

SPRITE: Secure and Private Routing in Payment Channel Networks

Gaurav Panwar, Roopa Vishwanathan, George Torres, Satyajayant Misra

New Mexico State University

{gpanwar, roopav, gtorresz, misra}@nmsu.edu

ABSTRACT

Payment channel networks are a promising solution to the scalability challenge of blockchains and are designed for significantly increased transaction throughput compared to the layer one blockchain. Since payment channel networks are essentially decentralized peer-to-peer networks, routing transactions is a fundamental challenge. Payment channel networks have some unique security and privacy requirements that make pathfinding challenging, for instance, network topology is not publicly known, and sender/receiver privacy should be preserved, in addition to providing atomicity guarantees for payments. In this paper, we present an efficient privacy-preserving routing protocol, SPRITE, for payment channel networks that supports concurrent transactions. By finding paths offline and processing transactions online, SPRITE can process transactions in just two rounds, which is more efficient compared to prior work. We evaluate SPRITE’s performance using Lightning Network data and prove its security using the Universal Composability framework. In contrast to the current cutting-edge methods that achieve rapid transactions, our approach significantly reduces the message complexity of the system by 3 orders of magnitude while maintaining similar latencies.

CCS CONCEPTS

- Security and privacy → Distributed systems security; Security protocols.

KEYWORDS

Privacy preserving protocols, Payment channel networks, Secure pathfinding

ACM Reference Format:

Gaurav Panwar, Roopa Vishwanathan, George Torres, Satyajayant Misra. 2024. SPRITE: Secure and Private Routing in Payment Channel Networks. In *Proceedings of ACM Conference (Conference’17)*. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

1 INTRODUCTION

Researchers have been devising efficient techniques to make cryptocurrency transactions more scalable, e.g., Bitcoin currently processes around seven transactions per second, and Ethereum around

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference’17, July 2017, Washington, DC, USA

© 2024 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

thirty transactions per second [6, 18], compared to centralized payment systems, such as Visa Inc., which, at a conservative estimate, can support up to 1700 transactions per second [5]. For addressing this, Layer-2 protocols, such as payment channels have been proposed as a workaround [31, 36, 45, 46], where several thousands of transactions can be processed with minimal blockchain writes and with no changes required to the blockchain’s underlying consensus mechanism (unlike other approaches such as sharding and alternate consensus mechanisms [22, 28, 29, 37, 41]). Payment channels also help enable *microtransactions*, which allow users to send small amounts of money, e.g., 10^{-4} Bitcoin, but without incurring high blockchain transaction fees [30].

Overview of payment channels: Two parties, Alice and Bob open a payment channel by depositing a certain amount of cryptocurrency into an address on a blockchain controlled by both parties’ signing keys. Say, Alice deposits x coins, and Bob deposits y coins. Alice and Bob can conduct several transactions by exchanging authenticated messages, thus changing the distribution of the coins in the channel, but without writing anything to the blockchain. The net worth of the channel remains $x + y$ coins. At a mutually agreed-upon time, they can close the channel by writing a blockchain transaction that commits the final, authenticated distribution of the coins to the blockchain. The coins are paid to Alice and Bob per the final transaction. To facilitate transactions between two parties that may not have a payment channel currently open between them, decentralized payment channel networks (PCNs) that enable transitive payments have been proposed [34, 35, 39, 47], where two unconnected users can send/receive payments if there exists a path comprising of several users with payment channels between them.

Motivation: Layer-2 protocols such as PCNs are gaining widespread acceptance. Lightning Network, which is a popular PCN based on the Bitcoin blockchain, had over 6 million users and 28 million payment channels open between June 2021 to July 2022 [31, 44]. Peer-to-peer (p2p) transactions between users in PCNs are becoming increasingly common, e.g., in 2021, another popular PCN, Ripple, had 15 million unique p2p transactions annually, with a maximum path length of 43 hops [4, 7]. Routing protocols which help discover payment paths between sender and receiver are at the core of PCNs. There could exist several paths between a sender and receiver in a PCN with differing channel balances. Each hop on a path incurs a routing fee, hence longer paths cost more.

Routing in PCNs is fundamentally different from traditional network routing in both, intent and security/privacy requirements, hence network routing protocols cannot be trivially ported to PCNs. Assuming a network graph with nodes and weighted links connecting them, in regular network routing, the intent is to transmit data, not route payments. Transmitting data does not alter the state of the nodes, but routing payments changes nodes’ available link balances. In network routing, bandwidth capacities and router/switch

identities are usually not considered private information, whereas, in PCNs, transaction amounts and node identities need to be kept private from all other nodes in the network. Transmission range and physical distance between devices are factors in network routing, but not in PCNs. PCNs reside entirely at the application layer, unlike network protocols in communication networks. Hence network routing protocols cannot be trivially ported.

Maximum flow algorithms such as Ford-Fulkerson [23] or Goldberg-Tarjan [24] would require either source routing or an external centralized, trusted entity to compute routes, besides having a high path computation overhead of $O(|V||E|^2)$ and $O(|V|^3)$ respectively, in a graph $G(V, E)$. While distributed versions of shortest path algorithms such as Dijkstra's algorithm exist [11, 12], they incur a computational complexity of $O(|V|^2) + O(|V|)$, which makes their scalability to large PCNs challenging.

Robust, scalable, decentralized PCN routing protocols hold the promise of making cryptocurrency transactions faster, hence, designing secure and efficient PCN routing protocols is a challenging research problem of practical significance. Such protocols can be used for on-demand pathfinding and routing in other decentralized networks, particularly edge networks, which have high node turnover rates.

Table 1: Routing Protocols in PCNs

PCN Routing protocols	Privacy of nodes	Decentralized	Atomicty
FSTR [33]	✗	✗	✗
Eckey <i>et al.</i> [20]	✗	✓	✓
Auto tune [26]	✗	✗	✗
Kadry <i>et al.</i> [27]	✗	✗	✗
MPCN-RP [17]	✗	✗	✓
SilentWhispers [34]	✓	✗	✓
SpeedyMurmurs [47]	✓	✓	✗
BLANC [39]	✓	✓	✓
Coinexpress [53]	✗	✓	✓
Vein [25]	✗	✗	✗
Spider [49]	✗	✗	✗
Flash [52]	✗	✗	✓
Robustpay [55]	✗	✗	✓
Robustpay+ [56]	✗	✗	✓
Webflow [54]	✓	✓	✗
SPRITE	✓	✓	✓

Related Work: Several early PCN routing protocols were centralized where routing relied on trusted entities [34, 38, 51]. Some protocols did not support concurrency [34], while others chose paths without knowing whether the chosen path can satisfy a minimum asking amount [47]. Some routing protocols do source routing [35, 49] where a sender constructs the entire path from itself to the receiver, while many protocols do not consider security and privacy aspects [17, 20, 21, 25–27, 55–57]. We provide a comparison of other relevant PCN routing protocols with SPRITE in Table 1, where our comparison metrics are informed by our security/privacy goals. The protocol in [39], while satisfying our

three comparison metrics, has a very high communication overhead, where every transaction requires blockchain writes, which defeats the idea of off-chain PCNs. Real-world PCNs such as Lightning Network (LN) [13, 31] implement a gossiping routing protocol, where each node gossips with its peers to build a local map of the network. This has issues such as nodes not being able to validate information given by peers, and often not finding the shortest path.

Our Contributions: In this paper, we design a decentralized routing protocol for PCNs, SPRITE, which helps reduce trust assumptions, takes into account network dynamics, and preserves key security/privacy goals, while supporting concurrent transactions with short paths. We formally prove the security of SPRITE in the Universal Composability framework. We experimentally evaluate the performance of SPRITE using Lightning Network datasets and compare its performance with two other state-of-the-art schemes, on several network topologies. Our analysis shows that SPRITE performs significantly better over a wide array of quantitative and qualitative metrics while improving security and privacy.

Outline: In Section 2, we define our system and threat models, in Section 3, we give an overview of the workflow of SPRITE. In Section 4, we describe the protocols that constitute SPRITE. In Section 5, we give the security analysis of SPRITE. In Section 6, we describe our experiments, and in Section 7 we conclude the paper.

2 SPRITE SYSTEM MODEL

In this section, we discuss the basics of a PCN, the parties involved in SPRITE and system parameters.

A PCN fundamentally can be conceptualized as a graph with users representing vertices and edges representing the payment channels between users. Figure 1 shows four parties and three two-party channels. The crossed-out number next to each party's name denotes that party's original balance in the channel, while the number above it denotes the new balance. The directionality of the arrows denotes the direction in which a payment can be processed.




Figure 1: George sending 25 coins to Ron via two intermediaries Alice and Bob in a PCN.

2.1 Parties

1) **Routing nodes:** In SPRITE some nodes with high number of connections will serve as publicly identifiable *routing nodes* (RN), in exchange for a fee, and denote the set of RNs by \mathbb{RN} . RNs are already in use in real-world PCNs, such as Lightning Network as liquidity providers [32], we leverage them for routing. In SPRITE, RNs help facilitate transactions: broadly, we segment the path between the sender and the receiver, with each segment checkpointed by an RN. If RN_s and RN_r are the RNs closest to sender and receiver respectively, the payment from sender to receiver will progress as: sender $\rightarrow RN_s \rightarrow RN_1 \rightarrow \dots \rightarrow RN_n \rightarrow RN_r \rightarrow receiver$. The sender need only tell RN_s the identity of the destination RN_r , RN_s will find the shortest path to RN_r , who will, in turn, be contacted by the receiver. Consequently, node disconnections/failures or malicious activities

on a segment are addressed and mitigated locally on each segment, and the rest of the path stays unaffected. Nodes volunteer to be RNs, and RNs are financially incentivized to help route transactions. RNs periodically broadcast messages about the available liquidity on their links to nodes within a radius, $\text{hopMax}_{\text{RN}}$.

SPRITE does not require any special security assumptions on which entities can choose to be RNs, and accounts for malicious RNs in the system (discussed further in Section 2.4 and Section 5). In a given transaction, RNs involved do not know the identities of Alice, Bob, or any other nodes on the path (except intermediate RNs or their immediate neighbors). RNs do not have a privileged position from a monitoring standpoint, except RN_s and RN_r will know that somebody in their $\text{hopMax}_{\text{RN}}$ radius is the sender/receiver, respectively. Additionally, intermediate RNs will neither know the identities of, nor the distances to RN_s and RN_r for a given transaction. Alice and Bob are free to choose the RN_s and RN_r per transaction based on the RNs available in their respective *routingTables*. If an Alice does not receive a broadcast message from any RN (indicating that she is outside the $\text{hopMax}_{\text{RN}}$ radius of all RNs in the system), she would need to connect either directly to a RN by forming a new payment channel or connect to another node in the network which is within $\text{hopMax}_{\text{RN}} - 1$ hops of some RN. Since RNs are economically incentivized to facilitate transactions, we assume RNs will be online, but SPRITE's functioning will not be impacted by any specific RN(s) going offline.

2) **Perimeter nodes:** Perimeter nodes are nodes that are located closer to the boundary of an RN's broadcast area where the area is determined by radius $\text{hopMax}_{\text{RN}}$. The idea of using perimeter nodes is to enable RNs that are spaced across the network to be able to communicate, without having to establish direct connections with each other. Two RNs that are far apart and want to route a transaction just need to find a common perimeter node in their local routing tables, and can route payments using that node. Since we want to preserve the perimeter nodes' privacy from RNs, in SPRITE, perimeter nodes are only identified by nonces they generate. The perimeter nodes will send a unique nonce to any RN that they receive a broadcast message from. If two RNs receive the same nonce, then they know they can reach each other through the perimeter node that sent the nonce. RNs with overlapping neighborhoods may have several common perimeter nodes.

3) **Regular nodes:** Any node that is not a routing node or a perimeter node is a regular node. We assume all nodes are rational and will act in their best economic interests. We assume the sender and receiver in a transaction can exchange messages out-of-band with each other, but payments are routed through nodes on the PCN. We use the terms users and nodes interchangeably.

4) **Blockchain:** SPRITE can work with any permission-less blockchain, and does not rely on blockchain-specific constructs such as hash time lock contracts (HTLCs) used in the Bitcoin blockchain, or smart contracts which are supported only by Turing complete blockchains, such as Ethereum. The blockchain is only used for opening/closing payment channels, thus avoiding excessive write/validator fees.

2.2 Setup

When a node joins the PCN, it establishes payment channels with other nodes who offer to connect with it or accept its connection offer. A node needs to connect to at least one other node to be part of the PCN. Nodes only reveal their identities to peers that they share a channel with. In this paper, we refer to peers sharing a channel as neighbors. Every node's identity is represented by a keypair denoted by $(\text{VK}_i, \text{SK}_i)$, of which VK_i is revealed to its neighbors. RNs will need to make their identities, i.e., verification keys, known to all nodes in the PCN, so nodes can use them for routing transactions.

Cryptographic Primitives: A *sequential aggregate signature* is a cryptographic primitive in which a series of users sign a message, where the final signature is computed sequentially by each user who adds her signature on her message. We use sequential aggregate signatures [43] (defined in Appendix 8.1) to maintain the privacy of non-RN nodes in the network (no need for publicly registered signing keypair) while still allowing for authentication of broadcast messages during the bootstrap phase. Furthermore, this helps from an efficiency perspective, since only one final signature needs to be verified rather than a series of signatures.

