Privacy-preserving Multi-hop Locks for Blockchain Scalability and Interoperability*

Giulio Malavolta†
Friedrich-Alexander-University
Erlangen-Nürnberg
malavolta@cs.fau.de

Pedro Moreno-Sanchez†‡
TU Wien
pedro.sanchez@tuwien.ac.at

Clara Schneidewind
TU Wien
clara.schneidewind@tuwien.ac.at

Aniket Kate
Purdue University
aniket@purdue.edu

Matteo Maffei
TU Wien
matteo.maffei@tuwien.ac.at

***Abstract*— Tremendous growth in cryptocurrency usage is exposing the inherent scalability issues with permissionless blockchain technology. *Payment-channel networks* (PCNs) have emerged as the most practically deployed solution to mitigate the scalability issues, allowing the bulk of payments between two users to be carried out off-chain. Unfortunately, as reported in the literature and further demonstrated in this paper, current PCNs do not provide meaningful security and privacy guarantees.**

In this work, we study and design secure and privacy-preserving PCNs. We start with a security analysis of existing PCNs, reporting a new attack that applies to all major PCNs, including the Lightning Network, and allows an attacker to steal the fees from honest intermediaries in the same payment path. We then formally define privacy-preserving multi-hop locks (PrivMuLs), a novel cryptographic primitive that serves as a cornerstone for the design of secure and privacy-preserving PCNs. We present several provably secure cryptographic instantiations that make PrivMuLs compatible with the vast majority of cryptocurrencies. In particular, we show that (linear) homomorphic one-way functions suffice to construct PrivMuLs for PCNs supporting such functions in their script language (e.g., Ethereum). We also propose a construction based on ECDSA signatures that *does not require scripts*, thus solving a prominent open problem in the field. PrivMuLs constitute a generic primitive whose usefulness goes beyond multi-hop payments in a single PCN and we show how to realize atomic swaps and interoperable PCNs from this primitive. Finally, our performance evaluation on a commodity machine finds that PrivMuLs operations can be performed in less than 100 milliseconds and require less than 500 bytes of communication overhead, even in the worst case. In fact, after acknowledging our attack, the Lightning Network developers are right now integrating ECDSA-based PrivMuLs into their PCN. This demonstrates the practicality of our approach and its impact on the security, privacy, interoperability, and scalability of today's cryptocurrencies.

*This is a draft (revision October 23, 2018)

†Both authors contributed equally and are considered to be co-first authors.

‡ This work was done while this author was at Purdue University.

I. Introduction

Cryptocurrencies are growing in popularity and are playing an increasing role in the worldwide financial ecosystem. In fact, the number of Bitcoin transactions grew by approximately 30% in 2017, reaching a peak of more than $420,000$ transactions per day in December 2017 [2]. This striking increase in demand has given rise to scalability issues [20], which go well beyond the rapidly increasing size of the blockchain. For instance, the permissionless nature of the consensus algorithm used in Bitcoin today limits the transaction rate to tens of transactions per second, whereas other payment networks such as Visa support peaks of up to 47,000 transactions per second [54].

Among the various proposals to solve the scalability issue [22], [23], [38], [49], *payment-channels* have emerged as the most widely deployed solution in practice. In a nutshell, two users open a payment channel by committing a single transaction to the blockchain, which is meant to lock their bitcoins in a deposit secured by a Bitcoin (smart) contract. These users can then perform several payments between each other without the need for additional blockchain transactions, by simply locally agreeing on the new deposit balance. A transaction is required only at the end in order to close the payment channel and unlock the final balances of the two parties, thereby drastically reducing the transaction load on the blockchain. Further research has proposed the concept of *payment-channel network* [49] (PCN), where two

users not sharing a payment channel can still pay each other using a path of open channels between them. Unfortunately, as we discuss below in detail, current PCNs fall short of providing adequate security, privacy, and interoperability guarantees.

A. State-of-the-art in PCNs

Several practical deployments of PCNs exist today [6], [10], [11] based on a common reference description for the Lightning Network (LN) [8]. Unfortunately, this proposal is neither privacy-preserving, as shown in recent works [30], [40], nor secure, which stays in contrast to what until now was commonly believed, as we show in this work. In fact, we present a new attack, the wormhole attack, which applies not only to the LN, the most widely deployed PCN, but also other PCNs based on the same cryptographic lock mechanism, such as the Raiden Network [9].

PCNs have attracted plenty of attention also from academia. Malavolta et al. [40] proposed a secure and privacy-preserving protocol for multi-hop payments. However, this solution is expensive as it requires non-trivial amount of data (i.e., around 5 MB) to be exchanged between the users in the payment path and it also hinders interoperability as it requires the Hash Time-Lock Contract (HTLC) to be available in the cryptocurrency.

Green and Miers presented BOLT, a hub-based privacy-preserving payment for PCNs [30]. BOLT requires cryptographic primitives only available in Zcash and it cannot be seamlessly deployed in Bitcoin. Moreover, this approach is limited to paths with a single intermediary and the extension to support of arbitrary length remains an open problem.

The rest of the existing PCN proposals suffer from similar drawbacks. Apart from not formalizing provable privacy guarantees, they are restricted to a setting with a trusted execution environment [36] or with a Turing complete scripting language [25], [26], [33], [43] so that they cannot seamlessly work with prominent cryptocurrencies today (except for Ethereum).

Poelstra introduced the notion of scriptless scripts, a modified version of a digital signature scheme so that a signature can only be created faithfully fulfilling a cryptographic condition [48]. The resulting signature is verifiable following the unmodified digital signature scheme. When applied to script-based systems like Bitcoin or Ethereum, they are accompanied by core scripts (e.g., script to verify the signature itself). This approach reduces the space required for cryptographic operations in the script, saving thus invaluable bytes in the blockchain. Moreover, it improves upon the fungibility of the cryptocurrency as transactions from payment channels no longer require a script different from other payments.

Although interesting, current proposals [48] lack formal security and privacy treatment and are based only on the Schnorr signature scheme, therefore being incompatible with major cryptocurrencies like Bitcoin. In fact, there are early proposals for Schnorr adoption in Bitcoin [55], but it is unclear if and when they will be realized.

In summary, existing proposals are neither generically applicable nor interoperable, since they rely on specific features (e.g., contracts) of individual cryptocurrencies or trusted hardware. Furthermore, there seems to be a gap between secure realization of PCNs and what is developed in practice, as we demonstrate with our attack, which affects virtually all PCNs deployed in practice.

B. Our contributions

In this work, we contribute to the rigorous understanding of PCNs and present the first interoperable, secure, and privacy-preserving cryptographic construction for PrivMuLs. Specifically,

- We analyze the security of existing PCNs, reporting a new attack (the *wormhole attack*) which allows dishonest users to steal the payment fees from honest users along the path (Section III). This attack applies to the LN, as well as any decentralized PCN where the sender does not know in advance the intermediate users along the path to the receiver. We communicated the attack to the LN developers, who acknowledged the issue.
- In order to construct secure and privacy-preserving PCNs, we introduce a novel cryptographic primitive called privacy-preserving multi-hop lock (PrivMuL). We model the security of such a primitive in the UC framework [19] to inherit the underlying composability guarantees (Section IV). Then we show that PrivMuLs can be generically combined with any blockchain to construct a fully-fledged PCN.
- As a theoretical insight emerging from the wormhole attack, we establish a lower bound on the communication complexity of secure PCNs that follow our UC definition: Specifically, we show that an extra round of communication to determine the path is necessary to have a secure transaction.
- We show how to realize PrivMuLs in different settings (Section V). In particular, we demonstrate that (linearly) homomorphic operations suffice to build any script-based PrivMuL. Furthermore, we show how to

realize PrivMuLs in a scriptless setting. This approach is of special interest because it reduces the transaction size, and, consequently, the blockchain load. We give a concrete construction based on the ECDSA signature, solving a prominent problem in the literature [48]. This makes PrivMuLs compatible with the vast majority of cryptocurrencies (including Bitcoin and Ethereum). In fact, PrivMuLs are being implemented right now in the LN [7], [28].

