Only if all attempts fail,  $s$  considers the transaction failed.  $s$  chooses the interval between two consecutive attempts uniformly at random within an interval of length  $tl$ . A repeated transaction attempt executes the above routing algorithm for the same sender, receiver, and value but uses different shares  $c_1, \dots, c_{|L|}$ . In addition to the parameters  $|L|$ ,  $a$ , and  $tl$ , the choice of the landmarks impacts the performance. Commonly, landmarks are nodes corresponding to financial institutions and hence have a large number of links, possibly leading to spanning trees of a lower depth and a higher performance. We characterize the impact of these parameters in our performance evaluation.

#### G. Privacy Analysis

Next, we argue that SpeedyMurmurs achieves the privacy goals proposed in Section III-C.

*Value Privacy:* Informally, we say that a PBT network achieves value privacy if the adversary cannot determine the value  $c$  of a  $routePay(c, u, v)$  operation between two non-compromised users, if the adversary is not sitting in any of the involved routing paths.

SpeedyMurmurs is a distributed PBT network and, in particular, the  $routePay$  is defined such that only users in the paths between the payer and payee are involved. Therefore, if the adversary does not compromise any such users, she does not get any information about the routed value (because the point-to-point communications are encrypted) and thereby value privacy is achieved.

An alternative scenario appears when the adversary corrupts users in some of the paths between payer and payee, but not all. In such case, we cannot prevent the adversary from estimating  $c$ . As we have  $c_i \geq 0$  for all  $i = 1 \dots |L|$ , knowing a subset of these values naturally reveals information about the total value  $c$ , namely that  $c \geq c_i$ . Moreover, as SpeedyMurmurs shares the value  $c$  uniformly among the paths and uses only positive shares, an adversary can estimate  $c$  as  $|L| * c_i$  upon observing  $c_i$ .

*Payer Privacy:* Informally, we say that a PBT network achieves payer privacy if an adversary compromising intermediate users between payer and payee cannot determine the actual payer of a  $routePay$  operation between non-compromised users.

An attacker sitting on the path between payer  $s$  and payee  $r$  might receive an anonymous routing address  $s$  (i.e., the adversary managed to corrupt the payer's neighbor). Nevertheless, as SpeedyMurmurs is a distributed PBT network, the adversary cannot determine whether the actual payer is  $s$  or another user  $s'$  connected to  $s$  through a direct link or a path of non-compromised users.

We note that achieving the notion of payer privacy in the presence of an adversary that has knowledge of the  $s$ 's complete topology (i.e., adversary knows the complete set  $N_{in}(s)$ ) and is aware of that, is not possible. However, in a large-scale distributed system, it seems unlikely that an adversary can be certain about all links a particular non-compromised user establishes. Hence, we consider it likely that an attacker not being aware of all links of a particular node is realistic.

*Payee Privacy:* Informally, we say that a PBT network achieves payee privacy if an adversary compromising intermediate users between payer and payee cannot determine the actual payee of a  $routePay$  between non-compromised users.

As before, the adversary compromising the user before the payee  $r$  might relay to  $r$  an anonymous return address. Nevertheless, as shown in the evaluation of VOUTE [31], an anonymous return address does not leak the corresponding user in the network. Therefore, the adversary cannot fully determine yet if  $r$  is the actual receiver, or the routing message is intended for another payee  $r'$  connected to  $r$  through a direct link or a path of non-compromised users. We note that payee privacy cannot be achieved by an adversary that knows the complete set  $N_{out}(r)$  and is aware of that fact. As above, we consider that unlikely given the distributed nature of SpeedyMurmurs.

#### H. Summary

In this section, we introduced SpeedyMurmurs, which proposes a privacy-preserving routing algorithm for PBT networks. Our key contributions in modifying VOUTE to the scenario of credit networks are i) the use of a two-phase construction protocol to account for the existing of unidirectional links (Algorithm 1), ii) the identification of criteria on when to apply on-demand maintenance (Algorithm 2), iii) the design of a path discovery algorithm that can adaptively choose links based on both the available credit and the coordinates of the neighboring nodes and can handle concurrency (Algorithm 3). Apart from using embedding-based routing, SpeedyMurmurs distinguishes itself from SilentWhispers by splitting the credit between paths before the path discovery. In this manner, nodes can base their forwarding decisions on the amount of credit they should forward rather than only their neighbors' distances to the destination. On the other hand, distributing funds before the path discovery prevents the algorithm from taking the overall available funds on the path into consideration. In the next section, we evaluate the impact of our design decisions on efficiency and effectiveness, analyzing in particular how the order of routing and fund distribution relates to the success ratio.