2.3 System Parameters

Transactions in a PCN might on occasion fail, e.g., due to abrupt node disconnections and insufficient liquidity along a path, thus necessitating retries. We set the number of times a transaction can be retried after a failure as a system-wide parameter, maxRetries . We also assume that each node i maintains a local state where it stores the number of times each transaction is retried, specifically, it maintains an arithmetic counter for each transaction (txid_i), $\text{retry}.\text{txid}_i \in \mathbb{Z}^+, i \in \mathbb{Z}^+$. If $\text{retry}.\text{txid}_i == \text{maxRetries}$, any new messages about that transaction will be rejected so the transaction can be tried on other paths. After transaction txid_i has been completed, times out, or is revoked, $\text{retry}.\text{txid}_i$ is deleted.

Hops: We define five parameters used in SPRITE: $\text{hopMax}_{\text{RN}}$, hopMax , pathStretch , hopCount , and hopBand . $\text{hopMax}_{\text{RN}}$ is the maximum number of hops an RN's broadcast message travels, hence defining the RN's neighborhood. hopBand is used for determining the distance of perimeter nodes. For example, if node r is an RN, $\text{hopMax}_{\text{RN}}$ is set to 20 hops and hopBand is set to 3, then all nodes that are at 18, 19, and 20 hops away from node r act as perimeter nodes. $\text{hopMax}_{\text{RN}}$ and hopBand are set individually by RNs. hopMax is a dynamic parameter that denotes the maximum number of hops a transaction can travel in a given segment. It is set by the sender for a given segment based on the estimated hopCount in the sender's *routingTable*. pathStretch , set by the sender, denotes an absolute upper bound on hopMax and is intended to be used only in case of routing problems that call for transaction retries within a segment. hopCount at a given node denotes the number of hops traveled by a message up until that node.

Timers: Transactions in SPRITE have two phases, *hold* and *pay*, and their corresponding segment-specific timers, $\text{te}_1.\text{txid}$ and $\text{te}_2.\text{txid}$, are maintained by each node participating in a given transaction designated by txid . These are internal countdown timers that are maintained by each node locally and are used by the nodes individually to determine when they should timeout the given transaction

and retry on a different path. Since each segment in the *hold* phase terminates at an RN, timer te_1 is cleared by nodes in a segment after a successful *hold* phase when the receiving downstream RN responds with an acknowledgment message for the transaction. Else, nodes will retry the transaction's *hold* phase on another path in the given segment after te_1 expires. Timer te_2 is cleared by all nodes in a transaction segment after a successful *pay* phase when they receive an acknowledgment that the payment has concluded successfully in their segment. Else, if te_2 expires, then the transaction is retried for *hold* and *pay* phases in the given segment. In SPRITE, we consider te_1 and te_2 to be system parameters set based on current network statistics and dynamics.

Fees: Similar to prior works, we assume RNs get paid a fixed amount periodically, contributed to by other nodes, and do not impose routing fees for transactions.¹ An economic analysis of routing fee models and optimal routing fee design is an orthogonal problem.

2.4 Threat Model and Security/Privacy Goals

Adversary actions: An adversary can adaptively corrupt any subset of users, including regular nodes, perimeter nodes and RNs, upon which the corrupted nodes' channels will be controlled by the adversary. The adversary can cause the corrupted users to behave in arbitrarily malicious ways, including misrouting payments and/or disseminating false information. We do not consider any node dropping/ignoring routing requests as malicious behavior, since that just means the node does not wish to participate in a given transaction, and an alternate path has to be found.

Adversary goals: An adversary wants to know nodes' identities that are not its immediate neighbors, including sender/receiver identities, and/or make people lose money, i.e., violate the atomicity of transactions.

Privacy-preservation: No node, not even RNs, know the identities of the sender, receiver, or any non-RN intermediaries for routing transactions, thus preserving sender/receiver privacy. SPRITE does not require the topology of the network to be known by any participating node in the system, as is standard in topology-hiding PCNs.² We assume the adversary cannot corrupt *all* PCN users.

Security/Privacy goals:

1) *Privacy of nodes*: Nodes should not know the identities of any nodes beyond their neighbors and RNs, nor garner any information (number of channels or balances) about other nodes.

2) *Transaction privacy*: No node should know the identities of the sender, receiver or the intermediaries in a transaction, unless it shares a channel with them. It should also not know amounts transferred in transaction paths it is not a part of.

3) Atomicity: Either a payment goes through in its entirety or not at all, i.e., either all link weights along a transaction path get updated by the transaction amount or none at all. In other words, no

honest party should lose credits because of the malicious behavior of other parties in the network.




Figure 2: Example of SPRITE protocol

3 WORKFLOW OF SPRITE

In this section, we provide an example run-through of the SPRITE protocol using Figure 2 where Alice is the sender and Bob is the receiver for a transaction. For presentation clarity, we do not pictorially depict multiple intermediary nodes between each of the parties in Figure 2. The bootstrap phase is used by RNs in the system to broadcast update messages that help nodes in their vicinity build routing tables. At the end of the broadcast phase, each node in the network will have a local routing table that indicates which RNs are reachable and through which of the node's neighbors. The nodes' routing tables also have estimates about the hop count and liquidity available to the corresponding RNs. The bootstrap phase also allows RNs to obtain information about what RNs are in their adjacent neighborhoods, the perimeter nodes that connect them, and how to reach RNs that are not in the adjacent neighborhoods.

When a transaction needs to occur, Alice and Bob coordinate out-of-band to confirm their closest reachable RNs (Charlie for Alice and Denise for Bob). From their routing table estimates, Alice and Bob decide the transaction amount based on the estimated liquidity available between Alice-Charlie and Denise-Bob, according to Alice's and Bob's corresponding routing tables. Alice sends a *hold_s* message to Charlie via one of her neighbors and this message is passed on by each node along the path including Mikaela, until it reaches Charlie (Figure 2, Steps 1-2 on Alice-Charlie segment). Simultaneously, Bob sends a *hold_d* message towards Denise through Hu (Figure 2, Steps 1-2 on Bob-Denise segment). Along the path, all nodes create pair-wise multisig hold contracts with their neighbors to reserve the transaction amount and set some local variables including hold phase timer (te_1) and pay phase timer (te_2).

When the messages reach Charlie and Denise, they reply with *holdACK* messages so that all nodes receiving the *holdACK* message clear their local te_1 timers and will no longer timeout and retry another path (Figure 2, Steps 3-4 in Alice-Charlie segment and Bob-Denise segment). Additionally, Charlie updates Alice's message so that it can be routed within the network through any intermediate RNs (Rajiv in this case) and is finally received by Denise (Figure 2,

¹In real-world PCNs such as LN, routing nodes currently get paid the same as other nodes, although there are proposals to update the fee structure [8–10].

²In LN, although edited snippets of the topology are made available for research purposes [19], one cannot extract the full network topology, as nodes' channel balances are not made public. Further, each payment channel funding transaction is a Pay-to-Witness-Script-Hash (P2WSH) address, and the nature of the script (a 2-of-2 multisig) will only be revealed once the funding transaction output is spent. Even if this were known/guessed, not all 2-of-2 multisig scripts on the Bitcoin blockchain correspond to payment channels. Finally, signing/verification keys are rotated by nodes for every channel (see [13]).

Steps 3-5 on Charlie-Rajiv and Rajiv-Denise segments). The message is updated by each perimeter node (Larry) and RN (Rajiv) on the path to facilitate forwarding the message towards Denise.

All nodes on the Alice to Denise path also set corresponding te_1 and te_2 timers (Figure 2, Steps 3-4 in Charlie-Rajiv segment and Step 5 in Rajiv-Denise segment) which are cleared when the corresponding RN in that segment is reached (Figure 2, Steps 5-6 in Charlie-Rajiv segment and Step 6 in Rajiv-Denise segment). In Figure 2, all nodes between Charlie and RN Rajiv, including perimeter node Larry will clear their te_1 timers after they receive a *holdACK* from Rajiv (Figure 2, Steps 6-8 on Charlie-Rajiv segment) and nodes between Rajiv and Denise will clear their timers when *holdACK* from Denise is received (Figure 2, Steps 6 on Rajiv-Denise segment). When Denise receives the two *hold_r* and *hold_s* messages, she sends Bob a *proceedPay* message (Figure 2, Steps 6-7 on Bob-Denise segment). On receiving *proceedPay*, Bob creates a *pay* message and sends it towards Denise (Steps 8-9 on Bob-Denise segment), which is then forwarded towards Charlie through intermediate RNs (Steps 10-12 on Charlie-Rajiv and Rajiv-Denise segment), and finally to Alice (Steps 13-14 on Alice-Charlie segment).

Each RN on the path replies with a *payACK* message when it receives a *pay* message and thus clearing timer te_2 for all nodes receiving the *payACK* message (Denise's *payACK* represented by Steps 10-11 on Bob-Denise segment, Rajiv's *payACK* represented by Step 11 on Rajiv-Denise segment, and Charlie's *payACK* represented by Steps 13-14 on Charlie-Rajiv segment). Finally, Alice sends out her own *payACK* when she receives the *pay* message (Steps 15-16 on Alice-Charlie segment), clearing the te_2 timers for nodes in the last segment, thus concluding the transaction.

4 CONSTRUCTION OF SPRITE

In the current Lightning Network, most new nodes connect to highly connected nodes in the network. This leads to a high concentration of nodes connected directly or with low hopcounts to well-connected nodes (RNs). This setup does not provide sender/receiver privacy from the highly connected nodes and there is the danger of highly connected nodes' link balances getting depleted quickly. Furthermore, if any RNs get disconnected or go offline, many other nodes would get disconnected from the network. In a network similar to Lightning, where RN nodes are closely located in terms of hop count, an RN-to-RN broadcast algorithm, which we refer to as R2RB (Algorithm 9) and define in Appendix 8.2, would work well.

However, if a PCN is built from the ground up with transaction security and node privacy as the focus, it is easy to assert that nodes would not necessarily always set up payment channels directly with well-known nodes (RNs) in the network since this would make the RN their next-hop neighbor and thus leak their identity as well as all their transactions' details. In a truly distributed network, new nodes would join other nodes in the periphery that they trust and not just RNs. In a system where RNs are located further apart, R2RB suffers from high message complexity due to long distances for *RN-Update* broadcast messages. We developed Algorithms 1, 2, henceforth referred to as R2NB, which reduces the distance each RN broadcasts to during the Setup phase, thus reducing the message complexity and adding to the efficiency of our scheme. The *hold* and *pay* phases remain the same for both approaches.

In practice, the first bootstrap phase in a given PCN will involve tuning of the $hopMax_{RN}$ parameter by the RNs to get an optimal overlap of perimeter nodes between neighboring RNs. The $hopMax_{RN}$ parameter is only used during the bootstrap phase of SPRITE and helps in limiting the number of broadcast messages from each RN; it is not used during a transaction. When new nodes join the network, they will receive their neighbors' *routingTables* regardless of their distance from any given RN and thus will join the neighborhood of the RN(s) that their neighbors occupy.

4.1 Bootstrap phase

This phase is described in Algorithm 1 and Algorithm 2. In the bootstrap phase the RNs first broadcast messages in the PCN within $hopMax_{RN}$ hops, advertising their available liquidity. The goal is to make nodes within $hopMax_{RN}$ aware that they can reach the respective RN, and help them construct their local routing tables.

RN broadcast to bootstrap neighborhood (Algorithm 1): In Algorithm 1, Lines 2-5, each RN k sets up the public parameters of an aggregate signature scheme, ppk and creates an aggregate signature keypair for itself (sk_k, vk_k). This is so all nodes in the RNs neighborhood can set up pseudonymous keypairs to hide their identity while propagating messages. It then composes an update message m_k , to be sent to all its neighbors. The message m_k contains k 's available liquidity in the outgoing direction, $currMax_s^k$, liquidity in the incoming direction $currMax_r^k$, and its real identity, VK_k . It also sets $hopCount$ to be zero and sets the $hopBand$. Perimeter nodes will be the farthest nodes from k in the band defined by nodes lying between $hopMax_{RN}$ hops and $(hopMax_{RN} - hopBand)$ hops from k . Each RN k can set its $hopBand$ independently. RN k timestamps and signs the message m_k using the signing key tied into its real identity, and produces a signature, σ'_k . It then again signs σ'_k and m_k using its aggregate signature signing key and creates an aggregate signature, σ_k , which is sent to k 's neighbors.

In Line 6, each node i within $hopMax_{RN}$ receives a set of messages (m'_k, \dots, m_j) and a set of verification keys (vk_k, \dots, vk_j) and a single aggregate signature σ_j which represents the aggregate signature of all nodes along the path from RN k to node j . Node i will then verify the signature, perform other checks (Lines 6-10), and update the values of $currMax_s^k$ and $currMax_r^k$ in its local routing table (Line 12). If node i is a non-perimeter and non-RN node, it then composes a new *RN-Update* message to forward to its neighbors. It increments the $hopCount$ by one, computes the new values of $currMax_s^k$, $currMax_r^k$ based on its local channel balances, appends its message to the message list and generates an aggregate signature on the appended list. It then sends the updated *RN-Update* message to its neighbors (Lines 19-22).