- We implemented our cryptographic constructions and show that they require at most 60 milliseconds to be computed and a communication overhead of less than 500 bytes in the worst case (Section VI). These results demonstrate that PrivMuLs are practical and ready to be deployed. In fact, PrivMuLs can be leveraged to design atomic swaps and interoperable (cross-currency) PCNs as well (Section VII).

II. Context: Payment Channel Networks

A. Payment Channels

A payment channel allows two users to exchange bitcoin without committing every single payment to the Bitcoin blockchain. For that, users first publish an on-chain transaction to deposit bitcoin into a multi-signature address controlled by both users. Such deposit also guarantees that all bitcoin are refunded at a possibly different but mutually agreed time if the channel expires. Users can then perform off-chain payments by adjusting the distribution of the deposit (that we will refer to as *balance*) in favor of the payee. When no more off-chain payments are needed (or the capacity of the payment channel is exhausted), the payment channel is closed with a *closing* transaction included in the blockchain. This transaction sends the deposited bitcoin to each user according the most recent balance in the payment channel. We refer the reader to [22], [23], [41], [49] for further details.

B. A Payment Channel Network (PCN)

A PCN can be represented as a directed graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, where the set $\mathbb{V}$ of vertices represents the Bitcoin accounts and the set $\mathbb{E}$ of weighted edges represents the payment channels. Every vertex $U \in \mathbb{V}$ has associated a non-negative number that denotes the fee it charges for forwarding payments. The weight on a directed edge $(U_1, U_2) \in \mathbb{E}$ denotes the amount of remaining bitcoin that U_1 can pay to U_2.

A PCN is used to perform off-chain payments between two users with no direct payment channel between them but rather connected by a path of open payment channels. For that, assume that S wants to pay α bitcoin to R, which is reachable through a path of the form $S \to U_1 \to \ldots \to U_n \to R$. For their payment to be successful, every link must have a capacity $\gamma_i \geq \alpha'_i$, where $\alpha'_i = \alpha - \sum_{j=1}^{i-1} \mathit{fee}(U_j)$ (i.e., the initial payment value minus the fees charged by intermediate users in the path). If the payment is successful, edges from S to R are decreased by α'_i. Importantly, to ensure that R receives exactly α bitcoin, S must start the payment with a value $\alpha^* = \alpha + \sum_{j=1}^{n} \mathit{fee}(U_j)$. We refer the reader to [30], [40], [41], [49] for further details.

The concepts of payment channels and PCNs have already attracted considerable attention from academia [23], [30], [31], [35], [40], [41], [43]. In practice, the Lightning Network (LN) [8], [49] has emerged as the most prominent example. Currently, there exist several independent implementations of the LN for Bitcoin [6], [10], [11]. Moreover, the LN is also considered as a scalability solution in other blockchain-based payment systems such as Ethereum [9].

C. Multi-Hop Payments Atomicity

A fundamental property that multi-hop payments have to fulfill is *atomicity*: Either the capacity of all channels in the path is updated or none of the channels is changed. Partial updates can lead to coin losses for the users on the path. For instance, a user could pay a certain amount of bitcoin to the next user in the path but never receive the corresponding bitcoin from the previous neighbour. The LN tackles this challenge by relying on a smart contract called *Hash Time-Lock Contract* (HTLC) [49]. This contract locks x bitcoin that can be released only if the contract's condition is fulfilled. The contract is defined in terms of a hash value $y := H(R)$ where R is chosen uniformly at random, the amount of bitcoin x, and a timeout t, as follows:

HTLC (Alice, Bob, y, x, t):
1) If Bob produces the condition R^* such that $H(R^*) = y$ before t days, Alice pays Bob x bitcoin.
2) If t days elapse, Alice gets back x bitcoin.

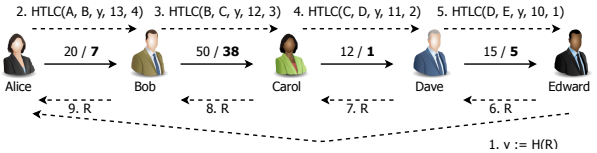

Fig. 1: Illustrative example of a payment from Alice to Edward for value 10 using HTLC contract. Non bold (bold) numbers represent the capacity of payment channels before (after) the payment. We assume all users to charge a fee of 1 bitcoin.

We depict in Fig. 1 an illustrative example of the use of HTLC in a payment. For ease of exposition, we assume that every user charges a fee of one bitcoin and the payment amount is 10 bitcoin. In this payment, Edward first sets up the payment by creating a random value R and sending $H(R)$ to Alice. Then, the *commitment phase* starts by Alice. She first sets on hold 13 bitcoin and then successively every intermediate user sets on hold the received amount minus his/her own fee. After Dave set 10 coins on hold with Edward, the latter knows that the corresponding payment amount is on hold at each channel and he can start the *releasing phase*. For that, he reveals the value R to Dave allowing him to fulfill the HTLC contract and settle the new capacity at the payment channel. The value R is then passed backwards in the path allowing the settlement of each payment channel in the path.

Privacy issues in PCNs. Recent works in the literature [30], [40] show that the current use of HTLC leaks a common identifier along the payment path (i.e., the condition $H(R^*)$) that can be used by an adversary to tell who pays to whom. Current solutions to this privacy issue are expensive in terms of computation and communication [40] or incompatible with major cryptocurrencies deployed today [30]. This calls for a more in-depth study of this crucial cryptographic tool.

III. Wormhole Attack in Existing PCNs

In a nutshell, the wormhole attack allows two colliding users on a payment path to exclude intermediate users from participating in the successful completion of a payment, thereby stealing the payment fees which were intended for honest path nodes.

In more detail, assume a payment path $(U_0, \ldots, U_i, \ldots, U_j, \ldots U_n)$ used by U_0 to pay an amount $\alpha + \sum_k \gamma_k$ to U_n, where $\gamma_k = fee(U_k)$ denotes the fee charged by the intermediate user U_k as a reward for enabling the payment. Further assume that U_i and U_j are two adversarial users that may deviate from the protocol if some economic benefit is at stake. The adversarial strategy is as follows.

In the commitment phase, every user behaves honestly. This, in particular, implies that every honest user has locked a certain amount of coins in the hope of getting rewarded for this. In the releasing phase, honest users $U_{j+1}, \ldots, U_n$ correctly fulfill their HTLC contracts and settle the balances and rewards in their corresponding payment channels.

The user U_j behaves honestly with U_{j+1} effectively settling the balance in their payment channel. On the other hand, U_j waits until the timeout set in the HTLC with U_{j-1} is about to expire and then agrees with U_{j-1} to cancel the HTLC and set the balance in their payment channel back to the last agreed one. Note that from U_{j-1}'s point of view, this is a legitimate situation (e.g., there might not be enough coins in a payment channel at some user after U_j and the payment had to be canceled). Moreover, the channel between U_{j-1} and U_j does not need to be closed, it is just rolled back to a previous balance, a feature present in the Lightning Network.

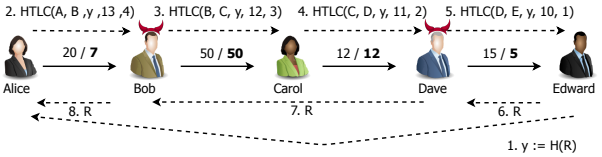

Fig. 2: Illustrative example of the attack in the LN. For simplicity, we assume that each intermediate node charges a payment fee of 1 bitcoin. Here, Bob and Dave are colluding. Dave and Bob exclude Carol from the successful completion of the payment, effectively stealing her payment fee.

As U_{j-1} believes that the payment did not go through, she also cancels the HTLC with U_{j-2}, who in turns cancels the HTLC with U_{j-3} and so on. This process continues until U_i is approached by U_{i+1}. Here, U_i cancels the HTLC with U_{i+1}. However, U_i gets the releasing condition R from U_j and can use it to fulfill the HTLC with U_{i-1} and therefore settle the new balance in that payment channel. Therefore, from the point of view of users $U_1, \ldots, U_{i-1}$, the payment has been successfully carried out. An illustrative example of this attack in the Lightning Network is shown in Fig. 2.