## V. PERFORMANCE EVALUATION

In this section, we evaluate the performance of SpeedyMurmurs in comparison to the related work, in particular

SilentWhispers’ landmark routing. For brevity, we write SilentWhispers to indicate the version of landmark routing applied in SilentWhispers.

More precisely, we aim to answer the following research questions:

- How do SpeedyMurmurs and SilentWhispers perform with regard to success ratio, delay, and overhead using a real-world dataset?
- SpeedyMurmurs and SilentWhispers differ in three major areas—routing algorithm, random credit assignment, and dynamic stabilization. What is the impact of each of these modifications on the above performance criteria?
- How do these results compare to the performance of other approaches?
- What is the impact of the landmark selection, the number of trees, and the number of transaction attempts?
- How does the long-term evolution of the credit network affect the performance?

We start by describing our simulation model and datasets. Afterwards, we specify the parameters of our simulation. We present and discuss our results.

Generally, our simulation executes the routing algorithm  $\mathcal{R}$  and performs the payment (if successful). We include the payment to realistically assess the stabilization overhead due to link changes. However, we did not implement any security measures that are usually part of the payment because they do not affect the routing algorithm and its performance. In particular, we do not execute the link setup protocol that ensures that neighboring nodes agree on the value of their link and later can settle disputes by providing signed statements of the change.

#### A. Simulation Model

We extended GTNA [32], a framework for graph analysis, to include our credit transaction mechanisms. In particular, GTNA offers templates for routing algorithms and performance metrics. We added functionality specific to PBT networks, in particular the functionality to dynamically update link weights.

Initially, our simulation constructs a credit network with nodes and links according to a provided description. Afterwards, we simulate a sequence of events in the credit network. A list of transactions, changes to links, and periodic re-computations of the spanning tree (only required for SilentWhispers), ordered by their temporal occurrence, determined the sequence of events. In the absence of realistic latency and bandwidth models, we did not model concurrency in our simulation. The simulation executed each event, including resulting changes to the spanning trees, before starting the next event.

We implemented two simulation modes. First, we considered a static credit network. In each step, the simulation executed a transaction and subsequently repaired the spanning tree if dynamic stabilization was applied. Afterwards, it returned the credit network to its original state. Second, we considered a dynamic network evolving over time. Transactions, node churn, and modifications of the extended credit changed the structure of the network and the weights on the links. While the

second mode was more realistic, it prevented a straightforward comparison of different approaches for individual transactions due to the differences in state at the time of the transaction.

We implemented the routing and stabilization algorithms of SilentWhispers and SpeedyMurmurs as specified in Sections II and IV, respectively. However, we disregard the cryptographic details for our evaluation, as they do not affect our performance metrics. Instead, the sender and receiver both send only one message to each landmark forwarded by all nodes on the shortest path to the landmark. In our implementation of SilentWhispers, each landmark then sends a message to all remaining landmarks, which is forwarded along the shortest paths, to account for the multi-party computation. When combining embedding-based routing with multi-party computation, the receiver sends messages to all landmarks. In addition to enabling the evaluation of each individual modification, the alignment of the two designs also resulted in a fairer comparison of overheads, as the original SilentWhispers sends all elements of a signature chain individually and thus results in a higher overhead as compared to sending them in one message. As SilentWhispers’ authors do not specify how the sender decides on the amount of partial credit  $c_i$  assigned to the  $i^{\text{th}}$  path, we decided to divide the total credit randomly between paths in agreement with the available minimum. In other words, if the sum of all minimal values was at least equal to the total transaction value  $c$ , we first divided  $c$  randomly upon the paths. We then randomly re-assigned all credit that exceeds the minimal value along a path to the remaining paths. We repeated the re-assignment step until the partial credit of each path was at most equal to the minimal credit on the path. During the simulation, we recorded all information necessary to derive the performance metrics described in Section III.

For SpeedyMurmurs and SilentWhispers, we consider the following parameters: i) the number of trees  $|L|$ , ii) the number of attempts  $a$  that nodes try to perform a transaction before declaring it failed, iii) the maximal interval  $tl$  between two attempts for the same transaction, and iv) the interval *epoch* between two periodic re-computations of the trees for SilentWhispers. For comparison, we expressed the stabilization overhead for SpeedyMurmurs in stabilization messages per *epoch*. In addition to the above parameters, we provided two approaches for choosing landmarks: choosing the nodes of maximal degree or choosing random nodes. Here, we define the maximal degree of a node as the number of connections with positive available credit in both directions. For the evolving credit network, we chose the nodes with the highest initial degree.

We implemented distributed versions of the Ford-Fulkerson max-flow algorithm [9] and tree-only routing for comparison. Tree-only routing only uses links in the spanning tree but chooses the shortest path rather than always passing through the landmarks. For Ford-Fulkerson, we replaced the centralized computation with a distributed version that discovers residual flows using a breadth-first search.