If node i happens to be a perimeter node (Lines 13-16) based on the $hopCount$ of the received message, it generates a nonce $Nonce_i$. It creates an *RN-UpdateReply* tuple that includes $Nonce_i$, updated values of $currMax_s^k$, $currMax_r^k$, and $hopCount$, and sends it to its previous node towards RN. All perimeter nodes also forward the *RN-Update* message until it reaches the node(s) at $hopMax_{RN}$ hops, who will send a reply but not broadcast the message further.

When nodes receive an *RN-UpdateReply* tuple, they act differently depending on whether they are an RN or a regular node. If the receiving node is an RN, then the message has traveled to

Algorithm 1: R2NB: Bootstrap broadcast from RN to perimeter nodes

```

1 Each node  $i$  initializes a table,  $routingTable$ , containing
  columns: (reachable RNs, next hop neighbor  $j$ ,  $currMax_s$ ,
   $currMax_r$ ,  $hopCount$ ,  $t_e$ ).
2 for each RN,  $k \in \mathbb{RN}$  do
3    $k$  does  $AS.Setup(1^\lambda) \rightarrow pp_k$  and runs
      $AS.KeyGen(pp_k) \rightarrow (sk_k, vk_k)$ .
4   Create a tuple  $m_k = (RN-Update, pp_k, VK_k, currMax_s^k,$ 
     $currMax_r^k, hopCount = 0, hopBand, hopMax_{RN}, ts)$ 
    for each neighbor  $j$ ,  $j \in [1..l]$  where  $l$  is the total
    number of neighbors of  $k$ . Create  $\sigma'_k \leftarrow \text{Sign}(SK_k, m_k)$ 
    and set  $m'_k = (m_k, \sigma'_k)$  Create signature
     $\sigma_k \leftarrow AS.Sign(sk_k, \perp, \perp, \perp, m'_k)$ .
5   return  $M = ((m'_k), (vk_k), \sigma_k)$  to each neighbor  $j$ .
6 for each node  $i$  in the network on receiving an  $RN-Update$ 
  message from neighbor  $j$  do
7   On receiving  $M = ((m'_k, \dots, m_j), (vk_k, \dots, vk_j), \sigma_j)$ ,  $i$ 
     parses  $(m_k, \sigma'_k) \leftarrow m'_k$  and  $(RN-Update, pp_k, VK_k,$ 
      $currMax_s^k, currMax_r^k, hopCount, hopBand,$ 
      $hopMax_{RN}, ts) \leftarrow m_k$ .
8   if  $(Verify(m_k, VK_k, \sigma'_k) \rightarrow 0) \vee$ 
      $(AS.Verify((m'_k, \dots, m_j), (vk_k, \dots, vk_j), \sigma_j) \rightarrow 0)$ 
     then
9     | Return  $\perp$ .
10     $i$  checks that  $hopCount$  value in all messages
      $(m'_k, \dots, m_j)$  are incremented by 1 in each message. If
     not, return  $\perp$ .
11     $i$  runs  $AS.KeyGen(pp_k) \rightarrow (sk_i, vk_i)$ .
12     $i$  updates its local  $routingTable$  for RN  $k$  and neighbor  $j$ 
     by updating the expiry time  $t_e = currTime + e$ ,
      $currMax_s^k$ , and  $currMax_r^k$ .
13    if  $((hopMax_{RN} - hopBand < hopCount) \wedge (hopCount \leq hopMax_{RN}))$  then
14      | Create a nonce  $Nonce_i \leftarrow \{0, 1\}^\lambda$ .
15      | Create return message  $mr'_i$  by updating contents of
         $m_j$  as  $mr'_i = (RN-UpdateReply, \cdot, \cdot, currMax_s^k,$ 
         $currMax_r^k, hopCount, \cdot, \cdot, \cdot, Nonce_i)$  where
         $hopCount = hopCount + 1$ ,
         $currMax_s^k = min(currMax_s^k, lw_{j,i})$ , and
         $currMax_r^k = min(currMax_r^k, lw_{i,j})$ .
16      |  $i$  creates signature  $\sigma_i \leftarrow AS.Sign(sk_i, \sigma_j, (m'_k,$ 
         $\dots, m_j), (vk_k, \dots, vk_j), mr'_i)$ .  $i$  sends  $MR = ((m'_k,$ 
         $\dots, m_j, mr'_i), (vk_k, \dots, vk_j, vk_i), \sigma_i)$  back to
        neighbor  $j$ .

```

the perimeter nodes and back. The receiving RN will store the reply information sent by the perimeter nodes in its $RNroutingTable$, indexed by the Nonce value sent by the perimeter node (Line 27-29). On the other hand, if the node receiving $RN-UpdateReply$ is a non-RN node, then it updates its local routing table again with the received information ($currMax_s^k$, etc.), adds the perimeter node's nonce to its local routing table, computes new values of

Algorithm 1: R2NB: Bootstrap broadcast from RN to perimeter nodes (continued)

```

17   if  $hopCount \geq hopMax_{RN}$  then
18     | Return  $\perp$ .
19   for each neighbors do
20     |  $i$  creates  $m_i$  by updating contents of  $m_j$  as
        $hopCount = hopCount + 1$ ,
        $currMax_s^k = min(currMax_s^k, lw_{i,s})$ , and
        $currMax_r^k = min(currMax_r^k, lw_{s,i})$ .
21     |  $i$  creates signature  $\sigma_i \leftarrow AS.Sign(sk_i, \sigma_j, (m'_k, \dots,$ 
        $m_j), (vk_k, \dots, vk_j), m_i)$ .
22     |  $i$  sets  $M = ((m'_k, \dots, m_j, m_i), (vk_k, \dots, vk_j, vk_i), \sigma_i)$ 
       and return  $M$  to neighbor  $s$ .
23 for each node  $j$  in the network on receiving an
   $RN-UpdateReply$  message from neighbor  $o$  do
24   On receiving  $MR =$ 
     $((m'_k, \dots, mr'_i, \dots, mr'_o), (vk_k, \dots, vk_i, \dots, vk_o), \sigma_o)$ .
25   if  $(AS.Verify((m'_k, \dots, mr'_o), (vk_k, \dots, vk_o), \sigma_o) \rightarrow 0)$ 
     then
26     | return  $\perp$ .
27   if  $j$  is the RN  $k$  then
28     | if  $(hopMax_{RN} - hopBand) \leq$ 
29       |  $|\{m'_k, \dots, mr'_i, \dots, mr'_o\}|/2 \leq hopMax_{RN}$  then
30       | | Add  $(Nonce_i, o, \cdot, \cdot, \cdot)$  to  $RNroutingTable_k$ .
31     | else
32       | Add  $Nonce_i$  and neighbor  $o$  to  $routingTable$ .
33       | Update contents of  $mr'_o$  as  $mr'_j = (RN-UpdateReply, \cdot,$ 
          $\cdot, currMax_s^k, currMax_r^k, hopCount, \cdot, \cdot, \cdot, Nonce_i)$ 
         where  $hopCount = hopCount - 1$ ,
          $currMax_s^k = min(currMax_s^k, lw_{j,o})$ , and
          $currMax_r^k = min(currMax_r^k, lw_{o,j})$ .
34       |  $j$  creates signature  $\sigma_j \leftarrow AS.Sign(sk_j, \sigma_o, (m'_k, \dots,$ 
          $mr'_o), (vk_k, \dots, vk_o), mr'_j)$ .
35       | Forward message  $MR = ((m'_k, \dots, mr'_o, mr'_j), (vk_k,$ 
          $\dots, vk_o, vk_j), \sigma_j)$  to neighbor from who
          $RN-Update$  message of  $k$  with timestamp  $ts$  was
         received.

```

$currMax_s^k$, $currMax_r^k$, decrements $hopCount$, and sends the signed message to the neighbor from whom it received the corresponding $RN-Update$ (Lines 31-34). Here e is the system-wide parameter for depicting the time duration after which a record is considered expired/stale in nodes' $routingTable$. In case node i had received the same message tuple with a lower $hopCount$ earlier, it drops the message to avoid loops. A possible optimization is nodes updating $currMax_s^k$, $currMax_r^k$ only once, instead of twice, i.e., on receipt of the $RN-Update$ tuple (Line 20) and not again after receipt of the $RN-UpdateReply$ tuple (Line 32). New nodes joining the PCN get $routingTables$ from their neighbors as soon as they join and will participate in $RN-Update$ broadcasts in the next time epoch. No

Algorithm 2: RNs exchanging nonces

- 1 Each $\text{RN}_i \in \text{RN}$ creates a table with rows $(\text{Nonce}_m, j, \cdot, \cdot, \cdot, \cdot)$, where j is the neighbor RN_i received Nonce_m from. Let \mathbb{N}_i be the set of all nonces obtained by RN_i .
- 2 Each RN_i then picks $\alpha \leftarrow \mathbb{Z}_p$, picks $d \in \mathbb{Z}^+$, and creates set $\mathcal{R}_i = \{r_i, \forall i \in [1..d]; r_i \leftarrow \{0, 1\}^\lambda\}, d = |\mathcal{R}_i|$. RN_i then sets $N_i = \mathbb{N}_i \cup \mathcal{R}_i$.
- 3 RN_i sends \mathbb{N}_i to all $\text{RN}_j \in \text{RN} \setminus \text{RN}_i$.
- 4 Each RN_i computes $\mathbb{N}_{ij} \leftarrow \mathbb{N}_i \cap \mathbb{N}_j$ for all $\text{RN}_j \in \text{RN} \setminus \text{RN}_i$, and builds its RNroutingTable locally.

re-calculation or broadcasts happen when new nodes join the network. For highly dynamic networks, the epoch value can be tuned or lowered so that the *RN-Update* broadcast messages account for significant changes in the topology. The cost of the *RN-Update* bootstrap phase is similar across epochs and depends on the current size of the network during the broadcast.

Algorithm 3: Alice- $\text{RN}_s - \dots - \text{RN}_r$ hold segments

- 1 Alice picks RN_s and Bob picks RN_r . Bob sets $\text{preimage} \leftarrow \{0, 1\}^\lambda$ and $\text{digest} = H(\text{preimage})$, and shares digest with Alice.
- 2 Let v be the amount of credits Alice wishes to send to RN_s . Alice picks token , $\text{preimage}_{txid} \leftarrow \{0, 1\}^\lambda$, $txid = H(\text{preimage}_{txid})$, and sends $txid$ to Bob.
- 3 Alice does $C_{\text{RN}_r} = E_{PK_{\text{RN}_r}}(\text{token}, v, txid)$ and $C_{\text{RN}_s} = E_{PK_{\text{RN}_s}}(VK_{\text{RN}_r}, v, txid, C_{\text{RN}_r})$.
- 4 Alice looks up her *routingTable* and picks a tuple $(\text{RN}_s, \text{node}_k, p_k)$, with $p_k = (\text{hopCount}, \text{currMax}_s, \text{currMax}_r, t_e)$ where $\text{currMax}_s \geq v$ and sets $\text{hopMax} = \text{hopCount} + \text{pathStretch}$. Alice creates a tuple $(\text{hold}_s, \text{RN}_s, VK_{\text{RN}_s}, v, txid, C_{\text{RN}_s}, \text{hopMax}, \text{digest}, t_{e1}, t_{e2})$ and sends it to node_k .
- 5 **for** Each node (node_i) in the network **do**
- 6 | Follow Algorithm 5

RNs exchanging nonces (Algorithm 2): After the PCN is bootstrapped, the RNs need to setup their local *RNroutingTables* which will help them find other RNs. At the end of Algorithm 1, each RN i would have received *RN-UpdateReply* tuples of the form $(\text{Nonce}_m, \cdot, \cdot, \cdot, \cdot, \cdot)$ from its neighbors, where m is a perimeter node within i 's hopMax_{RN} radius. RN i will receive several tuples containing nonces, we represent the set of unique nonces that i receives by \mathbb{N}_i (Line 1). RN i then pads the set \mathbb{N}_i with random strings and generates a larger set \mathcal{R}_i (Line 2). This is to ensure that other RNs cannot guess the size of \mathbb{N}_i , thus preserving privacy. SPRITE not only hides the identity of the perimeter nodes against all RNs in the system using the randomly generated nonces by perimeter nodes, but also hides the *number* of perimeter nodes each RN has within its hopMax_{RN} radius. All RNs then exchange their nonce sets and each RN finds the intersection of its set with other RNs' sets (Line 4). If even the nonce values need to be hidden for any reason, we can use more involved protocols such as private set intersection [42].