Discussion. An adversary controlling users U_i and U_j in a payment path that carries out the attack described in this section gets an overall benefit of $\sum_{k=i+1}^{j} \gamma_k$ bitcoins instead of only $\gamma_i + \gamma_j$ bitcoins in the case he behaves honestly. We make several observations here. First, the impact of this attack grows with the number of intermediate users between U_i and U_j as well as the number of payments that take both U_i and U_j in their path. Second, honest intermediate users cannot trivially distinguish the situation in which they are under attack from the situation where the payment is simply unsuccessful (e.g., there are not enough coins in one of the channels or one of the users is offline). In both cases, the view for the honest users is that the timeout established in the HTLC is reached, the payment failed and they get their initially committed coins reimbursed. In short, the wormhole attack allows an adversary to steal the fees from intermediate honest users without

leaving a inculpatory trace to them. Third, fees are the main incentive for users to act as intermediaries. The wormhole attack takes away this crucial benefit. In fact, this attack not only makes honest users lose their fees, but also incur collateral costs: Coins locked for the payment under attack cannot be used for another (possibly successful) payment simultaneously.

Fourth, while the Lightning Network is at its infancy, other well-established networks such as Ripple use paths with multiple intermediaries. For instance, in the Ripple network, more than 27% of the payments use more than two intermediaries [44]. Actually, paths with three intermediaries (e.g., sender $\rightarrow$ bank $\rightarrow$ currency-exchange $\rightarrow$ bank $\rightarrow$ receiver) are essential for currency exchanges, a key use case in LN itself [1]. When the LN grows to the scale of Internet, routes may consist of several intermediaries as in Internet today. Given these evidences, we expect to have long paths in the LN.

Responsible Disclosure. We notified this attack to the LN developers and they have acknowledged this issue. They are currently implementing our proposed solution to overcome the wormhole attack [28].

(In)evitability of the Wormhole Attack. The wormhole attack is not restricted to the LN, but generally applies to PCNs with multi-hop payments that involve only two rounds of communication. We assume a communication round to consist of traversing the payment path once, either forth (e.g., for setting up the payment) or back (e.g., for releasing the money). Additionally, we assume that in PCNs the communication between nodes is restricted to their direct neighbors, so in particular, there is no broadcast.[1] Consequently, using two rounds of communication for a payment implies that the payment is not preceded by a routing phase in which path-specific information is sent to nodes in the path. Under these assumptions, we state the lower bound in Theorem 1 and, due to the lack of space, we defer the proof to the extended version of the paper [5].

Theorem 1 (Inevitability of the wormhole attack)**.** *For all two-round (without broadcast channels) multi-hop payment protocols there exists a path prone to the wormhole attack.*

In this work we show that adding an additional round of communication suffices to overcome this impossibility result. In particular, with one additional round of communication, the sender of a payment can communicate path-specific secret information to the intermediate nodes. This information can then be used to make the release keys unforgeable for an attacker. The cryptographic protocols we introduce in the remainder of this paper adopt this approach.

[1]This is the case in the setting of off-chain protocols where users not sharing a payment channel do not communicate with each other.

IV. Definition

In the following we introduce a new cryptographic primitive called privacy-preserving multi-hop lock (PrivMuL). This primitive generalizes the locking mechanism used for payments in state-of-the art PCNs such as the Lightning Network. In Section VII we show that PrivMuLs are the only cryptographic component required to construct fully-fledged PCNs. As motivated in the previous section, we model the primitive such that it allows for an initial setup phase in which the first node of the path provides the other nodes on the path with some secret (path-specific) state. Formally, a PrivMuL is defined with respect to a universe of users $\mathbb{U}$ and it is a five-tuple of PPT algorithms and protocols $\mathbb{L} = (\mathsf{KGen}, \mathsf{Setup}, \mathsf{Lock}, \mathsf{Rel}, \mathsf{Vf})$ defined as follows:

Definition 1. *A* PrivMuL $\mathbb{L} = (\mathsf{KGen}, \mathsf{Setup}, \mathsf{Lock}, \mathsf{Rel}, \mathsf{Vf})$ *consists of the following efficient algorithms:*

$\{(\mathsf{sk}_i, \mathsf{pk}), (\mathsf{sk}_j, \mathsf{pk})\} \leftarrow \langle \mathsf{KGen}_{U_i}(1^\lambda), \mathsf{KGen}_{U_j}(1^\lambda) \rangle$: *On input the security parameter 1^λ the key generation protocol returns a shared public key* pk *and a secret key* sk_i *(*sk_j*, respectively) to U_i and U_j.*

$\{s_0^I, \ldots, (s_n^I, k_n)\} \leftarrow \langle \mathsf{Setup}_{U_0}(1^\lambda, U_1, \ldots, U_n) \ldots \mathsf{Setup}_{U_n}(1^\lambda) \rangle$: *On input a set of identities $(U_1, \ldots, U_n)$ and the security parameter 1^λ, the setup protocol returns, for $i \in [0, n]$, a state s_i to each party U_i. The receiver U_n additionally receives a key k_n.*

$\{(\ell, s_i^R), (\ell, s_{i+1}^L)\} \leftarrow \langle \mathsf{Lock}_{U_i}(s_i^I, \mathsf{sk}_i, \mathsf{pk}), \mathsf{Lock}_{U_{i+1}}(s_{i+1}^I, \mathsf{sk}_{i+1}, \mathsf{pk}) \rangle$: *On input two initial states s_i^I and s_{i+1}^I, two secret keys* sk_i *and* sk_{i+1}*, and a public key* pk*, the locking protocol is executed between two parties (U_i, U_{i+1}) and returns a lock ℓ and a right state s_i^R to U_i and the same lock ℓ and a left state s_{i+1}^L to U_{i+1}.*

$k' \leftarrow \mathsf{Rel}(k, (s^I, s^L, s^R))$: *On input an opening key k and a set of states (s^I, s^L, s^R), the release algorithm returns a new opening key k'.*

$\{0, 1\} \leftarrow \mathsf{Vf}(\ell, k)$: *On input a lock ℓ and a key k the verification algorithm returns a bit $b \in \{0, 1\}$.*

Correctness. A PrivMuL is *correct* if the verification algorithm Vf always accepts a honestly generated lock-key pair. For a more detailed and formal correctness definition, we refer the reader to the extended version [5].

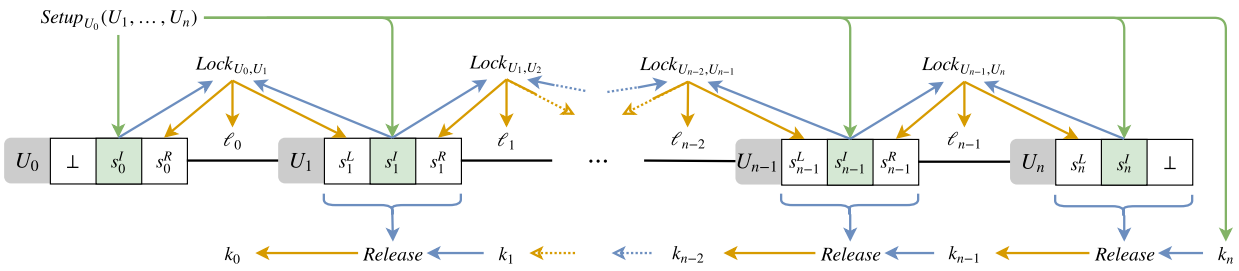

Fig. 3: Illustration of the usage of the PrivMuL primitive. It is assumed that links between the users on the path have been created upfront using the KGen algorithm and that the resulting public and secret keys are implicitly given as argument to the corresponding executions of the Lock protocol. Otherwise, the inputs (outputs) to (from) the Lock protocol and the Rel algorithm are indicated by blue (orange) arrows.

Key Ideas. Fig. 3 illustrates the usage of the different protocols underlying the PrivMuL primitive: We assume an initial phase where a network of users is generated using the (interactive) KGen protocol for establishing pairwise links between users. Consequently, all users in this network are assumed to hold a secret key and a shared public key for each link that they created. We recreate thereby the opening of payment channels that compose the PCN.