#### B. Dataset

We obtained datasets from crawling the PBT network Ripple [5]. In particular, we obtained crawls of the complete network from November 2016 and all link modifications and

transactions since its creation in January 2013. Based on these crawls, we derive datasets for both our simulation modes, the static and evolving network. In the following, we first describe our crawling method, followed by post-processing of the crawled data. Last, we present properties of the resulting datasets.

We restricted our evaluation to funded accounts: a Ripple account is *funded* when it owns a certain amount of XRP.<sup>2</sup> In April 2017, a user needed 20 XRP to fund an account. In this paper, we disregard transferring credit from one currency to another. Hence, we converted all values to US dollars and deleted all links and transactions in non-fiat currencies. After cleaning the dataset according to these three rules, we derived the credit network  $C'_{Nov16}$  for November 2016 and lists of both transactions and link value changes, sorted in temporal order. Based on the resulting transaction and link modifications lists, we then generated the credit network  $C'_0$  at the time of the first transaction as a starting point of our second mode, the evolving network. As our data does not reveal when nodes join and leave the network, we included all crawled nodes and links in our initial credit network but set the weight of links  $(u, v)$  that come into existence at a later point to 0. During the spanning tree construction, such links are ignored.

We resolved three inconsistencies between our model and the datasets. In rare cases, Ripple exhibits invalid credit arrangements; i.e., links  $(u, v)$  such that the balance  $bal(v, u)$  exceeds the available credit  $cred(v, u)$ . Usually, such occurrences result from changes to the extended credit agreement. We deleted all such links from the dataset. Furthermore, we removed self-transactions from the dataset, as they do not require routing algorithms according to our model. Last, landmark routing requires paths between all nodes and the landmarks, so that we restricted our evaluation to the giant component. These processing steps turned the initial snapshots  $C'_{Nov16}$  and  $C'_0$  into our final datasets  $C_{Nov16}$  and  $C_0$ . We obtained the final transaction and link modification lists by restricting the previous lists to entries involving only nodes in the final snapshots.

We now describe the properties of the resulting dataset.  $C_0$  contained 93,502 nodes and a total of 331,096 links, whereas  $C_{Nov16}$  contained 67,149 nodes and 199,574 links. The reason for the disparity is that  $C_0$  contained all active links and their adjacent nodes for a period of more than 3 years, whereas  $C_{Nov16}$  was a snapshot of the network on one particular date. Our final transaction lists had 970,472 and 692,737 entries for  $C_0$  and  $C_{Nov16}$ , respectively. We recorded a total of 652,216 link modifications for the evolving network  $C_0$ . The dataset and the code are publicly available.<sup>3</sup>

### C. Simulation Setup

Our first simulation setup realized the static simulation mode on the basis of the snapshot  $C_{Nov16}$ . We repeated simulations 20 times, using a different set of 50,000 transaction for each run. We chose these transactions pseudorandomly, seeded by the run number, from all transactions that were successful using Ford-Fulkerson, a total of 331,642 transactions. We then

evaluated all 8 possible combinations of routing algorithms (landmark routing or embedding-based), credit assignments to paths (multi-party computation or random assignment), and stabilization algorithms (periodic or on-demand) for the parameters  $|L| = 3$  and  $a = 2$ . We chose  $epoch = 1000$ , meaning we recomputed spanning trees each 1000 transactions. We choose the re-queuing interval as  $tl = 2 \cdot epoch$ . For the landmark selection, we considered both options: random and highest degree. Note that random choices were deterministic in the run number, ensuring comparability of all approaches under the same circumstances. For comparison with related approaches, we evaluated two versions of tree-only routing, using SilentWhispers' multi-party computation and periodic stabilization for the first version and SpeedyMurmurs' random credit assignment and on-demand stabilization for the second. We then evaluated the impact of the different parameters for SilentWhispers and SpeedyMurmurs. We vary the number of landmarks  $|L|$  between 1 and 7 and the number of attempts  $a$  between 1 and 10.

Our second simulation setup realized the evolution of the network under different algorithms: Ford-Fulkerson, SilentWhispers, and SpeedyMurmurs. Starting from the initial network  $C_0$ , the simulation initiated the transactions and changes link values according to the dataset. For SilentWhispers and SpeedyMurmurs, we set  $|L| = 3$ ,  $a = 2$ ,  $epoch = 1000\delta_{Av}$ , and  $tl = 2\delta_{Av}$  with  $\delta_{Av}$  denoting the average time between two transactions. In this manner, an epoch roughly corresponds to a day. We chose landmarks of the highest degree for SpeedyMurmurs and SilentWhispers. As Ford-Fulkerson is a deterministic algorithm, we only executed it once but averaged our results for SpeedyMurmurs and SilentWhispers over 20 runs.