Algorithm 4: Bob – RN_r hold segment

- 1 Bob generates $C'_{\text{RN}_r} = E_{PK_{\text{RN}_r}}(\text{token}, v, txid)$.
- 2 Bob looks up his *routingTable* and picks a tuple $(\text{RN}_r, \text{node}_k, p_k)$, with $p_k = (\text{hopCount}, \text{currMax}_s, \text{currMax}_r, t_e)$ where $\text{currMax}_r \geq v$ and sets $\text{hopMax} = \text{hopCount} + \text{pathStretch}$. Bob creates a tuple $(\text{hold}_r, \text{RN}_r, VK_{\text{RN}_r}, v, txid, C'_{\text{RN}_r}, \text{hopMax}, \text{digest}, t_{e1}, t_{e2})$ and sends it to node_k .
- 3 **for** Each node (node_i) in the network **do**
- 4 | Follow Algorithm 5

Determining t_{e1} and t_{e2} values: After Algorithm 2, RNs help senders determine t_{e1} and t_{e2} values for their transactions. A low value for t_{e1} and t_{e2} could result in premature timeout of a transaction when waiting a little longer would have resulted in the transaction completing successfully. t_{e1} and t_{e2} also shouldn't be so large that the liquidity in the network is locked up despite there being no viable paths via the involved RNs. The value of t_{e1} can be

Algorithm 5: Subroutine for every node for *hold* and *pay* phase

Each node (node_i):

Case 1: on receiving hold_x message, $x \in \{s, r\}$, $\text{msg} = (\text{hold}_x, Y, VK_{\text{RN}_{(\cdot)}}, v, txid, C_{\text{RN}_{(\cdot)}}, \text{hopMax}, \text{digest}, t_{e1}, t_{e2})$, calls $\text{hold}(\text{msg})$ defined in Algorithm 8.

Case 2: on receiving holdReject_x message $\text{msg} = (\text{holdReject}_x, Y, VK_{\text{RN}_{(\cdot)}}, v, txid)$ along with *routingTable* update, calls $\text{holdReject}(\text{msg})$ defined in Algorithm 8.

Case 3: on receiving holdACK_x message $\text{msg} = (\text{holdACK}_x, t, \sigma_{\text{RN}_{(\cdot)}})$ along with *routingTable* update, calls $\text{holdACK}(\text{msg})$ defined in Algorithm 8.

Case 4: that did not receive a holdACK_x tuple for a transaction $txid$, and current time $> t_{e1}$, calls $\text{holdACKTimeout}()$ defined in Algorithm 8.

Case 5: on receiving *pay* message $\text{msg} = (\text{pay}, \text{preimage}, v, txid)$, calls $\text{pay}(\text{msg})$ defined in Algorithm 8.

Case 6: on receiving *payACK* message $\text{msg} = (\text{payACK}, \cdot, \cdot)$, calls $\text{payACK}(\text{msg})$ defined in Algorithm 8.

Case 7: that did not receive a *payACK* tuple for a transaction $txid$, and current time $> t_{e2}$, calls $\text{payACKTimeout}()$ defined in Algorithm 8.

set by the sender based on a sampling of communication times with its next-hop neighbors. For setting the value of t_{e2} , each RN can estimate the communication time to its neighboring RNs, based on an estimate of number of hops per neighborhood and its estimated t_{e1} ; this can be built into the routing protocol with little overhead. This information can be broadcasted by RNs in their neighborhood (as part of the routing messages). When a sender sets the transaction's t_{e2} , they can use the aggregate statistic of t_{e2} values they receive from their RN, e.g., 3 times the aggregate t_{e2} . We assume a certain amount of trial and error in finding the right multiplier on the part of the sender.

4.2 Hold phase

This is the first phase of transaction processing. In this phase, all nodes along a path from Alice to Bob will reserve or "hold" the amount Alice wishes to send to Bob. For ease of discussion, we divide the path into three segments, Alice – RN_s , RN_s – RN_r , and Bob – RN_r . Since the Alice- RN_s hold segment (Algorithm 3) and Bob- RN_r hold segment (Algorithm 4) are self-explanatory, due to space constraints, we describe them in the full version [40].

Hold phase and Pay phase functions for intermediate nodes (Algorithm 5): This algorithm depicts the functions called by different nodes, i.e., regular/perimeter nodes and RNs, when they receive different messages during a SPRITE transaction (full details of the functions are in Appendix 8.2, Algorithm 8.) Let us now discuss when/why these functions are called by various nodes.

The *hold* function is called by a node on receiving a $hold_r$ or $hold_s$ message. The node checks its routing table and decides which neighbor the *hold* message needs to be forwarded to in order to route it to the target RN in the message. If no viable paths are available then the current node would forward a *holdReject* message to the neighbor from which it received the *hold* message originally. If a node in the network receives a *holdReject* message then it uses the *holdReject* function to process the message and make a decision about whether it should retry on other available paths or forward the *holdReject* message back in the direction of the sender.

holdACK and *payACK* functions are called by nodes in the network on receiving *holdACK* or *payACK* messages, respectively. These functions involve the verification of the received acknowledgment messages and forwarding them toward the sender on the transaction path. If a node in the network does not receive a *holdACK* or *payACK* message during the *hold* and *pay* phases, respectively, and the timers expire (te_1 for *hold* phase and te_2 for *pay* phase), then the respective nodes call the timeout functions, *holdACKTimeout* for *hold* phase and *payACKTimeout* for *pay* phase.

Hold phase and Pay phase RNs' actions (Algorithm 6): We now discuss how the RNs handle operations in the hold phase, described in Algorithm 6. We recollect that RN_s is the first RN in the path, and RN_r is the last one. When RN_s receives a $hold_s$ message from Alice, it retrieves the verification key of RN_r (Line 3). RN_s then constructs an onion consisting of successive encryptions for all the RNs, $\{RN_l, \dots, RN_k\}$ between RN_s and RN_r , with RN_r being the innermost layer of the onion. RN_s sends the onion to its next-hop neighbor along the path to RN_l (Line 3-8). Note that the intended recipient is the perimeter node common to RN_s and RN_l (since RN_s is not within $hopMax$ distance of RN_l). RN_s also sends a signed *holdACK* message to Alice whom it received the $hold_s$ message from (Line 9). This is done to give the sender assurance that RN_s has received her message, but without requiring any blockchain writes. If malicious nodes drop *holdACK* messages, Alice will re-send the $hold_s$ tuple after a timeout.

The honest intermediaries along the path will recognize the $hold_s$ message with the same $txid$ as a duplicate and will re-send the old, stored *holdACK* message along a different path. When RN_r receives the $hold_s$ tuple, she sends a *holdACK* tuple to Alice. Similarly, RN_r also sends a signed *holdACK* message back to Bob (Line 11-14). When an intermediate RN that is part of the onion created by RN_s receives $hold_s$, it peels off its layer, finds the identity of the next RN

Algorithm 6: RN operations in *hold* and *pay* phase.

```

1  if hold phase then
2    if nodei == RNs then
3      RNs on receiving (holds, RNs, VKRNs, v, txid, CRNs, hopMax, digest, te1, te2) tuple from a neighbor, does
4        mRNs ← DSKRNs(CRNs) where
5        mRNs = (VKRNr, v, txid, CRNr).
6        RNs looks up RNroutingTable to find a path
7          (RNk, RNk+1, ..., RNl) to RNr.
8        RNs creates mhold = (VKRNr, v, txid, CRNr)
9        for RNi in {RNl, ..., RNk+1, RNk} do
10          RNs does mhold = (PKRNi, v, txid,
11            EPKRNi(mhold))
12        RNs then sends (holds, Y, VKRNl, v, txid, mhold,
13          hopMax, digest, te1, te2), to its neighbor towards Y
14          according to RNroutingTable for selected path to Y
15          with hopMax = hopCount of the path.
16        RNs does t = (txid, holds, v), σRNs ←
17          Sign(skRNs, t), sends (holdACKs, t, σRNs) along
18          with local routingTable to neighbor that sent holds.
19      else if nodei == RNr then
20        if message is holds then
21          When RNr receives the holds message, then the
22            Alice – RNr segment is complete. RNr does
23            t = (txid, holds, v), σRNr ← Sign(skRNr, t),
24            sends (holdACKs, t, σRNr) along with local
25            routingTable to neighbor that sent holds.
26        else
27          When RNr receives the holdr message, then the
28            RNr – Bob segment is complete. RNr sends
29            t = (txid, holdr, v), σRNr ← Sign(skRNr, t),
30            sends (holdACKr, t, σRNr) along with local
31            routingTable to neighbor that sent holdr.
32      else if nodei = RNi, ∀RNi ∈ [RNk, RNk+1, ..., RNl]
33      then
34        RNi on receiving the tuple (holds, RNi, VKRNi, v,
35          txid, mhold, hopMax, digest, te1, te2) parses
36          mhold = (PKRNi, CRNi), sets
37          mhold = (PKRNi+1, CRNi+1) ← DSKRNi(CRNi).
38        RNi then sends (holds, Y, VKRNi+1, v, txid, mhold,
39          hopMax, digest, te1, te2) to its neighbor towards Y
40          according to RNroutingTable for selected path to Y
41          with hopMax = hopCount of the path.
42        RNi does t = (txid, holds, v), σRNi ← Sign(skRNi, t),
43          sends (holdACKs, t, σRNi) along with local
44          routingTable to neighbor that sent holds.
45      if pay phase then
46        RNi on receiving pay tuple, sets t = (pay, txid, vkRN, v),
47        does σRN ← Sign(skRN, t).
48        RNi then creates payACK tuple as (payACK, t, σRN) to
49        neighbor it had received pay tuple from.

```

it needs to forward the message to, and sends the signed tuple to the perimeter node it knows can reach the destination RN (it finds this information from its *RNroutingTable*) (Line 15-18).

Algorithm 7: Initialization of the *pay* phase.

- 1 At the end of *hold* phase (before te_2 expiry time) if RN_r has received a $hold_s$ tuple and a $hold_r$ tuple with matching *token*, v , and *txid* values, creates a message, $m = (proceedPay, txid, v)$, creates signature $\sigma_{proceedPay} = \text{Sign}(SK_{RN_r}, m)$ sends a tuple $(proceedPay, txid, v, \sigma_{proceedPay})$ towards Bob through Bob – RN_r segment.
- 2 On receiving the message Bob and Alice communicate out of band and Bob sends $(proceedPay, txid, v, \sigma_{proceedPay})$ tuple to Alice.
- 3 Bob creates a tuple $(pay, preimage, v, txid)$ and forwards it to its neighbor $node_o$ with *txid* towards RN_r .
- 4 **for** Each node $node_i$ on *txid* path on receiving *pay* message *msg* **do**
 - 5 | $node_i$ calls *pay*(*msg*) defined in Algorithm 8.
- 6 **for** Each node $node_i$ on *txid* path on receiving *payACK* message *msg* **do**
 - 7 | $node_i$ calls *payACK*(*msg*) defined in Algorithm 8.
- 8 **for** Each node $node_i$ on *txid* path that did not receive a *payACK* tuple and current time $> te_2$ **do**
 - 9 | $node_i$ calls *payACKTimeout()* defined in Algorithm 8.

Since PCNs are highly dynamic, there might be a situation during a transaction that an RN_i on the path between RN_s and RN_r cannot find a path to the next RN_{i+1} , even after the *maxRetries* number of retries. Neither the intermediate RNs nor any other non-RN nodes on the path can deviate from the original RN path defined by the onion created by RN_s . The intermediate nodes on each segment between two RNs do not know the next segment's target RN. In this case, the transaction needs to be failed all the way back to RN_s and then retried on a different path (different intermediate RNs) from RN_s to RN_r .

4.3 Pay phase

Algorithm 7: The *pay* phase is initialized by RN_r after it receives the $hold_s$ and $hold_r$ tuples originating from Alice and Bob respectively. Specifically, RN_r decrypts C_{RN_r} contained in $hold_s$ and C'_{RN_r} contained in $hold_r$, and compares the *token* contained in both of them. If the *token* is the same, that signifies to RN_r that some nodes Alice and Bob are sender and receiver in the transaction identified by *txid*, since only the two of them know *token*. RN_r then sends a signed *proceedPay* tuple to Bob, which signals the start of the *pay* phase. Bob forwards RN_r 's *proceedPay* tuple to Alice to let her know the *pay* phase has started (Line 2). If RN_r does not receive *token* in either $hold_s$ or $hold_r$, it sends a multisig($Rev, \cdot, RN_r, \cdot, \cdot, \cdot, txid, \cdot$) to its neighbor in the transaction path.

In the *pay* phase Bob's preceding neighbor along the path pays Bob first. Following this, each node pays its successor first, then gets paid back by its predecessor. Since nodes need some form of acknowledgment that the *pay* phase has gone through successfully, RN's that initiated the current segment send signed *payACK* tuples to the nodes in their segment (Algorithm 6, Lines 20, 21).

5 SECURITY ANALYSIS

We now discuss some potential attacks on SPRITE, and mitigation strategies, and then briefly discuss the formal analysis. We also give a phase-wise analysis of malicious activities in each of the bootstrapping, hold, and pay phases of SPRITE in Appendix 8.3.

5.1 Potential Attacks and Mitigation

Transaction malleability attack: A malicious RN_s colluding with a receiver Bob and RN_r might change the transaction amount v to v' . In the Alice- RN_s segment, the amount will be v , the change occurs in the segments after that, all the way up until Bob.

Case 1: Let us assume $v' > v$ and $\delta = (v' - v)$. At the end of this attack, Alice has paid Bob v coins and RN_s has paid Bob δ coins. None of the honest intermediaries will lose money: they get paid as many coins (by their successor) as they have paid along the path to their predecessor. The only entity losing money is RN_s since it will not get paid the δ amount and will only get paid v coins, tied to the tuple it received. *Case 2:* Let $v' < v$. If Bob, RN_s , and RN_r are all malicious, Bob will get paid v' and RN_s will get paid the difference ($\delta = v - v'$) tied to the tuple received from Alice with v coins, but since they were both collaborating malicious entities, this does not affect honest intermediaries. If Bob is honest, then Bob will get paid v , but RN_s will send a lower amount v' to RN_r , thus making malicious RN_r lose money. In both cases the adversaries end up losing money but none of the honest nodes get less coins than what they paid, hence we do not consider these to be successful attacks on SPRITE.