In the setup phase (depicted in green), the sender U_0 decides upon a path of users $(U_0, \ldots, U_n)$ and executes the Setup protocol with the nodes of the path. Each user U_i on the path learns its initial state s_i^I and the receiver U_n additionally learns the initial opening key k_n. The introduction of the initial state at each intermediate user is crucial for security and privacy. As shown in Section III, the lack of initial state at intermediate users inevitably enables the wormhole attack. Intuitively, we can use this initial state as "rerandomization factor" to ensure that locks in the same path are unlinkable for the adversary.

Next, in the locking phase, each pair of users jointly executes the pairwise Lock protocol, starting from U_0. Two users U_i and U_{i+1} run the Lock protocol on their initial states s_i^I and s_{i+1}^I to generate a lock ℓ_i. The creation of this lock represents the commitment from U_i to perform an application-dependent action if a cryptographic problem is solved by U_{i+1}. In the case of LN, this operation represents the commitment of U_i to pay a certain amount of coins to U_{i+1} if U_{i+1} solves the cryptographic condition. Each user also learns some secret information s_i^R (resp. s_{i+1}^L) that will be needed for releasing the lock later on. As the locks are directed, the usage of this information will differ depending on whether it was gained while creating a lock with a preceding (left) or with a subsequent (right) neighbor on the path: As U_i creates a lock with its right neighbor, it learns the state s_i^R denoting the state of its *right lock*. Correspondingly, U_{i+1} learns the state s_i^L denoting the state of its *left lock*. While these extra states are not present in the LN (i.e., every lock is based on the same cryptographic puzzle $H(R)$), having them is crucial for security. They make the releasing of different locks in the path independent from each other and therefore ensure that a lock ℓ_i can only be released if ℓ_{i+1} has been released before.

Finally, after the entire path is locked, the receiver U_n can use the information it received during the setup phase (k_n and s_n^I) together with the information it learned from creating its left lock (s_n^L) to generate a key for releasing its left lock. In the same fashion each intermediate node can use its state (containing the initial state as well as the states from locking) to derive a valid key for its left lock from a valid key for its right lock using the Rel algorithm. This last phase resembles the opening phase of the LN where each pair of users settles the new balances for their deposit at each payment channel in the payment path.

A. Security and Privacy Definition

To model security and privacy in the presence of concurrent executions we resort to the universal composability framework from Canetti [19]. We allow thereby the composition of PrivMuLs with other application-dependent protocols while maintaining security and privacy guarantees. We model the players in our protocol as interactive Turing machines that communicate with a trusted functionality $\mathcal{F}$ via secure and authenticated channels. We model the attacker $\mathcal{A}$ as a PPT machine that corrupts a subset of users prior the execution of the interaction. Upon corruption of a user U, the attacker is provided with the internal state of U and the incoming and outgoing communication of U is routed through $\mathcal{A}$.

KeyGen$(U_j, \{L, R\})$

Upon invocation by U_i:
send $(U_i, \{L, R\})$ to U_j
receive b from U_j
if $b = \perp$ send $\perp$ to U_i and abort
if L insert (U_i, U_j) into $\mathcal{U}$ and send (U_i, U_j) to U_i
if R insert (U_j, U_i) into $\mathcal{U}$ and send (U_j, U_i) to U_i

Lock(*lid*)

Upon invocation by U_i:
if *getStatus*(*lid*) $\neq$ Init or *getLeft*(*lid*) $\neq U_i$ then abort
send (lid, Lock) to *getRight*(*lid*)
receive b from *getRight*(*lid*)
if $b = \perp$ send $\perp$ to U_i and abort
updateStatus(*lid*, Lock)
send (lid, Lock) to U_i

GetStatus(*lid*)

Upon invocation by U_i:
return $(lid,$ *getStatus*$(lid))$ to U_i

Setup$(U_0, \ldots, U_n)$

Upon invocation by U_0:
if $\forall i \in [0, n-1] : (U_i, U_{i+1}) \notin \mathcal{U}$ then abort
$\forall i \in [0, n-1]$: $lid_i \leftarrow\!\!\$\ \{0,1\}^\lambda$
insert $(lid_0, U_0, U_1, \mathsf{Init}, lid_1), (lid_{n-1}, U_{n-1}, U_n, \mathsf{Init}, \perp)$ into $\mathcal{L}$
send $(\perp, lid_0, \perp, U_1, \mathsf{Init})$ to U_0
send $(lid_{n-1}, \perp, U_{n-1}, \perp, \mathsf{Init})$ to U_n
$\forall i \in [1, n-1]$: insert $(lid_i, U_i, U_{i+1}, \mathsf{Init}, lid_{i+1})$ into $\mathcal{L}$
send $(lid_{i-1}, lid_i, U_{i-1}, U_{i+1}, \mathsf{Init})$ to U_i

Release(*lid*)

Upon invocation by U_i:
if *getRight*(*lid*) $\neq U_i$ or *getStatus*(*lid*) $\neq$ Lock or
getStatus(*getNextLock*(*lid*)) $\neq$ Rel
and *getNextLock*(*lid*) $\neq \perp$ then abort
updateStatus(*lid*, Rel)
send (lid, Rel) to *getLeft*(*lid*)

Fig. 4: Ideal functionality for PrivMuLs

Let $\mathsf{EXEC}_{\tau,\mathcal{A},\mathcal{E}}$ be the ensemble of the outputs of the environment $\mathcal{E}$ when interacting with the adversary $\mathcal{A}$ and parties running the protocol τ (over the random coins of all the involved machines).

Definition 2 (Universal Composability)**.** *A protocol τ UC-realizes an ideal functionality $\mathcal{F}$ if for any PPT adversary $\mathcal{A}$ there exists a simulator $\mathcal{S}$ such that for any environment $\mathcal{E}$ the ensembles $\mathsf{EXEC}_{\tau,\mathcal{A},\mathcal{E}}$ and $\mathsf{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{E}}$ are computationally indistinguishable.*

Ideal Functionality. We formally define the ideal world functionality $\mathcal{F}$ for PrivMuLs in the following. For a more modular treatment, our UC definition models only the cryptographic lock functionality, rather than aiming at a comprehensive characterization of PCN. In Section VII we show how one can construct a full PCN (e.g., as defined in [40]) by composing this functionality with time locks, balance updates, and on-chain channel management. For ease of exposition we assume that each pair of users establishes only a single link per direction. The model can be easily extended to handle the more generic case. $\mathcal{F}$ works in interaction with a universe of users $\mathbb{U}$ and initializes two empty lists $(\mathcal{U}, \mathcal{L}) := \emptyset$, which are used to track the users and the locks, respectively. The list $\mathcal{L}$ represents a set of lock chains. The entries are of the form $(lid_i, U_i, U_{i+1}, f, lid_{i+1})$ where lid_i is a lock identifier that is unique even among other lock chains in $\mathcal{L}$, U_i and U_{i+1} are the users connected by the lock, $f \in \{\mathsf{Init}, \mathsf{Lock}, \mathsf{Rel}\}$ is a flag that represents the status of the lock, and lid_{i+1} is the identifier of the next lock in the path. For sake of better readability, we define functions operating on $\mathcal{L}$ extracting lock-specific information given the lock's identifier, such as the lock's status (*getStatus*$(\cdot)$), the nodes it is connecting (*getLeft*$(\cdot)$, *getRight*$(\cdot)$), and the next lock's identifier (*getNextLock*$(\cdot)$). In addition we define an update function *updateStatus*$(\cdot, \cdot)$ that changes the status of lock to a new flag.

The interfaces of the functionality $\mathcal{F}$ are specified in Fig. 4. The KeyGen interface allows a user to establish a link with another user (specifying whether it wants to be the left or the right part of the link). The Setup interface allows a user U_0 to setup a path (starting from U_0) along previously established links. The Lock interface allows a user to create a lock with its right neighbor on a previously created path and the Release algorithm allows a user to release the lock with its left neighbor, in case that the user is either the receiver or its right lock has been released before. Finally, the GetStatus interface allows one to check the current

status of a lock, i.e., whether it is initialized, locked or released. Internally, the locks are modeled by identifiers that are unique across all paths that have been created. Consequently, each lock identifier also identifies the path along which it was established.