### D. Results

We start by comparing a wide range of algorithms for the static simulation setup. Table I displays the results for different combinations of the three proposed modifications to SilentWhispers as well as our implementations of tree-only routing and Ford-Fulkerson. Note that Ford-Fulkerson is a deterministic algorithm but its delays and overheads vary as the set of transactions varies between runs.

As expected, greedy embeddings led to shorter paths due to finding shortcuts between different branches of the tree. Hence, all settings using greedy embeddings exhibited lower delays and transaction overheads than the corresponding landmark-based protocols. Indeed, greedy embeddings reduced the path length and the transaction overhead by nearly a factor of 2. Greedy embeddings also increased the success ratio due to the shorter paths and the lower probability of encountering a link with low available credit.

The impact of the random assignment of credit was less clear-cut: While removing the need to involve landmarks into the routing process reduced the delay and the transaction overhead for all parameter settings, the impact on the success ratio differed between embedding-based and landmark routing. When combined with landmark routing, random credit assignments resulted in a definite drop in success from more than 60% to only 8%. The reason for the low success ratio was the high probability of encountering at least one link

<sup>2</sup>XRP is the symbol of the Ripple currency.

<sup>3</sup>We will provide them anonymously to the reviewers via the program chairs upon request.

TABLE I: Performance of different transaction schemes, varying the routing algorithm (T-Tree, GE-greedy embedding), the stabilization method (PER-periodic, OND-on-demand), the assignment of credit on paths (MUL-multi-party computation, RAND-random), and the landmark selection (HD - highest degree, RL-random landmark) for five metrics: success ratio: fraction of successful transactions (higher is better), delay: longest chain of messages (lower is better), transaction: messages sent per transaction (lower is better), path length: length of discovered paths between sender and receiver (lower is better), stabilization: messages for stabilizing the trees sent per epoch (lower is better). SilentWhispers corresponds to the setting LM-MUL-PER whereas SpeedyMurmurs is GE-RAND-OND.

Setting	Success Ratio	Delay	Transaction	Path Length	Stabilization
SilentWhispers-HD	$0.651 \pm 0.005$	$15.01 \pm 0.08$	$82.0 \pm 0.2$	$5.30 \pm 0.01$	$598722 \pm 0$
LM-MUL-OND-HD	$0.62 \pm 0.03$	$14.7 \pm 0.5$	$81 \pm 2$	$5.3 \pm 0.1$	$8000000 \pm 2000000$
LM-RAND-PER-HD	$0.09 \pm 0.01$	$8.3 \pm 0.1$	$35.1 \pm 0.5$	$3.23 \pm 0.05$	$598722 \pm 0$
LM-RAND-OND-HD	$0.09 \pm 0.09$	$9 \pm 1$	$37 \pm 4$	$3.4 \pm 0.4$	$2000 \pm 2000$
GE-MUL-PER-HD	$0.908 \pm 0.001$	$11.52 \pm 0.03$	$49.0 \pm 0.1$	$1.951 \pm 0.003$	$598722 \pm 0$
GE-MUL-OND-HD	$0.905 \pm 0.004$	$11.5 \pm 0.2$	$49.0 \pm 0.5$	$1.954 \pm 0.007$	$4000 \pm 4000$
GE-RAND-PER-HD	$0.913 \pm 0.001$	$6.016 \pm 0.009$	$18.30 \pm 0.04$	$1.867 \pm 0.003$	$598722 \pm 0$
SpeedyMurmurs-HD	$0.906 \pm 0.006$	$6.02 \pm 0.04$	$18.3 \pm 0.1$	$1.87 \pm 0.01$	$300 \pm 300$
TREEONLY-SW-HD	$0.863 \pm 0.003$	$15.9 \pm 0.1$	$81.9 \pm 0.3$	$3.17 \pm 0.01$	$598722 \pm 0$
TREEONLY-SM-HD	$0.54 \pm 0.04$	$6.7 \pm 0.3$	$23.5 \pm 0.7$	$2.01 \pm 0.07$	$5000 \pm 5000$
Ford-Fulkerson	$1.00 \pm 0.00$	$49500 \pm 900$	$49500 \pm 900$	$3.2 \pm 0.1$	$0 \pm 0$
SilentWhispers-RL	$0.1 \pm 0.2$	$15 \pm 2$	$130 \pm 10$	$7.6 \pm 0.6$	$598722 \pm 0$
SpeedyMurmurs-RL	$0.912 \pm 0.007$	$5.99 \pm 0.06$	$18.2 \pm 0.2$	$1.863 \pm 0.009$	$1000 \pm 1000$

with insufficient credit to satisfy the random assignment. In contrast, greedy embeddings exhibited much shorter paths and the flexibility to potentially choose between several neighbors. These two properties negated the disadvantageous impact of the random credit assignment, so that greedy embedding in combination with random assignment resulted in the same success ratio of 91% as in combination with multi-party computation.