Transaction forgery attack: We assume no honest users in the system will share their signing keys related to SPRITE with other users. This avoids any situations where an adversary can communicate on a channel created between two neighbors on behalf of one of them (e.g., Alice/Bob → Craig, where Alice and Bob share a channel and Bob is malicious), or the adversary can sign contracts on behalf of an honest Alice without Alice's knowledge (e.g., Bob → Alice → Craig, where Bob is malicious). If any user's keys are leaked then that user will generate a new set of keys and notify all her neighbors about the new keys. One could use forward-secure signatures [14] for invalidating the old leaked keys.

Sybil/Counting-based attack: An adversary could intercept network communications over time, isolating $hold_x$ messages, and associating messages sharing the same *txid* and *digest*. The adversary will try to identify the sender/receiver in a transaction by isolating messages with the highest *hopMax* or lowest *timer* values.

Counting number of hops based on *hopMax* does not reveal the identity of sender/receiver since each RN resets the *hopMax* value for each segment. The *hopMax* value is decremented by each node and is an estimate of the expected *hopCount* to the target routing helper in the current segment, and tells how far the current message should go before being dropped. This does not leak to a node in the network information about how far the sender of the current received message was from it (intermediate nodes do not know which segment they are a part of). Since te_1 and te_2 are system parameters and are included in the hold messages, all nodes in the network will receive the same value of te_1 and te_2 . On receiving the hold message, each node locally computes its timeout values $te_1.txid$ and $te_2.txid$, and does not forward the local values further.

Sender refusing to pay: Whenever there are timeouts in the hold phase for a specific segment, the sender RN for that segment will retry the hold phase on a different path. If there are timeouts in the pay phase the nodes that timed out in that specific segment, will publish their hold and pay contracts on a public repository or blockchain. Since the hold and pay contracts are signed with pseudonymous identities, this does not leak information about nodes to the public, but neighbors know each others' identities and if a node does not post a pay contract associated with a hold contract then this identifies the malicious activity to the whole network. Any honest neighbors will then avoid the malicious node for subsequent retries and transactions. If the sender is the malicious node, then all nodes on the path need to discard the hold and pay contracts and roll back the transaction since the sender has been identified as malicious and the sender-RN_s segment will not be retried.

5.2 Formal Security Analysis

We analyze the security of SPRITE in the Universal Composability framework [15]. To this end, we define an ideal functionality, $\mathcal{F}_{\text{SPRITE}}$, consisting of three functionalities, $\mathcal{F}_{\text{setup}}$, $\mathcal{F}_{\text{hold}}$, and \mathcal{F}_{pay} . We use two helper functionalities from [15], \mathcal{F}_{sig} and \mathcal{F}_{smt} , to model ideal functionalities for digital signatures and secure/authenticated channels, respectively.

$\mathcal{F}_{\text{setup}}$ models the broadcast phase where nodes register and establish payment channels and RNs register and make known their verification key to other nodes in the network. It also provides functionality for broadcasting messages such as *RN-Update* and *RN-UpdateReply*. $\mathcal{F}_{\text{hold}}$ provides interfaces for creating a *hold_s* message from sender and *hold_r* message from receiver, RN-specific *hold* phase functionalities, and the pairwise contract multisig functionality. \mathcal{F}_{pay} provides interfaces specific to the *pay* phase, creation and verification of a *pay* message, pairwise contracts creation and signing in *pay* phase, etc. We assume that all functionalities in $\mathcal{F}_{\text{sprite}}$ have access to a global clock from which they can obtain the current time. We give the proof of the following theorem along with the functionalities in the full version [40].

THEOREM 5.1. *Let $\mathcal{F}_{\text{sprite}}$ be an ideal functionality for SPRITE. Let \mathcal{A} be a probabilistic polynomial-time (PPT) adversary for SPRITE, and let \mathcal{S} be an ideal-world PPT simulator for $\mathcal{F}_{\text{sprite}}$. SPRITE UC-realizes $\mathcal{F}_{\text{sprite}}$ for any PPT distinguishing environment \mathcal{Z} .*

6 EXPERIMENTAL ANALYSIS

6.1 Experimental Setup

We compared R2RB and R2NB with BIAnC [39] and Speedy Murmurs [48] (referred to as SM in this section), across two topology types, ten topologies each [50]. The first topology, referred to as LT, was taken from the publicly available Lightning gossip dataset [2] from May 31, 2022. The network has 15833 nodes and 156072 channels. We removed any nodes that did not have any outgoing connections along with 80% of the nodes which had one incoming or outgoing connection (these nodes are not involved in routing), leaving 8995 nodes and 129724 channels in LT. We designated the top 10 highly connected nodes as RNs for evaluating R2RB and BIAnC, and to act as landmarks in SM. As channel capacity is not present in the gossip messages from the Lightning data, we choose

the maximum allowed amount for a single transaction as the link weight as this value should correlate to a realistic channel capacity. We compare the performance of BIAnC, SM, and R2RB on LT. R2NB is not applicable to LT due to the closely located RN nodes.

We constructed a second privacy-preserving network topology (PPNT) as described in Section 4, to evaluate R2RB, R2NB, BIAnC, and SM. We start by taking the RNs in LT and start adding nodes to the PCN where the initial few nodes set up payment channels with RNs but subsequent nodes joining the network connect to other regular nodes, thus forming layers around the RNs. We add nodes until each RN has a diameter of about 7 hops and a neighborhood of roughly 800 nodes. The perimeter nodes of each neighborhood are randomly connected to perimeter nodes belonging to other RN neighborhoods. PPNT had 7978 nodes and 25302 channels. The link weights used in this topology are similar to LT. We categorized the link weights from LT into two groups, the first group contained channels with at least one highly connected node (RN), and the second group was made up of links between two regular nodes. The link weights were then randomly sampled from these two groups and assigned to the links in PPNT based on the channel type.

We randomly chose senders and receivers with at least 3 and 8 hops between them for LT & PPNT respectively. Although publicly available data for the Lightning network claims an average of $22k$ transactions per day [1] – significantly lower than 10 transactions/second, we set a transaction rate at 10 transaction/second. This high rate was used to assess the scalability of SPRITE. In SM, each transaction gets split into 10 uniform sub-transactions, one for each RN (referred to as landmarks in SM).

We implemented R2RB, R2NB, SM and BIAnC, and deployed the generated topologies in the *ns-3* simulator [3] for our experiments. The results were averaged over 10 runs with PPNT for a total of $100k$ transactions. The simulations were run on a Desktop class machine with Intel(R) Core(TM) i7-10700 @ 3.8 GHz CPU and 64 GB of RAM. The metrics for comparison are: path stretch (ratio of the hop-count of a completed transaction to the optimal hop-count), end-to-end transaction processing time (latency), transaction success rate, set up costs during Bootstrap phase (message complexity and duration), and the overall message complexity of the entire simulation.

6.2 Experimental Results

LT Topology results: Figure 3a shows the growth of the message complexity within LT over time. BIAnC inundates the network with broadcasts for each transaction and given the interconnected nature of LT this results in a dramatic increase in message complexity, growing at a rate roughly 100 times that of SM. SM, while having only a fraction of the number of messages compared to BIAnC, still grows at a much faster rate than R2RB. This is attributable to the splitting of each transaction and the acknowledgments sent back on the payment path in the routing phase.

Figure 3b shows the growth of latency with respect to hop-count. BIAnC has a higher latency compared to R2RB and SM. This is attributed to BIAnC having three phases as opposed to two in SM and R2RB. Note that given its sub-optimality, BIAnC never chooses a 3-hop sender-receiver path. SM and R2RB have similar latencies. We model cryptographic operations for both R2RB and BIAnC, but not for SM (they didn't have any). We also do not model the delay




Figure 3: Results for simulations in the Lightning Topology (LT).

imposed by blockchain operations for BIAnC. The hop-counts of transactions in LT range between 3-10 hops, with BIAnC, SM, and R2RB having average hop-counts of 7, 5, and 6, respectively. The hop-counts for SM represent the highest across the hop-counts of all the split transactions.

Figure 3c shows the cumulative distribution function (CDF) of latencies for all transactions. Both R2RB and SM outperform BIAnC significantly, which had an average latency of 113.5 ms, while SM and R2RB had average latencies of 79.6 ms and 95.3 ms, respectively. The additional delay in BIAnC is on account of the extra broadcast-based *Find* phase. R2RB is able to perform almost as well as SM in terms of real-world delays while providing significantly more security and privacy guarantees. It also has a significantly higher transaction success rate at 97.17% compared to SM's 81.3%. R2RB outperforms SM in terms of success rate due to our in-network retry mechanism, as well as *routingTable* updates that are propagated within the network for each *holdACK* and *holdReject* message. In LT 9.864% of transactions required a retry attempt for R2RB. Due to the design of BIAnC, the sender can only send the maximum available credits on the fastest path to the receiver, hence, only 69.06% of transactions sent the full amount of required credits. For practical applications, these transactions can be repeated by splitting the larger ones into sub-transactions, similar to SM.

Figure 3d shows the total number of messages required to bootstrap the network with routing information while Figure 3e shows the duration of the phase. BIAnC is excluded from this comparison as it does not have a Bootstrap phase. SM requires more messages

for its bootstrapping phase in LT than R2RB but takes about 10 ms less than R2RB to complete this phase.

The path stretch of transactions is shown in Figure 3f, it should be noted that BIAnC always finds the most optimal path in terms of hop-count due to its broadcast-based pathfinding mechanism. The path stretch for SM was calculated by taking the average amount of hops taken by each sub-transaction and comparing that against the optimal path (obtained from Dijkstra's algorithm) between the sender and the receiver. For R2RB and BIAnC, the number of hops taken by a transaction were compared against the total hops in the corresponding optimal paths between the sender & RN_s , RN_s & RN_r , and RN_r & the receiver. SM incurs the worst path stretch with a median of 1.075, while R2RB has a median path stretch of 1.0. The variation in path stretch for transactions in R2RB is due to the *routingTable* of nodes becoming stale as the simulation progresses with new transactions. The routing tables can remain fresh by issuing periodic broadcasts from RNs, similar to the Bootstrap phase, to update the *routingTable* of nodes. The higher path stretch in SM can be attributed to its embedded prefix routing and splitting of transactions among different paths.

PPNT Topology results: All four schemes show linear growth in the number of messages as seen in Figure 4a. With R2RB and R2NB the number of messages is the lowest and continues to grow linearly at these low values. As with LT, BIAnC's *Find* phase results in thousand times more messages than R2RB and R2NB while SM results in a ten times higher number of messages in comparison.

Figure 4b shows the growth of latency with respect to hop-count. Both R2RB and R2NB have a slightly larger latency for each




Figure 4: Results for simulations in the Privacy Preserving Network Topology (PPNT).

transaction of a given hop-count when compared to SM due to the cryptographic operations between pairs of nodes on the path.

Figure 4c shows a CDF where it can be observed that for the majority of transactions, R2RB and R2NB have lower latencies than BIAnC and SM while maintaining transaction success rates of 97.01% and 96.02%. SM on the other hand, has a success rate of 76.08%. Roughly five percent of transactions in SPRITE (R2RB and R2NB) have higher latencies than those found in BIAnC and SM due to SPRITE's in-network retries that would otherwise fail.

The number of messages and the duration of the Bootstrap phase were averaged over ten runs; results shown in Figure 4d and 4e. In contrast to LT, SM has high number of messages, with an average of around 348k messages when compared to R2RB with 190k messages and R2NB with 88k messages respectively. The high complexity of R2RB and SM is due to the more distributed nature of the PPNT network, where the landmarks (RNs in SPRITE) have a much lower degree than LT. In SM, this results in more nodes receiving multiple messages for each landmark advertisement compared to LT. Transactions in SM take the least amount of time, while R2RB takes the most, but similar to LT the difference is negligible.

The transaction path stretch in Figure 4f shows that BIAnC is the most efficient in terms of transaction path length. This is because it finds the most optimal path in terms of hop-count due to its broadcast-based pathfinding mechanism. This optimal path stretch does come at the cost of higher overhead and much higher latencies.

The median path stretch value for SM is 1.45 and is significantly higher than R2RB and R2NB with respective values of 1.07 and 1.15. Due to the distributed topology, the prefix-based embedding

system in SM does not identify the shortest path when landmarks are far from the sender or receiver. R2NB's inefficiency is due to the unknown distance of chosen perimeter node to the next RN.

7 CONCLUSION

In this paper, we present SPRITE, a secure, privacy-preserving, and efficient routing protocol for payment channel networks. SPRITE can support concurrent transactions and takes just two rounds of communication for pathfinding and routing transactions, which is the most optimal till date. One direction for future work is investigating the design of economic models for estimating and optimizing routing fees for both, regular nodes and routing nodes in a PCN. Another direction for future work is to mechanically verify the proof of security of SPRITE (and potentially other PCN protocols) using interactive theorem provers such as EasyUC [16].