B. Discussion

We discuss how the security and privacy properties of interest for PrivMuLs are modeled by the ideal functionality.

Atomicity. Loosely speaking, atomicity means that every user in a path is able to release its left lock in case that his right lock was already released. This is enforced by $\mathcal{F}$ as i) it is keeping track of the chain of locks and their current status in the list $\mathcal{L}$ and ii) the Release interface of $\mathcal{F}$ allows one to release a lock *lid* (changing the flag to Rel) if *lid* is locked and the follow-up lock (*getNextLock*(*lid*)) was already released.

Consistency. A PrivMuL is consistent if no attacker can release his left lock without its right lock being released before. This prevents scenarios where some PrivMuL is released before the receiver is reached and, more generically, the wormhole attack described in Section III. To see why our ideal functionality models this property, observe that the Release interface allows a user to release the left lock only if the right lock has already been released or the user itself is the receiver. In this context, no wormhole attack is possible as intermediate nodes cannot be bypassed.

Relationship Anonymity. Relationship anonymity [13] requires that each intermediate node does not learn any information about the set of users in a PrivMuL beyond its direct neighbors. This property is satisfied by $\mathcal{F}$ as the lock identifiers are sampled at random and during the locking phase a user only learns the identifiers of its left and right lock as well as its left and right neighbor. We further discuss this privacy notion in the full version [5].

V. Constructions

A. Cryptographic Building Blocks

Throughout this work we denote by $\lambda \in \mathbb{N}^+$ the security parameter. Given a set S, we denote by $x \leftarrow\!\!\$\, S$ the sampling of an element uniformly at random from S, and we denote by $x \leftarrow A(in)$ the output of the algorithm A on input *in*. We denote by $\mathsf{min}(a, b)$ the function that takes as input two integers and returns the smaller of the two. In the following we briefly recall the cryptographic building blocks of our schemes.

Homomorphic One-Way Functions. A function $g : \mathcal{D} \to \mathcal{R}$ is one-way if, given a random element $x \in \mathcal{R}$, it is hard to compute a $y \in \mathcal{D}$ such that $g(y) = x$. We say that a function g is homomorphic if $\mathcal{D}$ and $\mathcal{R}$ define two abelian groups and for each pair $(a, b) \in \mathcal{D}^2$ it holds that $g(a \circ b) = g(a) \circ g(b)$, where $\circ$ denotes the group operation. Throughout this work we denote the corresponding arithmetic group additively.

Commitment Scheme. A commitment scheme COM consists of a commitment algorithm $(\mathsf{decom}, \mathsf{com}) \leftarrow \mathsf{Commit}(1^\lambda, m)$ and a verification algorithm $\{0, 1\} \leftarrow \mathsf{V}_{\mathsf{com}}(\mathsf{com}, \mathsf{decom}, m)$. The commitment algorithm allows a prover to commit to a message m without revealing it. In a second phase, the prover can convince a verifier that the message m was indeed committed by showing the unveil information decom. The security of a commitment scheme is captured by the standard ideal functionality $\mathcal{F}_{\mathsf{com}}$ [19].

Non-Interactive Zero-Knowledge. Let R be an NP relation and let L be the set of positive instances, i.e., $L := \{x \mid \exists w \text{ s.t. } R(x, w) = 1\}$. A non-interactive zero-knowledge proof [16] scheme NIZK consists of an efficient prover algorithm $\pi \leftarrow \mathsf{P}_{\mathsf{NIZK}}(w, x)$ and an efficient verifier $\{0, 1\} \leftarrow \mathsf{V}_{\mathsf{NIZK}}(x, \pi)$. A NIZK scheme allows the prover to convince the verifier about the existence of a witness w for a certain statement x without revealing any additional information. The security of a NIZK scheme is modeled by the following ideal functionality $\mathcal{F}_{\mathsf{NIZK}}$: On input $(\mathsf{prove}, \mathit{sid}, x, w)$ by the prover, check if $R(x, w) = 1$ and send $(\mathsf{proof}, \mathit{sid}, x)$ to the verifier if this is the case.

Homomorphic Encryption. One of the building blocks of our work is the additive homomorphic encryption scheme $\mathsf{HE} := (\mathsf{KGen}_{\mathsf{HE}}, \mathsf{Enc}_{\mathsf{HE}}, \mathsf{Dec}_{\mathsf{HE}})$ from Paillier [45]. The scheme supports homomorphic operation over the ciphertexts of the form $\mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, m) \cdot \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, m') = \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, m + m')$. We assume that Paillier's encryption scheme satisfies the notion of ecCPA security, as defined in the work of Lindell [37].

ECDSA Signatures. Let $\mathbb{G}$ be an elliptic curve group of order q with base point G and let $H : \{0, 1\}^* \to \{0, 1\}^{|q|}$ be a collision resistant hash function. The key generation algorithm $\mathsf{KGen}_{\mathsf{ECDSA}}(1^\lambda)$ samples a private key as a random value $x \leftarrow\!\!\$\, \mathbb{Z}_q$ and sets the corresponding public key as $Q := x \cdot G$. To sign a message m, the signing algorithm $\mathsf{Sig}_{\mathsf{ECDSA}}(\mathsf{sk}, m)$ samples some $k \leftarrow\!\!\$\, \mathbb{Z}_q$ and computes $e := H(m)$. Let $(r_x, r_y) := R = k \cdot G$, the algorithm computes $r := r_x \bmod q$ and $s := \frac{e + rx}{k} \bmod q$. The signature consists of (r, s). The verification algorithm $\mathsf{Vf}_{\mathsf{ECDSA}}(\mathsf{pk}, \sigma, m)$ recomputes $e = H(m)$ and returns 1 if and only if $(x, y) = \frac{e}{s} \cdot G + \frac{r}{s} \cdot Q$ and $r = x \bmod q$. It is a well known fact that

for every valid signature (r, s), also the pair $(r, -s)$ is a valid signature. To make the signature *strongly* unforgeable we augment the verification equation with a check that $s \leq \frac{q-1}{2}$. We assume the existence of an interactive protocol $\Pi^{\text{ECDSA}}_{\mathsf{KGen}}$ executed between two users where the one receives (x_0, Q, sk), where sk is a Paillier secret key and $Q = x_0 \cdot x_1 \cdot G$, whereas the other obtains $(x_1, Q, \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, x_0 \cdot x_1))$, where pk is the corresponding Paillier public-key. An efficient protocol that fits these requirements has been recently proposed by Lindell [37].

Anonymous Communication. We assume an anonymous communication channel Π_{anon} available among peers of the network, which is modelled by the ideal functionality $\mathcal{F}_{\mathsf{anon}}$ which anonymously delivers messages to users in the network (e.g., the onion routing functionality as described in [18]).

B. Generic Construction

An interesting question related to PrivMuLs is under which class of hard problems such a primitive exists. A generic construction using trapdoor permutation was given (implicitly) in [40]. Here we propose a scheme from any homomorphic one-way function. Examples of homomorphic one-way functions include discrete logarithm and the learning with error problem [50]. Let $g : \mathcal{D} \to \mathcal{R}$ be a homomorphic one-way function, and let Π_{anon} be an anonymous communication channel. The algorithms of our construction are given in Fig. 5. Note that the key generation algorithm simply returns the identities of the users and therefore it is omitted.

In the setup algorithm, the user U_0 initializes the PrivMuL by sampling n values $(y_0, \ldots, y_{n-1})$ from the domain of g. Then it sends (via Π_{anon}) a triple $(g(\sum_{j=0}^{i-1} y_j), g(\sum_{j=0}^{i} y_j), y_i)$ to each intermediate user. The intermediate user U_i can then check that the triple is well formed using the homomorphic properties of g. Two contiguous users U_i and U_{i+1} can agree on the shared value of $\ell_i := Y_i = g(\sum_{j=0}^{i} y_j)$ by simply comparing the second and first element of their triple, respectively. Note that publishing a valid opening key k such that $g(k) = \ell$ corresponds to inverting the one-way function g. The opening of the locks can be triggered by the last node in the chain U_n: The initial key $k_n := \sum_{i=0}^{n-1} y_i$ consists of a valid pre-image of $\ell_{n-1} := Y_{n-1}$. As soon as the "right" lock is released, each intermediate user U_i has enough information to release its "left" lock. To see this, observe that $g(k_{i+1} - y_i) = g(\sum_{j=0}^{i} y_i - y_i) = g(\sum_{j=0}^{i-1} y_i) = Y_{i-1}$. For the security of the construction, we state the following theorem. Due to space constraints, the proof is deferred to the extended version [5].