On-demand stabilization reduced the stabilization overhead (abbreviated by Stabilization in Table I) drastically: While rebuilding the spanning trees periodically resulted in more than half a million messages per epoch, on-demand stabilization only required a few thousands of messages, as shown in the last column of Table I when comparing the algorithms using ‘PER’ and ‘OND’. On-demand stabilization induced high variance because the value of links close to the root of a spanning tree rarely drops to 0 but incurred enormous overhead in these rare occurrences. The simulation showed a clear advantage of on-demand stabilization. We admit that the considerable advantage of on-demand stabilization is partially due to lack of link value changes and actual dynamics in the static simulation. In the second part of this section, we therefore evaluate the stabilization overhead in a dynamic environment.

Second, we compare SilentWhispers and SpeedyMurmurs with tree-only routing and Ford-Fulkerson based on the results in Table I. As expected, Ford-Fulkerson exhibited prohibitive delays and transaction overheads. The fact that Ford-Fulkerson also results in a longer average path length seems counterintuitive at first. However, the result was a side effect of Ford-Fulkerson discovering long paths to maximize the available credit that the other approaches failed to discover. As illustrated in Fig. 1 and detailed in Sec. II, tree-only routing finds the shortest route in the spanning tree, possibly without passing a landmark, but does not include links that are not contained in the tree. Thus, tree-only routing is a compromise between SilentWhispers’ routing algorithm and embedding-based routing. As a consequence, the performance results

when using only tree links are in between the performance of SilentWhispers and SpeedyMurmurs.

Next, we evaluate the impact of different configuration parameters on the performance. As indicated in the last two rows of Table I, choosing random landmarks did not considerably affect the performance of SpeedyMurmurs but reduced the performance of SilentWhispers due to the existence of longer paths to a landmark with few connections. In contrast, increasing the number of trees  $|L|$  affected the success ratio of SpeedyMurmurs negatively and SilentWhispers positively, as Fig. 2a indicates. The reason for the observed decrease in success was the increased likelihood that at least one path did not have sufficient credit. An increased  $|L|$  further increased the delays, as shown in Fig. 2b. The impact was more pronounced for SilentWhispers due to the fact that landmarks had to wait until all messages for the multi-party computation arrived. The number of attempts  $a$  had a slight positive effect on the success ratio, as indicated by Fig. 2c. Yet, as the transaction overhead is linear in the number of attempts, the slight increase might not warrant multiple attempts.

For all algorithms but Ford-Fulkerson, the success ratio was considerably below 100%. It stands to reason that a lot of users might not be willing to accept a failure rate of 10% or more. Note that a failure to route does not reduce the funds of any user, so there is no loss in funds associated with a routing failure. Furthermore, in a non-static environment, users can retry the transaction at a later point in time after the network has sufficiently changed for it to work. If neither failure nor waiting is an option, we could apply Ford-Fulkerson on failure. By reducing the transactions that require Ford-Fulkerson to 10%, we still considerably improve the efficiency in comparison to a network relying exclusively on Ford-Fulkerson at the price of a slight increase in delay due to the preceding use of SpeedyMurmurs. In addition, we hope that with increasing popularity, both the connectivity of the PBT networks and the amount of available funds increase beyond the current state of the Ripple network, which is bound to entail a higher probability of success.




Fig. 2: SilentWhispers vs. SpeedyMurmurs: Impact of number of parallel trees and attempts at performing a transaction




Fig. 3: Comparing SpeedyMurmurs and SilentWhispers in a dynamic setting based upon Ripple transaction and link changes from 2013 to 2016 on per-epoch scale; success is computed as the ratio of the actual success ratio and the success ratio of the Ford-Fulkerson algorithm as a baseline; for c), we present moving averages over 50 epochs to increase readability

Last, we evaluated the impact of dynamics on the performance of SpeedyMurmurs and SilentWhispers. As stated above, the impact of dynamics is particularly of interest to decide if on-demand stabilization is indeed more efficient than periodic stabilization. To better comprehend the reasons underlying our results, Fig. 3a displays the number of transactions and link changes per epoch for the Ripple dataset. While the number of transactions did not vary greatly over the period of three years, link creations and modifications were frequent in some short intervals but rare during the remaining observation period. The frequency of link changes directly relates to the stabilization overhead of SpeedyMurmurs, as indicated by Fig. 3b. Whereas the stabilization overhead was usually below 100 messages per epoch, the overhead increased to about  $10^9$  messages during periods of frequent change. Note that only the first two of the four batches of link changes resulted in a drastically increased need for stabilization. After the first two batches, spanning trees had formed and new link additions mostly created shortcuts that did not require changes to the trees. In contrast, the stabilization overhead of SilentWhispers only depended on the number of edges in the network and hence increased as the graph grows over time. During intervals of frequent change, the stabilization overhead of SilentWhispers was considerably lower than SpeedyMurmurs' stabilization overhead. However, during 'normal' operation, SilentWhispers' stabilization overhead exceeded the overhead