ACKNOWLEDGEMENTS

The authors thank Ryan Gentry and Alex Bosworth from Lightning Labs for their insights and discussion, which helped us improve the paper. The authors also thank Kartick Kolachala for his help with Table 1, and the anonymous reviewers for their feedback. This research was partially funded by the US National Science Foundation under grants #2148358 and #1914635, and the US Department of Energy grant #DE-SC0023392. Any opinions, findings, conclusions, or recommendations expressed in this material are solely those of the authors and do not necessarily reflect the views of the US federal agencies.

REFERENCES

- [1] [n. d.]. The growth of the Lightning Network. <https://k33.com/research/archive/articles/the-growth-of-the-lightning-network>.
- [2] [n. d.]. Lightning network gossip datasets and topology. <https://github.com/lmresearch/topology>.
- [3] [n. d.]. ns-3 network simulator. <https://www.nsnam.org/>.
- [4] [n. d.]. Ripple data. <https://data.ripple.com/>.
- [5] [n. d.]. Visa fact sheet. <https://www.visa.co.uk/dam/VCOM/download/corporate/media/visanet-technology/aboutvisafactsheet.pdf>.
- [6] [n. d.]. What is the Lightning Network in Bitcoin and how does it work? <https://cointelegraph.com/bitcoin-for-beginners/what-is-the-lightning-network-in-bitcoin-and-how-does-it-work>.
- [7] [n. d.]. XrpScan. <https://xrpScan.com/>.
- [8] 2023. *Phoenix Wallet 4: Trampoline payments*. Accessed: 2023-12-19.
- [9] 2023. *Trampoline routing*. Accessed: 2023-12-19.
- [10] 2023. *What are trampoline payments*. Accessed: 2023-12-19.
- [11] Abdelrahaman Aly, Edouard Cuvelier, Sophie Mawet, Olivier Pereira, and Mathieu Van Vyve. 2013. Securely Solving Simple Combinatorial Graph Problems. In *Financial Cryptography and Data Security - 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 7859)*, Ahmad-Reza Sadeghi (Ed.). Springer, 239–257.
- [12] Abdelrahaman Aly and Mathieu Van Vyve. 2014. Securely Solving Classical Network Flow Problems. In *Information Security and Cryptology - ICISC 2014 - 17th International Conference, Seoul, Korea, December 3-5, 2014, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 8949)*, Jooyoung Lee and Jongsung Kim (Eds.). Springer, 205–221.
- [13] Andreas Anotonopoulos, Olaoluwa Osuntokun, and Rene Pickhardt. [n. d.]. Mastering the Lightning Network. <https://github.com/lmbook/lmbook>.
- [14] Mihir Bellare and Sara K. Miner. 1999. A Forward-Secure Digital Signature Scheme. In *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings (Lecture Notes in Computer Science, Vol. 1666)*, Michael J. Wiener (Ed.). Springer, 431–448.
- [15] Ran Canetti. 2004. Universally composable signature, certification, and authentication. In *Proceedings. 17th IEEE Computer Security Foundations Workshop, 2004*.
- [16] Ran Canetti, Alley Stoughton, and Mayank Varia. 2019. EasyUC: Using EasyCrypt to Mechanize Proofs of Universally Composable Security. In *32nd IEEE Computer Security Foundations Symposium, CSF 2019, Hoboken, NJ, USA, June 25-28, 2019*. IEEE, 167–183.
- [17] Yanjiao Chen, Yuyang Ran, Jingyue Zhou, Jian Zhang, and Xueluan Gong. 2022. MPCN-RP: A Routing Protocol for Blockchain-Based Multi-Charge Payment Channel Networks. *IEEE Transactions on Network and Service Management* 19, 2 (2022), 1229–1242.
- [18] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed E. Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. 2016. On Scaling Decentralized Blockchains - (A Position Paper). In *Financial Cryptography and Data Security - FC 2016 International Workshops, BITCOIN, VOTING, and WAHC, Christ Church, Barbados, February 26, 2016, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 9604)*, Jeremy Clark, Sarah Meiklejohn, Peter Y. A. Ryan, Dan S. Wallach, Michael Brenner, and Kurt Rohloff (Eds.). Springer, 106–125.
- [19] Christian Decker. [n. d.]. Lightning Network Research - Topology Datasets. <https://github.com/lmresearch/topology>. <https://doi.org/10.5281/zenodo.4088530>
- [20] Lisa Eckey, Sebastian Faust, Kristina Hostákova, and Stefanie Roos. 2020. Splitting Payments Locally While Routing Interdimensionally. *IACR Cryptol. ePrint Arch.* 2020 (2020), 555.
- [21] Felix Engelmann, Henning Kopp, Frank Kargl, Florian Glaser, and Christof Weinhardt. 2017. Towards an Economic Analysis of Routing in Payment Channel Networks. In *Proceedings of the 1st Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers (SERIAL '17)*. Article 2, 6 pages.
- [22] Ittay Eyal, Adem Efe Gencer, Emin Gün Sirer, and Robbert van Renesse. 2016. Bitcoin-NG: A Scalable Blockchain Protocol. In *13th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2016, Santa Clara, CA, USA, March 16-18, 2016*, Katerina J. Argyraki and Rebecca Isaacs (Eds.). USENIX Association, 45–59.
- [23] L.R. Ford and D.R. Fulkerson. 1954. Maximal flow through a network. *Canadian Journal of Mathematics* 8 (1954).
- [24] A.V. Goldberg and R. E. Tarjan. 1988. A new approach to the maximum flow problem. *J. of ACM* 35 (1988), 921–940.
- [25] Qianyun Gong, Chengjin Zhou, Le Qi, Jianbin Li, Jianzhong Zhang, and Jingdong Xu. 2021. VEIN: High Scalability Routing Algorithm for Blockchain-based Payment Channel Networks. In *20th IEEE International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom 2021, Shenyang, China, October 20-22, 2021*. IEEE, 43–50.
- [26] Hsiang-Jen Hong, Sang-Yoon Chang, and Xiaobo Zhou. 2022. Auto-Tune: Efficient Autonomous Routing for Payment Channel Networks. In *2022 IEEE 47th Conference on Local Computer Networks (LCN)*. 347–350.
- [27] Heba Kadry and Yasser Gadallah. 2021. A Machine Learning-Based Routing Technique for Off-chain Transactions in Payment Channel Networks. In *2021 IEEE International Conference on Smart Internet of Things (SmartIoT)*. 66–73. <https://doi.org/10.1109/SmartIoT52359.2021.00020>
- [28] Angelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10401)*, Jonathan Katz and Hovav Shacham (Eds.). Springer, 357–388.
- [29] Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Ewa Syta, and Bryan Ford. 2018. OmniLedger: A Secure, Scale-Out, Decentralized Ledger via Sharding. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*. IEEE Computer Society, 583–598.
- [30] MIT Media lab. Digital currency initiative. [n. d.]. Layer 2: The Lightning Network. <https://dci.mit.edu/lightning-network>.
- [31] lightning [n. d.]. Lightning network. <https://lightning.network/>.
- [32] Lightning Network [n. d.]. Lightning Network Routing Nodes. <https://docs.lightning.engineering/the-lightning-network/multihop-payments/what-makes-a-good-routing-node>.
- [33] Siyi Lin, Jingjing Zhang, and Weigang Wu. 2020. FSTR: Funds Skewness Aware Transaction Routing for Payment Channel Networks. In *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2020, Valencia, Spain, June 29 - July 2, 2020*. IEEE, 464–475.
- [34] G. Malavolta, P. Moreno-Sánchez, A. Kate, and M. Maffei. 2017. SilentWhispers: Enforcing Security and Privacy in Decentralized Credit Networks. In *Annual Network and Distributed System Security Symposium, NDSS*.
- [35] G. Malavolta, P. Moreno-Sánchez, A. Kate, M. Maffei, and S. Ravi. 2017. Concurrency and Privacy with Payment-Channel Networks. In *Proceedings ACM SIGSAC Conference on Computer and Communications Security, CCS*. 455–471.
- [36] Andrew Miller, Iddo Bentov, Surya Bakshi, Ranjit Kumaresan, and Patrick McCorry. 2019. Sprites and State Channels: Payment Networks that Go Faster Than Lightning. In *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*. 508–526.
- [37] Andrew Miller, Ari Juels, Elaine Shi, Bryan Parno, and Jonathan Katz. 2014. Permacoin: Repurposing Bitcoin Work for Data Preservation. In *2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*. IEEE Computer Society, 475–490.
- [38] P. Moreno-Sánchez, A. Kate, M. Maffei, and K. Pecina. 2015. Privacy Preserving Payments in Credit Networks: Enabling trust with privacy in online marketplaces. In *Annual Network and Distributed System Security Symposium, NDSS*.
- [39] Gaurav Panwar, Satyajayant Misra, and Roopa Vishwanathan. 2019. BLAnC: Blockchain-based Anonymous and Decentralized Credit Networks. In *Proceedings of the Ninth ACM Conference on Data and Application Security and Privacy, CODASPY*. 339–350.
- [40] Gaurav Panwar, Roopa Vishwanathan, George Torres, and Satyajayant Misra. [n. d.]. SPRITE: Secure and Private Routing in Payment Channel Networks (Full Version). <https://eprint.iacr.org/2024/122>.
- [41] Sunoo Park, Albert Kwon, Georg Fuchsbauer, Peter Gazi, Joël Alwen, and Krzysztof Pietrzak. 2018. SpaceMint: A Cryptocurrency Based on Proofs of Space. In *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 10957)*, Sarah Meiklejohn and Kazue Sako (Eds.). Springer, 480–499.
- [42] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. 2019. SpOT-Light: Lightweight Private Set Intersection from Sparse OT Extension. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 11694)*, Alexandra Boldyreva and Daniele Micciancio (Eds.). Springer, 401–431.
- [43] David Pointcheval and Olivier Sanders. 2016. Short Randomizable Signatures. In *Topics in Cryptology - CT-RSA 2016 - The Cryptographers' Track at the RSA Conference 2016, San Francisco, CA, USA, February 29 - March 4, 2016, Proceedings (Lecture Notes in Computer Science, Vol. 9610)*, Kazue Sako (Ed.). Springer, 111–126.
- [44] Joseph Poon and Thaddeus Dryja. [n. d.]. The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments. <https://lightning.network/lightning-network-paper.pdf>.
- [45] raiden [n. d.]. Raiden network. <https://raiden.network/>.
- [46] ripple [n. d.]. Ripple. <https://ripple.com>.
- [47] S. Roos, P. Moreno-Sánchez, A. Kate, and I. Goldberg. 2018. Settling payments fast and private: efficient decentralized routing for path-based transactions. In *Annual Network and Distributed System Security Symposium, NDSS*.
- [48] Stefanie Roos, Pedro Moreno-Sánchez, Aniket Kate, and Ian Goldberg. 2018. Settling Payments Fast and Private: Efficient Decentralized Routing for Path-Based Transactions. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*. The

Internet Society.

[49] Vibhaalakshmi Sivaraman, Shailesh Bojja Venkatakrishnan, Kathleen Ruan, Parimjan Negi, Lei Yang, Radhika Mittal, Giulia C. Fanti, and Mohammad Alizadeh. 2020. High Throughput Cryptocurrency Routing in Payment Channel Networks. In *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020, Santa Clara, CA, USA, February 25-27, 2020*, Ranjita Bhagwan and George Porter (Eds.). USENIX Association, 777–796.

[50] George Torres and Gaurav Panwar. [n. d.]. SPRITE Implementation. <https://github.com/nsol-nmsu/sprite>.

[51] B. Viswanath, M. Mondal, P. K. Gummadi, A. Mislove, and A. Post. 2012. Canal: scaling social network-based Sybil tolerance schemes. In *Proceedings of EuroSys*. 309–322.

[52] Peng Wang, Hong Xu, Xin Jin, and Tao Wang. 2019. Flash: Efficient Dynamic Routing for Offchain Networks (*CoNEXT '19*). Association for Computing Machinery, New York, NY, USA, 370–381. <https://doi.org/10.1145/3359989.3365411>

[53] Ruozhou Yu, Guoliang Xue, Vishnu Teja Kilari, Dejun Yang, and Jian Tang. 2018. CoinExpress: A Fast Payment Routing Mechanism in Blockchain-Based Payment Channel Networks. In *27th International Conference on Computer Communication and Networks, ICCCN 2018, Hangzhou, China, July 30 - August 2, 2018*. IEEE, 1–9.

[54] Xiaoxue Zhang, Shouqian Shi, and Chen Qian. 2021. WebFlow: Scalable and Decentralized Routing for Payment Channel Networks with High Resource Utilization. *CoRR* abs/2109.11665 (2021). <https://arxiv.org/abs/2109.11665>

[55] Yuhui Zhang and Dejun Yang. 2019. RobustPay: Robust Payment Routing Protocol in Blockchain-based Payment Channel Networks. In *2019 IEEE 27th International Conference on Network Protocols (ICNP)*. 1–4.

[56] Yuhui Zhang and Dejun Yang. 2021. RobustPay⁺: Robust Payment Routing With Approximation Guarantee in Blockchain-Based Payment Channel Networks. *IEEE/ACM Trans. Netw.* 29, 4 (2021), 1676–1686.