Theorem 2. *Let g be a homomorphic one-way function and let Π_{anon} be an anonymous communication channel, then the construction in Fig. 5 UC-realizes the ideal functionality $\mathcal{F}$.*

The generic construction presented here requires a cryptocurrency supporting scripts that define (linearly) homomorphic operations. This construction is therefore of special interest in blockchain technologies such as Ethereum [4] and Hyperledger Fabric [12], where any user can freely deploy a smart contract without restrictions in the cryptographic operations available. We stress that any function with homomorphic properties is suitable to implement our construction. For instance, lattice-based functions (e.g., from the learning with errors problem) can be used for applications where post-quantum cryptography is required. However, many cryptocurrencies, led by Bitcoin, do not support unrestricted scripts and the deployment of generic PrivMuLs requires non-trivial changes (i.e., a hard fork). To overcome this challenge, we turn our attention to scriptless PrivMuLs, where a signature scheme can simultaneously be used for authorization and locking.

C. Scriptless Schnorr-based Construction

The crux of a scriptless locking mechanism is that the lock can consist only of a message m and a public key pk of a given signature scheme and can be released only with a valid signature σ of m under pk. Scritpless locks stem from an idea of Poelstra [47], who proposed a way to embed contracts into Schnorr signatures. In this work we cast Poelstra's informal idea in our framework and we formally characterize its security and privacy guarantees. We further optimize this scheme in order to save one round of communication.

Recall that a public key in a Schnorr signature consists of an element $Q := x \cdot G$ and a signature $\sigma := (k \cdot G, s)$ on a message m is generated by sampling $k \leftarrow_{\$} \mathbb{Z}_q$, computing $e := H(Q \| k \cdot G \| m)$, and setting $s := k - xe$. On a very high level, the locking mechanism consists of an "incomplete" distributed signing of some message m: Two users U_i and U_{i+1} agree on a randomly chosen element $R_0 + R_1$ using a coin tossing protocol, then they set the randomness of the signature to be $R := R_0 + R_1 + Y_i$. Next they jointly compute the value $s := r_0 + r_1 + e \cdot (x_0 + x_1)$ as if Y_i was not part of the randomness, where e is the hash of the transcript so far. The resulting (R, s) is *not* a valid signature on m, since the additive term y^* (where $y^* \cdot G = Y_i$) is

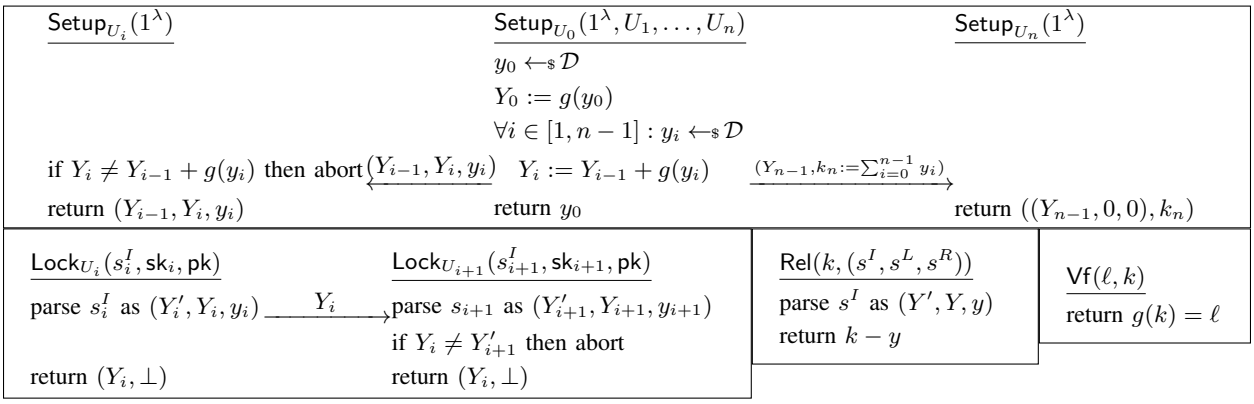

Fig. 5: Algorithms and protocols for the generic construction

missing from the computation of s. However, once the discrete logarithm of Y_i is revealed, a valid signature m can be computed by U_{i+1}. Leveraging this observation, we can enforce an *atomic* opening: The subsequent locking (between U_{i+1} and U_{i+2}) is conditioned on some $Y_{i+1} = Y_i + y_{i+1} \cdot G$. This way, the opening of the right lock reveals the value $y^* + y_{i+1}$ and U_{i+1} can immediately extract y^* and open its left lock with a valid signature on m. We defer the formal description and the analysis of the scheme to the extended version [5].

D. Scriptless ECDSA-based Construction

The Schnorr-based scheme is limited to cryptocurrencies that use Schnorr signatures to authorize transactions and thus is not compatible with those systems, prominently Bitcoin, that implement ECDSA signatures. Therefore, an ECDSA-based scriptless PrivMuL is interesting both from a practical and a theoretical perspective as to whether it can be done at all. Prior to our work, the existence of such a construction was regarded an open question [48]. The core difficulty is that the Schnorr-based construction exploits the linear structure of the signature, whereas the ECDSA signing algorithm completely breaks this linearity feature (e.g., it requires to compute multiplicative shares of a key and inverse of elements within a group). In the following, we show how to overcome these problems, introducing an ECDSA-based construction for PrivMuLs: Locks are of the form (pk, m) and can only be opened with an ECDSA signature σ on m under pk.

Let $\mathbb{G}$ be an elliptic curve group of order q with base point G and let $H : \{0,1\}^* \to \{0,1\}^{|q|}$ be a hash function. The ECDSA-based construction is shown in Fig. 6. Each pair of users (U_i, U_j) generates a shared ECDSA public key $\mathsf{pk} = (x_i \cdot x_j) \cdot G$ via the $\Pi_{\mathsf{KGen}}^{\mathrm{ECDSA}}$ protocol. Additionally, U_i receives his share x_0 and a Paillier secret key, whereas and U_j receives the share x_1 and an encryption c of $x_0 \cdot x_1$. The corresponding key generation protocol is fully described in [37].

The setup of a PrivMuL is very similar to the setup of the generic construction in Fig. 5 except that the one-way function g is now instantiated with discrete logarithm over elliptic curves. Each intermediate user U_i receives a triple (Y_{i-1}, Y_i, y_i) such that $Y_i := Y_{i-1} + y_i \cdot G$, from Π_{anon}. For technical reasons, the initiator of the PrivMuL also includes a proof of wellformedness for each Y_i.