of SpeedyMurmurs by more than 2 orders of magnitude. We evaluated the success in relation to Ford-Fulkerson and hence divided the actual success ratio of each epoch by the success ratio of Ford-Fulkerson for the corresponding epoch. As can be seen from Fig. 3c, the success could exceed 1 if an alternative routing algorithm exhibited a higher success ratio. Note that higher success ratios were indeed possible due to the fact that different routing algorithms resulted in different payments and hence different network states. Different network states implied a different set of possible transactions, so that a transaction could fail for Ford-Fulkerson but succeed for SilentWhispers or SpeedyMurmurs. In comparison, SpeedyMurmurs and SilentWhispers achieved similar success ratios for most of the time; however, at the end of the simulation interval, SilentWhispers outperformed SpeedyMurmurs. The sudden increase in success correlates with the addition or change of many links, as can be seen from Fig. 3a. The additional links increase the density of the graph, leading to shorter paths, and hence a higher success probability. The fact that SilentWhispers achieves a higher success ratio than SpeedyMurmurs could be due to the tree structure: SilentWhispers maintains breadth-first search trees whereas SpeedyMurmurs initially constructs breadth-first search trees but does not change the parent of a node if a new neighbor offers a shorter path to the root. The longer paths to the root could have negative effects on the probability of success. As the actual success ratio of all considered

algorithms is low during later epochs, e.g., frequently below 5%, the result might be an artifact of our dataset and post-processing method.

In summary, we answered our five initial research questions as follows:

- SpeedyMurmurs achieved a higher performance than SilentWhispers with regard to all considered metrics for the static scenario.
- On-demand stabilization and embedding-based routing had a positive effect on all 5 performance metrics. In contrast, the use of random credit assignment might decrease the success ratio slightly. However, when used in combination with the other two modifications, the effect was mostly negated.
- Ford-Fulkerson usually achieved a higher success than both SpeedyMurmurs and SilentWhispers. However, the algorithm resulted in an enormous transaction overhead, exceeding the overhead of the other algorithms by 2 to 3 orders of magnitude.
- An increased number of trees or attempts to perform a transaction did not considerably increase the success ratio of SpeedyMurmurs but incurred increased overheads.
- The evolution of the PBT network affects the performance of the proposed algorithms. Stabilization overhead and success ratio vary considerably depending on the frequency of transactions and link changes.

The dynamic evaluation suggests working on the design of an alternative spanning tree maintenance algorithm. In particular, the results raise the question of suitable criteria for dynamically switching between on-demand and periodic stabilization. Indeed, as SilentWhispers is more efficient during periods of frequent change but results in higher overhead otherwise, such a switching mechanism could further reduce the communication overhead and hence increase scalability.

## VI. RELATED WORK

Maximizing the set of possible transactions in a credit network is NP-hard [11]. Instead, many existing systems have opted for considering one transaction at a time and applying the max-flow approach [9] as a routing algorithm. Nevertheless, existing algorithms [8] run in  $O(V^3)$  or  $O(V^2 \log(E))$  time and hence do not scale to a growing number of users and transactions [26], [35].

Prihodko et al. recently proposed Flare [27], a routing algorithm for the Bitcoin Lightning Network, a network of Bitcoin payment channels among Bitcoin users that enables off-chain transactions [25]. In Flare, all nodes keep track of their  $k$ -neighborhood; i.e., nodes at a hop distance of at most  $k$  and all links between them. In addition, each node maintains paths to a set of nearby *beacon* nodes.

This routing algorithm reveals the weight of all links in the  $k$ -neighborhood, usually for  $k \geq 3$ . This results in a privacy concern as the weight of a link between two users is exposed to users other than those two. Furthermore, nodes spread all updates to the  $k$ -neighborhood, meaning each credit change results in possibly hundreds of messages, which is highly inefficient for frequent transactions and hence changes in available credit.

Canal [35] presents the first efficient implementation of tree-only routing applied to looking for paths in credit networks. A trusted central party computes the shortest paths in the spanning trees between sender and receiver. If these paths provide enough credit to settle a transaction, the routing terminates successfully. Otherwise, it fails. In the face of network dynamics, the central server re-computes spanning trees constantly.

This approach has severe privacy implications as a central server maintains the complete PBT network and can keep track of the route taken by each payment. Moreover, the central server constitutes a single point of failure.