[57] Yuhui Zhang, Dejun Yang, and Guoliang Xue. 2019. CheaPay: An Optimal Algorithm for Fee Minimization in Blockchain-Based Payment Channel Networks. In *2019 IEEE International Conference on Communications, ICC 2019, Shanghai, China, May 20-24, 2019*. IEEE, 1–6.

8 APPENDIX

8.1 AS Function Definitions

DEFINITION 1. (*Sequential Aggregate Signatures [43]*). Let $\mathbb{G}_1, \mathbb{G}_2$ be prime-order cyclic groups of size p , such that $g \in \mathbb{G}_1, \tilde{g} \in \mathbb{G}_2$, and $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$.

- **AS.Setup(1^k):** Given a security parameter k , this algorithm selects a random $x \in \mathbb{Z}_p$ and outputs $pp \leftarrow (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g, X, \tilde{g}, \tilde{X})$, where $X = g^x$ and $\tilde{X} = \tilde{g}^x$.
- **AS.KeyGen(pp):** This algorithm selects a random $y \leftarrow \mathbb{Z}_p$, computes $\tilde{Y} \leftarrow \tilde{g}^y$ and sets sk as y and pk as Y .
- **AS.Sign($sk, \sigma, (m_1, \dots, m_r), (pk_1, \dots, pk_r), m$)** proceeds as follows:
 - If $r = 0$, then $\sigma \leftarrow (g, X)$;
 - If $r > 0$ but $AS.Verify((pk_1, \dots, pk_r), \sigma, (m_1, \dots, m_r)) = 0$, then it halts;
 - If $m = 0$, then it halts;
 - If for some $j \in \{1, \dots, r\}$ $pk_j = pk$, then it halts.
If the algorithm did not halt, then it parses sk as y and σ as (σ_1, σ_2) , selects $t \leftarrow \mathbb{Z}_p$ and computes $\sigma' = (\sigma'_1, \sigma'_2) \leftarrow (\sigma_1^t, (\sigma_2 \cdot \sigma_1^{y \cdot m})^t)$. It eventually outputs σ' .
- **AS.Verify($(pk_1, \dots, pk_r), (m_1, \dots, m_r), \sigma$)** parses σ as (σ_1, σ_2) and pk_j as \tilde{Y}_j , for $j = 1, \dots, r$, and checks whether $\sigma_1 \neq 1_{\mathbb{G}_1}$ and $e(\sigma_1, \tilde{X} \cdot \prod \tilde{Y}_j^{m_j}) = e(\sigma_2, \tilde{g})$ are both satisfied. In the positive case, it outputs 1, and 0 otherwise.

8.2 Algorithms

Subroutine for intermediate node (Algorithm 8): This algorithm details the functions called by nodes in the network during *hold* and *pay* phases discussed in Algorithm 5. *hold* function is called by different nodes when they receive a $hold_s$ or $hold_r$ tuple.

If the receiving node happens to be a perimeter node, it checks if the tuple was addressed to it (line 5). If so, and if it has a path to transmit the asking amount, it updates the $hold_s$ tuple with the destination RN as the receiver and forwards the message. Nodes in the network that have a path to the target RN or perimeter node, set the transaction's local retry counter to be zero. They also store the transaction digest and identity of the destination RN, initialize the timers te_1 and te_2 , forward the message to their neighbor, and sign a multisig *hold* contract with their neighbor. Each node also sets its transaction-specific timers, $txid.te1$ and $txid.te2$. If the node is the target RN node in the *hold* message, it follows Algorithm 6 to process the message.

If the node is a non-perimeter, non-RN node, it signs a multisig contract with its neighbor and forwards the message to its next-hop neighbor after checking if it has a viable path with sufficient liquidity to the destination RN/perimeter node (lines 11-16). Let us now focus on nodes setting their timers. Each Node computes two transaction-specific timers $txid.te1$ and $txid.te2$ as functions of the current time *currTime* and the global te_1 and te_2 (line 13, 15). We do this since each node that is not a destination (either perimeter node or an RN) cannot simply use te_1 and te_2 for timeouts, since we need to ensure that timeouts of nodes are staggered, i.e., the node closest to a destination times out first, whereas the nodes farthest away from the destination time out last. If a node receives multiple copies of a $hold_x$ message (where $x \in \{s, r\}$ corresponding to sender and receiver respectively) associated with the same $txid$, it ignores subsequent messages. When a node cannot find any viable path to a destination (even after retries), it sends a *holdReject* message back to its predecessor from whom it received the $hold_x$ message. It also sends a copy of its current *routingTable* so the predecessor can update its own *routingTable*.

The *holdReject* function is called by a node in the network on receiving a *holdReject* message. Nodes in the network could choose to retry on other paths upto *maxRetries* limit (lines 18-19) or choose to forward *holdReject* message back towards the sender if path is unavailable to a target RN (line 23). If Alice or Bob receive a *holdReject*, that means that there is no viable path from their neighbor to RN_s (for Alice) or RN_r (for Bob), hence they need to retry the transaction with different neighbors and possibly new RNs (line 21). The way RNs handle $hold_x$, *holdReject* messages is slightly different; we discuss RNs' actions in these phases separately in Algorithm 6. The *holdACK* function is called by a node on receiving a $holdACK_x$ message. When the target RN in a given segment receives a $hold_x$ tuple, and is willing to process the transaction, it sends back a *holdACK* tuple containing its signature to the neighboring node it received the $hold_x$ tuple from.

When a node receives a $holdACK_x$ tuple, it deletes its timer te_1 , and forwards the $holdACK_x$ to its predecessor. Ultimately, the sender should receive the $holdACK_x$ tuple, which will tell her that the $hold_x$ messages along that segment were successful, and reached the destination RN. If a node does not receive a *holdACK* tuple until te_1 expires, it calls the *holdACKTimeout* function, which retries the transaction if possible, else it drops the transaction and sends a $holdReject_x$ tuple towards the sender (line 28). If any node including an RN receives a *hold* tuple on a different path, after it has already forwarded a *holdACK* towards Alice on another path, then it should replay the *holdACK* tuple on the new path and send a *holdReject*

message on the previous path. This accounts for a malicious node on any path not forwarding *holdACK* tuples downstream thus timing out nodes preceding it, and prompting the formation of another path to the target RN. Due to space constraints, we discuss the pay phase actions in the full version [40].

Algorithm 8: Subroutine for every node for *hold* and *pay* phase

```

def hold ( $hold_x$ ,  $Y$ ,  $VK_{RN(\cdot)}$ ,  $v$ ,  $txid$ ,  $C_{RN(\cdot)}$ ,  $hopMax$ ,  $digest$ ,  

 $te_1, te_2$ )
  if  $\nexists (Y, node_j, p_j)$  in  $routingTable$  with  

 $p_j = (hopCount, currMax_s, currMax_r, t_e)$  where  

 $(hopCount \leq hopMax) \wedge ((currMax_x \geq v) \vee ((lw_{ij} \geq v) \wedge (currTime \geq t_e)))$  then
    Create tuple ( $holdReject_x$ ,  $Y$ ,  $VK_{RN(\cdot)}$ ,  $v$ ,  $txid$ ) and  

    send with  $routingTable$  to neighbor that sent  $hold_x$ .  

    Call multisig( $Rev, \perp, i, j, lw_{ij}, v, txid, ts$ ) and delete  

 $retry.txid$ ,  $digest.txid$ ,  $segTarget.txid = RN(\cdot)$ ,  

 $te_1.txid$ , and  $te_2.txid$ .  

    return
  if  $Nonce_i \in (hold_s, Nonce_i, VK_{RN(\cdot)}, v, txid, C_{RN(\cdot)},$   

 $hopMax, digest, te_1, te_2)$  tuple belongs to  $node_i$  then
    Lookup  $routingTable$  for tuple ( $RN(\cdot)$ ,  $node_j, p_j$ )  

    with  $p_j = (hopCount, currMax_s, currMax_r, t_e)$   

    where  $(hopCount \leq hopMax) \wedge ((currMax_x \geq v) \vee ((lw_{ij} \geq v) \wedge (currTime \geq t_e)))$ , update  $hold_s$   

    tuple to ( $hold_s, RN(\cdot), VK_{RN(\cdot)}, v, txid, C_{RN(\cdot)}$ ,  

 $hopMax, digest, te_1, te_2$ ) with  $hopMax = hopCount$ ,  

    and forward to  $node_j$ . Set  $retry.txid = 0$ ,  

 $digest.txid = digest$ ,  $segTarget.txid = RN(\cdot)$ ,  

 $te_1.txid = currTime + (te_1 * hopCount)$ , and  

 $te_2.txid = currTime + te_2$ . Call multisig( $\perp, hold_s, i, j, lw_{ij}, v, txid, ts$ ).  

    Update  $t_e = currTime + e$  and  $currMax_s = currMax_s - v$  for  $p_j$  in  $routingTable$ . return
  else if  $node_i == RN(\cdot)$  then
    Follow Alg. 6.
  else
    Update  $hold_x$  tuple  $hopMax = hopMax - 1$  and  

    forward tuple to  $node_j$ .
    if  $Y = Nonce(\cdot)$  then
      Set  $retry.txid = 0$ ,  $digest.txid = digest$ ,  

 $segTarget.txid = RN(\cdot)$ ,  $te_1.txid = currTime + (te_1 * (hopCount + hopMax_{RN}))$ ,  

      and  $te_2.txid = currTime + te_2$ .
    else
      Set  $retry.txid = 0$ ,  $digest.txid = digest$ ,  

 $segTarget.txid = RN(\cdot)$ ,  

 $te_1.txid = currTime + (te_1 * hopCount)$ , and  

 $te_2.txid = currTime + te_2$ .
    Call multisig( $\perp, hold_x, i, j, lw_{ij}, v, txid, ts$ ). Update  

 $t_e = currTime + e$  and  $currMax_x = currMax_x \pm v$  for  $p_j$  in  $routingTable$ .
  
```

R2RB Bootstrap protocol (Algorithm 9) describes the operations during the bootstrap phase of the R2RB protocol.

```

def holdReject ( $holdReject_x$ ,  $Y$ ,  $VK_{RN(\cdot)}$ ,  $v$ ,  $txid$ )
  Update local  $routingTable$  with new info received.
  if  $(\exists (Y, node_j, p_j) \text{ with } p_j = (hopCount, currMax_s,$   

 $currMax_r, t_e) \text{ where } (hopCount \leq hopMax) \wedge$   

 $(retry.txid < maxRetries) \wedge ((currMax_x \geq v) \vee$   

 $((lw_{ij} \geq v) \wedge (currTime \geq t_e)))$  then
    Update  $hold_x$  tuple  $hopMax = hopMax - 1$  and  

    forward tuple to  $node_j$ . Call multisig( $\perp, hold_x, i, j, lw_{ij}, v, txid, ts$ ). Set  $retry.txid = retry.txid + 1$ .
  else if  $(node_i == Alice \wedge \nexists (RN_s, node_j, p_j)) \vee$   

 $(node_i == Bob \wedge \nexists (RN_r, node_j, p_j))$  where  

 $p_j = (hopCount, currMax_s, currMax_r, t_e)$  and  

 $currMax_x \geq v$  then
    Choose new  $v'$  and restart Algorithm 3 and 4.
  else
    Forward tuple ( $holdReject_x$ ,  $Y$ ,  $VK_{RN(\cdot)}$ ,  $v$ ,  $txid$ ) along  

    with local  $routingTable$  to neighbor that sent  $hold_x$ .  

    Call multisig( $Rev, \perp, i, j, lw_{ij}, v, txid, ts$ ) and delete  

 $retry.txid$ ,  $digest.txid$ ,  $te_1.txid$ , and  $te_2.txid$ .
def holdACK ( $holdACK_x$ ,  $t$ ,  $\sigma_{RN(\cdot)}$ )
  Update local  $routingTable$  with new info received.
  Parse  $t = (txid, hold_x, v)$ . Verify( $VK_{RN(\cdot)}, \sigma_{RN(\cdot)}, t \rightarrow 1$ ,  

 $RN(\cdot) == segTarget.txid$ , and delete timer  $te_1$  for  $txid$ .
   $node_i$  then forwards the  $holdACK_x$  tuple with  

 $routingTable$  to neighbor that sent  $hold_x$ .
def holdACKTimeout()
   $node_i$  calls multisig( $Rev, \perp, i, j, lw_{ij}, v, txid, ts$ ) to  $node_j$   

  that it had sent  $hold_x$  tuple to, and retries send  $hold_x$  to  

  other neighbors for target  $Y$  for  $txid$ . If no such  

  neighbors exist, create  $holdReject_x$  tuple, call  

  multisig( $Rev, \perp, i, o, lw_{io}, v, txid, ts$ ), and send along  

  with  $routingTable$  to  $node_o$  that sent  $hold_x$  message.  

  Delete  $retry.txid$ ,  $digest.txid$ ,  $segTarget.txid = RN(\cdot)$ ,  

 $te_1.txid$ , and  $te_2.txid$ .
def pay( $pay, preimage, v, txid$ )
  if  $H(preimage) \neq digest.txid$  then
    return  $\perp$ .
  if  $node_i$  is an RN then
    Follow Alg. 6.
  if  $node_i$  is Alice then
    Create  $t = (pay, txid, preimage_{txid}, v)$ , set  $payACK =$   

 $(payACK, t, \perp)$  and send to neighbor that sent  $pay$   

    tuple. return
  Forward  $pay$  tuple to next neighbor  $node_o$  on  $txid$  path  

  along with multisig( $\perp, pay, i, o, lw_{io}, v, txid, ts$ ).