The locking algorithm is initiated by two users U_i and U_{i+1} who agree on a message m (which encodes a unique id) and on a value $Y_i := y^* \cdot G$ of unknown discrete logarithm. The two parties then run a coin tossing protocol to agree on a randomness $R = (r_0 \cdot r_1) \cdot Y_i$. When compared to the Schnorr instance, the crucial technical challenge here is that the randomnesses are composed multiplicatively due to the structure of the ECDSA signature and therefore, the trick applied in the Schnorr construction no longer works here. R is computed through a Diffie-Hellman-like protocol, where the parties exchange $r_0 \cdot Y_i$ and $r_1 \cdot Y_i$ and locally recompute R. As before, the shared ECDSA signature is computed by "ignoring" the term Y_i, since the parties are unaware of its discrete logarithm. The corresponding tuple $\left(r_x, s' := \frac{r_x \cdot (x_i \cdot x_{i+1}) + H(m)}{r_0 \cdot r_1}\right)$ is jointly computed using the encryption of $x_i \cdot x_{i+1}$ and the homomorphic properties of Paillier encryption. This effectively means that $(r_x, s') = (r_x, s^* \cdot y^*)$, where (r_x, s^*) is a valid ECDSA signature on m. In order to check the validity of s', the parties additionally need to exchange the value $R^* := (r_0 \cdot r_1) \cdot G = (y^*)^{-1} \cdot R$. The computation of R^* (together with the corresponding consistency proof) is piggybacked in the coin tossing. Given R^*, the validity

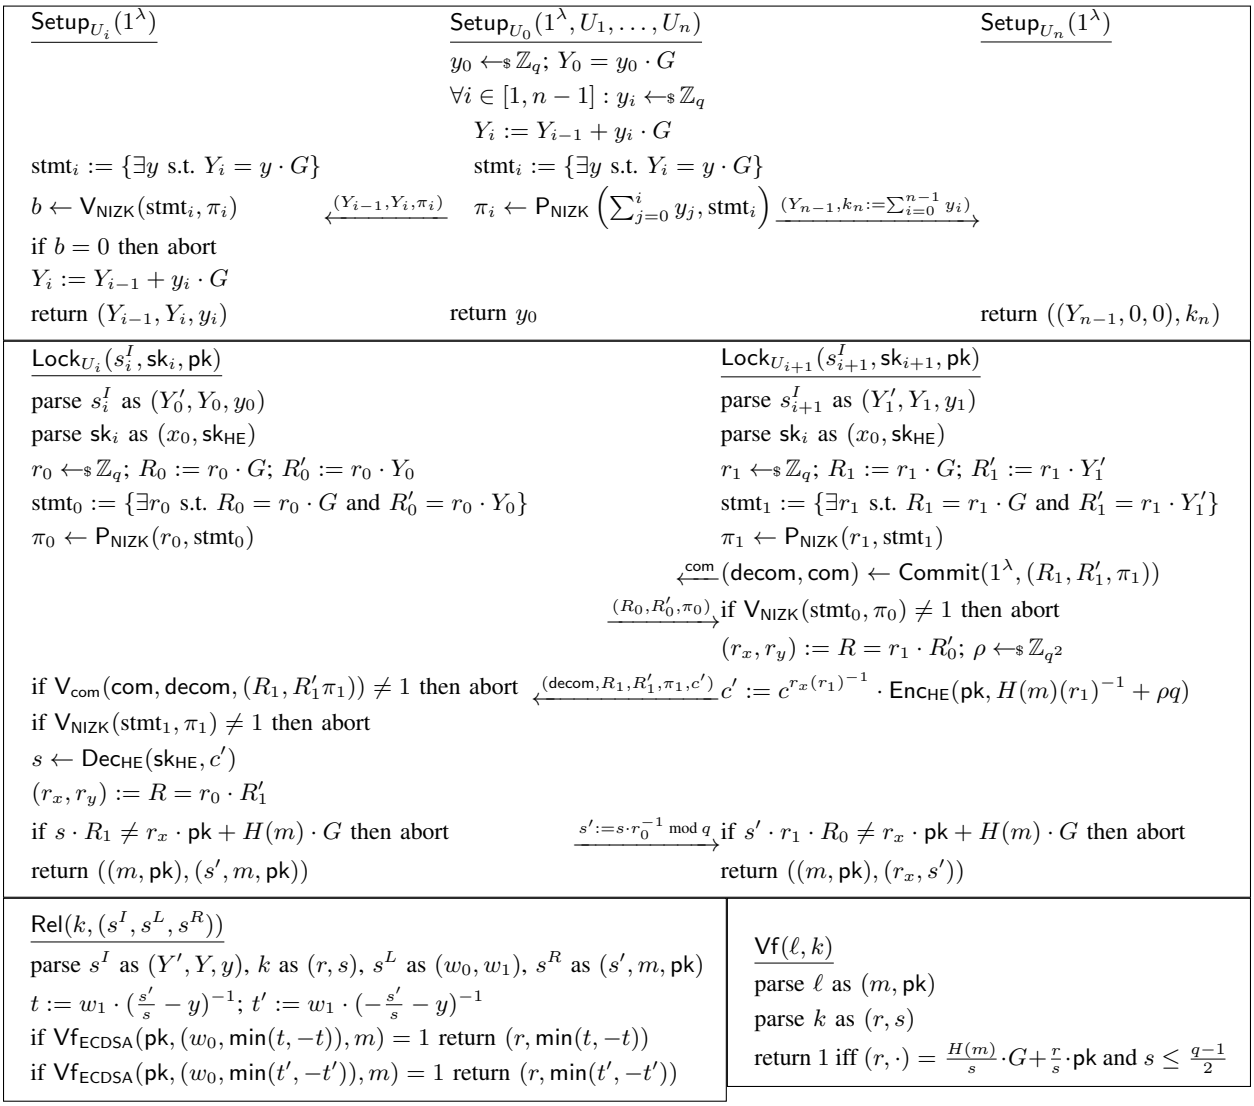

Fig. 6: Algorithms and protocols for the ECDSA-based construction.

of s' can be easily verified by both parties by recomputing it "in the exponent".

From the perspective of U_{i+1}, releasing his left lock without a key for his right lock implies solving the discrete logarithm of Y_i. On the converse, once the right lock is released, the value $y^* + y_{i+1}$ is revealed (where y_{i+1} is part of the state of U_{i+1}) and a valid signature can be computed as $\left(r_x, \frac{s'}{y^*}\right)$. The security of the construction is established by the following theorem (see [5] for a full proof).

Theorem 3. *Let* COM *be a secure commitment scheme, let* NIZK *be a non-interactive zero knowledge proof, let* $\Pi_{\mathsf{KGen}}^{ECDSA}$ *be a secure shared key generation protocol, and let* Π_{anon} *be an anonymous communication channel. If ECDSA signatures are strongly existentially unforgeable and Paillier encryption is* ecCPA *secure, then the construction in Fig. 6 UC-realizes the ideal functionality* $\mathcal{F}$.

E. Hybrid PrivMuLs

We observe that, when instantiated over the same elliptic curve $\mathbb{G}$, the setup protocols of the Schnorr and ECDSA constructions are identical. This means that the initiator of the lock does not need to know whether each intermediate lock is computed using the ECDSA or Schnorr method. This opens the doors to hybrid

PrivMuLs: Given a unified setup, the intermediate pair of users can generate locks using an arbitrary locking protocol. The resulting PrivMuL is a chaining of (potentially) different locks and the release algorithm needs to be adjusted accordingly. For the case of ECDSA-Schnorr the user needs to extract the value y^* from the right Schnorr signature (R^*, s^*) and his state $s^R := s' = s^* - y^* + y_{i+1}$ and $s^I := (Y_i, Y_{i+1}, y_{i+1})$. Given y^*, he can factor it out of its left state $s^L = ((r, s \cdot y^*), m, \mathsf{pk})$ and recover a valid ECDSA signature.

The complementary case (Schnorr-ECDSA) is handled mirroring this algorithm. Similar techniques also apply to the generic construction, when the one-way function is instantiated appropriately (i.e., with discrete logarithm over the same curve). This flexibility enables atomic swaps and cross-currency payments (see Section VII). The security for the hybrid PrivMuLs follows similar to the standard case.

VI. Performance Analysis

A. Implementation Details

We have developed a prototypical Python implementation to demonstrate the feasibility of our construction and evaluated its performance in terms of computation time, computation cost, and communication overhead. We have used the Charm library [3] for the cryptographic operations. We have instantiated ECDSA signatures over the elliptic curve *secp256k1* (the one used in Bitcoin) and we have implemented the homomorphic one-way function with the discrete logarithm function $g(x) := x \cdot G$ over the same curve. Zero-knowledge protocols for discrete logarithms have been implemented using Σ protocols [21] and made non-interactive using the Fiat-Shamir heuristic [27]. For a commitment scheme we have used SHA-256 modeled as a random oracle [14].

B. Evaluation

Testbed. We conducted our experiments on a machine with an Intel Core i7, 3.1 GHz and 8 GB RAM. We consider the following four algorithms: Setup, Lock, Rel, Vf. We do not consider KGen as we use off-the-shelf algorithms without modification. Moreover, the key generation is executed only once upon creating a link and thus does not affect the online performance of PrivMuLs. We refer to [37] for a detailed performance evaluation of the ECDSA key generation. The results of our performance evaluation are shown in Table I.