PrivPay [21] increases the privacy of Canal by using trusted hardware at the central server. However, PrivPay relies on a similar landmark technique as Canal and is also a centralized solution, therefore the scalability is still low and the issue of a single point of failure remains unsolved. Additionally, the PrivPay paper introduces for the first time the notions of value privacy and sender/receiver privacy for payments in a credit network. In this work, we define the privacy notions for routing in a PBT network as a building block not only for credit networks but also for any PBT network.

SilentWhispers [15] uses landmark routing in a fully distributed credit network. Both sender and receiver send messages in the direction of the landmarks, which constitute rendezvous nodes. In other words, paths in SilentWhispers are concatenations of the sender’s path to a landmark and the path from said landmark to the receiver. All paths pass a landmark, even if sender and receiver happen to be in the same branch, potentially leading to performance issues. However, as we discuss throughout this paper, SpeedyMurmurs, the routing algorithm proposed in this work, outperforms the routing approach proposed in SilentWhispers while achieving the privacy notions of interest.

Malavolta et al. [16] recently proposed Rayo and Fulgor, two payment-channel networks (i.e., PBT networks) that provide a necessary tradeoff between privacy and concurrency. Their study of concurrency could be leveraged to extend the concurrency in SpeedyMurmurs. Nevertheless, they do not tackle the path selection problem. Therefore, SpeedyMurmurs becomes an excellent candidate to complement Rayo and Fulgor.

*Summary:* Existing routing approaches often disregard privacy. Most of them require centralization or shared public information; SilentWhispers is the only existing distributed system focusing on privacy. However, it relies on a distributed landmark routing technique that is potentially inefficient. Our in-depth performance and privacy evaluation shows that SpeedyMurmurs provides higher overall performance when compared to state-of-the-art routing approaches, while achieving the privacy notions of interest.

## VII. CONCLUSION AND FUTURE WORK

In this work, we designed SpeedyMurmurs, an efficient routing algorithm for completely distributed payment systems. Our extensive simulation study and analysis indicate that SpeedyMurmurs is highly efficient and achieves a high probability of success while still providing transaction privacy

against a strong network adversary. As transaction privacy is essential for PBT applications, SpeedyMurmurs is an ideal routing algorithm for credit networks and payment channel networks, as well as for emerging inter-blockchain protocols.

As our results indicate that on-demand and periodic stabilization are suitable for different phases of a PBT network's evolution, future work can extend upon our results by investigating the option of dynamically switching between on-demand and periodic stabilization.

#### ACKNOWLEDGEMENTS

This work benefited from the use of the CrySP RIPPLE Facility at the University of Waterloo.