```

We recall that R2RB differs from R2NB in the distance each RN has to broadcast the *RN-Update* message which is depicted in Figure 5. This distance is larger in R2RB because of the absence of perimeter nodes in the network. The broadcasted messages from each RN travel a certain number of hops away from the RN, allowing nodes in the given area to route transactions to the corresponding RN. Due to the larger broadcast area, neighboring RNs will receive each other's broadcast messages and be able to route transactions

```

def payACK(payACK, t,  $\sigma_{RN_{(.)}}$ ) then
36   if Received (payACK, t,  $\sigma_{RN_{(.)}}$ ) then
37     Parse t = (pay, txid,  $vk_{RN_{(.)}}$ , v), Verify( $vk_{RN_{(.)}}$ ,
38      $\sigma_{RN_{(.)}}$ , t)  $\rightarrow$  1, verify  $RN_{(.)} \in RN$ , delete  $te_2.txid$ .
39     nodei then forwards the payACK tuple to the
40     neighbor it had received the pay tuple from.
41   else if Received (payACK, t,  $\perp$ ) then
42     Parse t = (pay, txid,  $preimage_{txid}$ , v), verify
43      $H(preimage_{txid}) \stackrel{?}{=} digest.txid$ , if true delete  $te_2.txid$ .
44     If nodei ==  $RN_s$ , return, else forward the payACK to
45     neighbor that sent pay tuple.
def payACKTimeout()
46   nodei calls multisig(Rev,  $\perp$ , i, j,  $lw_{ij}$ , v, txid, ts) to the
47   neighbor nodej that it had originally sent pay tuple to
48   and to the other neighbor that it had originally
49   received pay tuple from.

```




Figure 5: RNs neighborhoods in R2RB.

between them directly. The key advantage for R2RB is the elimination of perimeter nodes, with the trade-off of larger message complexity in the system due to larger broadcast distances for the *RN-Update* message.

Multisig contracts (Algorithm 10): The hold and pay phases involve neighboring nodes signing multisig contracts between them. In the hold phase, the contract stipulates that two neighboring nodes *j* and *k* agree to decrease/increase their link weights lw_{jk} and lw_{kj} respectively, by the sender's asking amount (*v*) in the future when the pay tuple comes through. The multisig contract in the pay phase actually updates the link weights, and both neighboring nodes need to sign the new balances. Note that in the RN_r -Bob segment, since the payment goes in the $RN_r \rightarrow$ Bob direction, the link weights are updated in the opposite direction compared to the Alice - RN_s and RN_s - RN_r segments. If a multisig contract signed in the hold or pay phases needs to be revoked, the contract and signatures on them are discarded.

8.3 Informal Security Analysis

Bootstrapping phase: For verifying if RNs set up correct AS parameters, *pp*, all nodes along a path can individually check if they can produce a valid signature on a test message, else discard the

Algorithm 9: R2RB: Bootstrap broadcast from RN to RN

```

1 Each node i initializes a table, routingTablei containing
2   columns: (reachable RNs, next hop neighbor j, currMaxs,
3   currMaxr, hopCount, te).
2 for each RN, k  $\in RN$  do
3   k does AS.Setup( $1^\lambda$ )  $\rightarrow pp_k$  and runs
4     AS.KeyGen(ppk)  $\rightarrow (sk_k, vk_k)$ .
4   Create a tuple  $m_k = (RN\text{-}Update, pp_k, vk_{RI_k}, currMax_k^s,$ 
5   currMaxrk, hopCount = 0, hopMax, ts) for each
6   neighbor j, j  $\in [1..l]$  where l is the total number of
7   neighbors of k. Create  $\sigma'_k \leftarrow \text{Sign}(sk_{RI_k}, m_k)$  and set
8    $m'_k = (m_k, \sigma'_k)$  Create signature
9    $\sigma_k \leftarrow \text{AS.Sign}(sk_k, \perp, \perp, \perp, m'_k)$ .
5   k sends  $M = ((m'_k), (vk_k), \sigma_k)$  to each neighbor j.
6 for each node i in the network on receiving an RN-Update
7   message from neighbor j do
8   On receiving  $M = ((m'_k, \dots, m_j), (vk_k, \dots, vk_j), \sigma_j)$ , i
9   parses  $(m_k, \sigma'_k) \leftarrow m'_k$  and  $(RN\text{-}Update, pp_k, VK_k,$ 
10  currMaxsk, currMaxrk, hopCount, hopMax, ts) \leftarrow m_k.
11  if Verify( $m_k, VK_k, \sigma'_k$ )  $\rightarrow 0$   $\vee$  (AS.Verify( $(m_k, \dots,$ 
12   $m_j), (vk_k, \dots, vk_j), \sigma_j$ )  $\rightarrow 0$ ) then
13    | Return  $\perp$ .
14    i checks that hopCount value in all messages
15     $(m'_k, \dots, m_j)$  are incremented by 1 in each message. If
16    not, return  $\perp$ .
17    i runs AS.KeyGen(ppk)  $\rightarrow (sk_i, vk_i)$ .
18    i updates its local routingTable for RN k and neighbor j
19    by updating the expiry time  $t_e = ts + e$ , currMaxsk, and
20    currMaxrk.
21  if hopCount of received message is equal to hopMax in
22   $m_k$  then
23    | Return  $\perp$ .
24  else
25    for each neighbor s do
26      i creates  $m_i$  by updating contents of  $m_j$  as
27      hopCount = hopCount + 1,
28      currMaxsk = min(currMaxsk,  $lw_{i,s}$ ), and
29      currMaxrk = min(currMaxrk,  $lw_{s,i}$ ).
30      i creates signature  $\sigma_i \leftarrow \text{AS.Sign}(sk_i, \sigma_j,$ 
31       $(m'_k, \dots, m_j), (vk_k, \dots, vk_j), m_i$ ).
32      i sets  $M = ((m'_k, \dots, m_j, m_i), (vk_k, \dots,$ 
33       $vk_j, vk_i), \sigma_i)$  and sends it to neighbor r.

```

pp (we have not shown this simple step for presentation clarity). If RNs do not selectively forward to certain neighbors, we do not consider it as malicious behavior. The regular nodes within a given RN's hopMax radius will receive the RN's broadcasted messages from other neighbors in the neighborhood.

The next issue is nodes underreporting or overreporting currMax_s and currMax_r. We do not consider nodes underreporting currMax_s and currMax_r as malicious behavior since every node can individually decide the amount of funds to commit on its own links. If nodes overreport currMax_s, currMax_r to a value greater than that of their

Algorithm 10: Multisig Exchange

Input : $o \in \{\perp, Rev\}$, $t \in \{hold_s \mid hold_r \mid pay\}$, j, SK_j, VK_j ,
 $k, SK_k, VK_k, lw_{jk}, v, txid, ts$

- 1 **if** $o == Rev$ **then**
- 2 j and k discard currently stored contracts for $txid$ and
 delete $fw_{jk}.txid$ and $fw_{kj}.txid$.
- 3 **return**
- 4 **if** $t == pay$ **then**
- 5 j and k set $lw_{jk} = fw_{jk}.txid$ and $lw_{kj} = fw_{kj}.txid$.
- 6 **return**
- 7 **if** $t == hold_s$ **then**
- 8 Set $fw_{jk} = lw_{jk} - v$. Set $fw_{kj} = lw_{kj} + v$.
- 9 **if** $t == hold_r$ **then**
- 10 Set $fw_{jk} = lw_{jk} + v$. Set $fw_{kj} = lw_{kj} - v$.
- 11 j sends $\sigma_j \leftarrow \text{Sign}_{SK_j}$ (contract = $(lw_{jk}, lw_{kj}, fw_{jk}, fw_{kj}, txid, digest, ts)$ to k).
- 12 k sends $\sigma_k \leftarrow \text{Sign}_{SK_k}$ (contract = $(lw_{jk}, lw_{kj}, fw_{jk}, fw_{kj}, txid, digest, ts)$ to j).
- 13 **if** $\text{Verify}_{VK_k}(\text{contract}, \sigma_k) \stackrel{?}{=} 1$ **then**
- 14 j stores $(\sigma_j, \sigma_k, \text{contract})$, $fw_{kj}.txid = fw_{kj}$ and
 $fw_{jk}.txid = fw_{jk}$.
- 15 **if** $\text{Verify}_{VK_j}(\text{contract}, \sigma_j) \stackrel{?}{=} 1$ **then**
- 16 k stores $(\sigma_j, \sigma_k, \text{contract})$, $fw_{kj}.txid = fw_{kj}$ and
 $fw_{jk}.txid = fw_{jk}$.
- 17 **if** $t == hold_s$ **then**
- 18 j updates the $\text{currMax}_s = \min(fw_{jk}, \text{currMax}_s)$ for all
 paths going through k . k updates the
 $\text{currMax}_r = \min(fw_{kj}, \text{currMax}_r)$ for all paths going
 through j .
- 19 **if** $t == hold_r$ **then**
- 20 j updates the $\text{currMax}_r = \min(fw_{jk}, \text{currMax}_r)$ for all
 paths going through k . k updates the
 $\text{currMax}_s = \min(fw_{kj}, \text{currMax}_s)$ for all paths going
 through j .

own links, that is malicious behavior. Due to privacy concerns, nodes' link weights cannot, of course, be verified by anyone, but overreporting will eventually cause transaction failure (since there was no actual liquidity) and result in revoked hold/pay contracts with penalties for the misbehaving node. In any case, no node will lose money. The AS scheme helps verify that the currMax values do not increase in the series of aggregated messages to help identify malicious nodes in the network as well. A malicious node cannot increase the $\text{currMax}_s, \text{currMax}_r$ value signed by the RN as part of the first aggregated message because the first message is signed by the RN using its publicly verifiable signing key.

The other potential source of malicious behavior is nodes underreporting or overreporting hopCount values. First note that the hopCount is contained in every message $((m'_k, \dots, m_j), \text{Line 7 in Algorithm 1})$ that is aggregated in the signature. Any honest node along a path can verify that the hopCount contained within every message is incremented by one, starting with $m'_k = 0$ (thus reducing hopCounts would be immediately detected, and the *RN-Update*

message discarded). Inflating hopCounts would not be in the best interest of the malicious node(s) because honest nodes could have alternative shorter paths to the intended target node.

Concerning perimeter nodes, two situations could arise: Case 0: A regular node pretends to be a perimeter node by overreporting its hopCount. In this case, that node's nonce will not figure in the set intersection of two RNs since the node was not actually a perimeter node. The node cannot do anything further. Case 1: A perimeter node underreports its hopcount or drops a message. We do not consider this malicious behavior, since it just means that the node does not wish to participate in transactions. Since the *RN-Update* messages are broadcasted, RNs will get replies from other perimeter nodes. Even if an RN does not pad its nonce list with random nonces (Algorithm 2, Line 2), it will not leak the identity of its perimeter nodes to other RNs, although it will reveal the number of perimeter nodes that RN has paths to.

Hold phase: If an honest node along a path does not receive *holdACK* or *holdReject* messages for a given transaction before the expiry of its timer te_1 , the transaction will time out and will have to be retried. Malicious nodes can try to change the message type (the first field), but unknown message types will get dropped by honest nodes along a path. Malicious nodes might also try to change the "Y" parameter denoting the identity of the next RN or perimeter node to forward messages to (Algorithm 5, Case 1, 2). The message will be held at the misdirected RN/perimeter node which could also be potentially malicious. But eventually, the *hold* phase for that segment will timeout, and the *hold* contracts will be rolled back. Other parameters such as *hopMax*, *digest* being modified, or $C_{RN}(\cdot)$ being re-encrypted (Algorithm 5, Case 1) will result in the *hold* messages being misdirected, but the *hold* phase times out, and we will not get to the *pay* phase.

A malicious RN_s cannot misroute a hold message tuple to an RN'_r instead of the sender's selected RN_r, e.g., by creating an incorrect onion. This is because Bob's *hold_r* will be sent to RN_r, and since RN'_r never received it, the misrouted transaction will eventually time out, and any signed contracts will be rolled back. Similarly, no malicious node, including RNs can increase/decrease the transaction amount *v* to an arbitrary value, because: 1) since the receiver knows the correct amount, the hold will eventually timeout at the last hop and fail. 2) All honest nodes along the path will have to commit to paying the amount in the hold phase. Any honest nodes which receive a *pay* message with a transaction amount different from the original *hold* message will refuse to proceed with the *pay* phase, hence timing out the transaction and causing a rollback of contracts.

The one thing that a malicious RN_s could potentially do is increase the path length to RN_r by several more RNs than is required. The transaction will eventually reach RN_r via a longer path in the RN_s – RN_r segments. Potential solutions include the sender specifying a maximum number of layers in the onion encryption at RN_s, based on periodic network statistics released by the RNs. We leave incorporating such mechanisms into SPRITE as future work.

Pay phase: If a node intentionally misroutes the pay tuple or does not forward it, resulting in the pay tuple not reaching the target node on time, te_2 timer will expire, causing nodes to time out and rollback their pay contracts. In case of any other malicious activity, the *hold* contract signed in the previous phase can be enforced.