#### REFERENCES

- [1] Raiden network. <http://raiden.network/>. Accessed July 2017.
- [2] Real-time gross settlement. [https://en.wikipedia.org/wiki/Real-time\\_gross\\_settlement](https://en.wikipedia.org/wiki/Real-time_gross_settlement). Accessed July 2017.
- [3] Ripple website. <https://ripple.com/>. Accessed July 2017.
- [4] Stellar website. <https://www.stellar.org/>. Accessed July 2017.
- [5] Frederik Armknecht, Ghassan O Karame, Avikarsha Mandal, Franck Youssef, and Erik Zenner. Ripple: Overview and Outlook. In *International Conference on Trust and Trustworthy Computing*, pages 163–180. Springer, 2015.
- [6] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. *On Scaling Decentralized Blockchains*, pages 106–125. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [7] Dimitri DeFigueiredo and Earl T. Barr. TrustDavis: A Non-Exploitable Online Reputation System. In *E-Commerce Technology*, 2005.
- [8] Yefim Dinitz. Dinitz's Algorithm: The Original Version and Even's Version. In *Theoretical Computer Science*. 2006.
- [9] Lester R Ford and Delbert R Fulkerson. Maximal Flow Through a Network. *Canadian journal of Mathematics*, 8(3), 1956.
- [10] Ethereum Foundation. Ethereum Project. <https://www.ethereum.org/>. Accessed July 2017.
- [11] Arpita Ghosh, Mohammad Mahdian, Daniel M. Reeves, David M. Pennock, and Ryan Fugger. Mechanism design on trust networks. In *Internet and Network Economics, Third International Workshop, WINE 2007, San Diego, CA, USA, December 12-14, 2007, Proceedings*, pages 257–268, 2007.
- [12] Andreas Hofer, Stefanie Roos, and Thorsten Strufe. Greedy Embedding, Routing and Content Addressing for Darknets. In *NetSys*, 2013.
- [13] Amrit Kumar, Clément Fischer, Shruti Tople, and Prateek Saxena. A traceability analysis of Monero's blockchain. Cryptology ePrint Archive, Report 2017/338, 2017. <http://eprint.iacr.org/2017/338>.
- [14] Lightning Network Team. Atomic cross-chain trading. [https://en.bitcoin.it/wiki/Atomic\\_cross-chain\\_trading](https://en.bitcoin.it/wiki/Atomic_cross-chain_trading). Accessed July 2017.
- [15] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. SilentWhispers: Enforcing Security and Privacy in Credit Networks. In *NDSS*, 2017.
- [16] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Srivatsan Ravi. Concurrency and Privacy with Payment-Channel Networks. In *CCS*, 2017.
- [17] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the 2013 Conference on Internet Measurement Conference, IMC '13*, pages 127–140, New York, NY, USA, 2013. ACM.
- [18] Andrew Miller, Malte Möser, Kevin Lee, and Arvind Narayanan. An empirical analysis of linkability in the Monero blockchain. *CoRR*, abs/1704.04299, 2017.
- [19] Alan Mislove, Gaurav Oberoi, Ansley Post, Charles Reis, Peter Druschel, and Dan S. Wallach. AP3: Cooperative, decentralized anonymous communication. In *ACM SIGOPS European Workshop, EW 11*, 2004.
- [20] Prateek Mittal and Nikita Borisov. Shadowwalker: Peer-to-peer anonymous communication using redundant structured topologies. In *ACM CCS*, pages 161–172, 2009.
- [21] Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Kim Pecina. Privacy Preserving Payments in Credit Networks: Enabling trust with privacy in online marketplaces. In *NDSS*, 2015.
- [22] Pedro Moreno-Sanchez, Muhammad Bilal Zafar, and Aniket Kate. Listening to Whispers of Ripple: Linking Wallets and Deanononymizing Transactions in the Ripple Network. *PoPETS*, 4/2016, 2016.
- [23] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. <http://fastbull.dl.sourceforge.net/project/bitcoin/Design%20Paper/bitcoin.pdf/bitcoin.pdf>, 2011.
- [24] Christos H Papadimitriou and David Ratajczak. On a Conjecture Related to Geometric Routing. In *Algorithmic Aspects of Wireless Sensor Networks*, 2004.
- [25] Joseph Poon and Thaddeus Dryja. The Bitcoin Lightning Network: Scalable Off-chain Instant Payments. <https://lightning.network/lightning-network-paper.pdf>, 2015.
- [26] Ansley Post, Vijit Shah, and Alan Mislove. Bazaar: Strengthening user reputations in online marketplaces. In *NSDI*, 2011.
- [27] Pavel Prihodko, Slava Zhigulin, Mykola Sahno, Aleksei Ostrovskiy, and Olaoluwa Osuntokun. Flare: An Approach to Routing in Lightning Network. [http://bitfury.com/content/5-white-papers-research/whitepaper\\_flare\\_an\\_approach\\_to\\_routing\\_in\\_lightning\\_network\\_7\\_7\\_2016.pdf](http://bitfury.com/content/5-white-papers-research/whitepaper_flare_an_approach_to_routing_in_lightning_network_7_7_2016.pdf), 2016.
- [28] Fergal Reid and Martin Harrigan. *An Analysis of Anonymity in the Bitcoin System*, pages 197–223. Springer New York, New York, NY, 2013.
- [29] Michael K. Reiter and Avi D. Rubin. Crowds: Anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1):66–92, November 1998.
- [30] Stefanie Roos, Martin Beck, and Thorsten Strufe. Anonymous Addresses for Efficient and Resilient Routing in F2F Overlays. In *INFOCOM*, 2016.
- [31] Stefanie Roos, Martin Beck, and Thorsten Strufe. VOUTE-Virtual Overlays Using Tree Embeddings. *CoRR*, abs/1601.06119, 2016.
- [32] Benjamin Schiller and Thorsten Strufe. GTNA 2.0-A Framework for Rapid Prototyping and Evaluation of Routing Algorithms. In *Summer Computer Simulation Conference*, 2013.
- [33] Stefan Thomas and Evan Schwartz. A Protocol for Interledger Payments. <https://interledger.org/interledger.pdf>, 2015.
- [34] P. F. Tsuchiya. The Landmark Hierarchy: A New Hierarchy for Routing in Very Large Networks. *SIGCOMM*, 18(4), 1988.
- [35] Bimal Viswanath, Mainack Mondal, Krishna P. Gummadi, Alan Mislove, and Ansley Post. Canal: Scaling Social Network-based Sybil Tolerance Schemes. In *EuroSys*, 2012.
- [36] XRP Chat. Interledger payment through 7 ledgers. <https://www.xrpchat.com/topic/5660-interledger-payment-through-7-ledgers/>, 2017. Accessed July 2017.