Computation Time. We measure the computation time required by the users to perform the different algorithms. For the case of two-party algorithms (e.g., Setup and Lock) we consider the time for the two users together. We make two main observations: First, the script-based construction based on discrete logarithm is faster than scriptless PrivMuLs. Second, all the algorithms require computation time of at most one millisecond on a commodity hardware.

		Generic	**Schnorr**	**ECDSA**
Setup	Time (ms)	$0.3 \cdot n$	$1 \cdot n$	$1 \cdot n$
	Comm (bytes)	$96 \cdot n$	$128 \cdot n$	$128 \cdot n$
Lock	Time (ms)	–	2	60
	Comm (bytes)	32	256	416
Rel	Time (ms)	–	0.002	0.02
	Comm (bytes)	0	0	0
Vf	Time (ms)	–	0.6	0.06
	Comm (bytes)	0	0	0
Comp Cost (gas)		$350849 \cdot n$	0	0
Lock size (bytes)		32	$32 + \|m\|$	$32 + \|m\|$
Open size (bytes)		32	64	64

TABLE I: Comparison of the resources required to execute the algorithms for the different PrivMuLs. We denote by n the length of the path. We denote the negligible computation times by – (e.g., single memory read). We denote the size of an application-dependent message by $|m|$ (e.g., a transaction in a payment-channel network).

Communication Overhead. We measure the communication overhead as the amount of information that users need to exchange during the execution of interactive protocols, in particular, Setup and Lock. As expected, the generic construction based on discrete logarithm requires less communication overhead than scriptless constructions. The scriptless construction based on ECDSA requires a higher communication overhead. The higher communication overhead required by the ECDSA approach is mainly due to having the signing key distributed multiplicatively and a more complex structure of the final signature when compared to the Schnorr approach.

Computation Cost. We measure the computation cost in terms of the gas required by a smart contract implementing the corresponding algorithm in Ethereum. Naturally, we consider this cost only for the generic approach based on discrete logarithm. We observe that setting up the corresponding contract requires 350849 unit of gas per hop. At the time of writing, each PrivMuL therefore costs considerably less than 0.01 USD.

Application Overhead. We measure the overhead incurred by the application in terms of the memory required to handle application-dependent data, i.e., information defining the lock and the opening. In tune with the rest of measurements, the generic construction based on discrete logarithms requires the smallest amount of

memory, both for lock and opening information. The different scriptless approaches require the same amount of memory from the application.

Scalability. We study the running time and communication overhead required by each of the roles in a multi-hop lock protocol (i.e., sender, receiver and intermediate user). We consider only the generic approach and the ECDSA construction as representative of the scriptless approach. In the absence of significant metrics from current PCNs, we consider a path length of ten hops is suggested for similar payment networks such as the Ripple credit network [39].

Regarding the computation time, the sender requires 3 ms with the generic approach and 10 ms with the ECDSA scriptless approach. The computation time at intermediate users remain below 1 ms for ECDSA and negligible with the generic approach as they only have to check the consistency of the locks with the predecessor and the successor, independently of the length of the path. Similarly, the computation overhead of the receiver remains below 1 ms as she only checks if a given key is valid to open the lock according to the verify algorithm. In summary, a non-private payment over a path of 5 users takes over 600 ms as reported in [40]. Extending it with the constructions presented in this work provides formal privacy guarantees at virtually no overhead.

Regarding the communication overhead, the sender must send a message of about 960 bytes for the generic approach while about 1280 bytes are required instead if ECDSA scriptless locks are used. Since Sphinx, the anonymous communication network used in the LN, requires padded messages at each node to ensure anonymity, we foresee that every intermediate user must forward a message of the same size.

Comparing these results with other multi-hop and privacy-preserving PCNs available in the literature, we make the following observations. First, the overhead for the constructions presented in this work are in tune with TeeChain [36], where the overhead per hop is about 0.4 ms in a setting where cryptographic operations required for the multi-hop locks have been replaced by a trusted execution environment. Second, our constructions significantly reduce the communication and computation overhead required by multi-hop HTLC [40]: While a payment using multi-hop HTLC requires approximately 5 seconds and 17MB of communication, our approach requires only few milliseconds and less than 1MB.

In summary, the evaluation results show that even with an unoptimized implementation, our constructions offer significant improvements on computation and communication overhead and are ready to be deployed in practice.

VII. Applications

A. Payment-Channel Networks

PrivMuLs can be generically combined with a blockchain B to construct a fully-fledged PCN. Loosely speaking, the transformation works as follows: In the first round the sender sets up the locks running the Setup algorithm, then each pair of intermediate users executes the Lock protocol and establishes the following PrivMuL contract.

PrivMuL (Alice, Bob, ℓ, x, t):
1) If Bob produces the condition k such that $\mathsf{Vf}(\ell, k) = 1$ before t days, Alice pays Bob x coins.
2) If t days elapse, Alice gets back x coins.

Where ℓ is the output lock and x and t are chosen as specified in Section II. Note that we have to assume that B supports the Vf algorithm in its script language. The rest of the payment is unchanged except that the intermediate users execute the Rel algorithm to extract a valid key k to claim the corresponding payment. In the extended version [5], we provide the exact description of the algorithms and we prove the following theorem.

Theorem 4 (Informal)**.** *Let B a secure blockchain and let $\mathbb{L}$ be a secure PrivMuL, then we can construct a secure PCN (as defined in [40]).*

This shows that PrivMuLs are the only cryptographic primitive (except for the blockchain) needed to construct PCNs. The only limitation is that the blockchain needs support the verification of the corresponding contract in their scripting language (see the discussion above). For this reason, the scriptless-construction are preferred for those blockchains where the scripting language does not support the evaluation of a homomorphic one-way function (such as Bitcoin).

Application to the Lightning Network. When applied to the LN, the ECDSA PrivMuL construction conveys several advantages: First, it eliminates the security issues existing in the current LN due to the use of the HTLC contract. Second, it reduces the transaction size as a single signature is required per transaction. This has the benefit of lowering the communication overhead, the transaction fees, and the blockchain memory requirements for closing a payment channel. In fact, we have received initial feedback from the LN community indicating the suitability of our ECDSA-based construction and that initial implementation and testing has begun.

The applicability of our proposals are not restricted to the LN or Bitcoin: There exist other PCNs that could similarly take advantage of the scriptless PrivMuLs presented in this work. For instance, the Raiden Network has been presented as a payment channel network for the scalability issue in Ethereum. The adoption of our ECDSA scriptless PrivMuLs would bring the same benefits to the Raiden Network as it would to the LN.

B. Atomic Swaps

Assume two users U_0 and U_1 holding coins in two different cryptocurrencies that want to exchange them. An *atomic swap* protocol ensures that either the coins are swapped or the balances are untouched, i.e., the exchange must be performed atomically. The widely used protocol for atomic swaps described in [17] leverages the HTLC contract to perform the swap. In a nutshell, an atomic swap can be seen as a multi-hop payment over a path of the form (U_0, U_1, U_0). This approach inherits the security concerns of HTLC contract. Scriptless PrivMuLs also enhance this application domain with formally proven security guarantees.

Additionally, our constructions contribute to the *fungibility* of the coins, a crucial aspect of any currency and therefore also of cryptocurrencies. Current protocols rely on transactions that are clearly distinguishable from regular payments (i.e., one-to-one payments). In particular, atomic swap transactions contain the HTLC contract, in contrast with regular transactions. Scriptless PrivMuLs eliminate this issue since even atomic swaps transactions only require a single signature from a public key, making them indistinguishable from regular payments. Similar arguments also apply for multi-hop payments in PCNs.

C. Interoperable PCNs

In the plethora of cryptocurrencies existing today, an interesting problem consists of performing a multi-hop payment where each link represents a payment channel defined in a different cryptocurrency. In this manner, a user with a payment channel funded in a given cryptocurrency can use it to pay to another user with a payment channel in a different cryptocurrency. Currently, the InterLedger protocol [53] tackles this problem and proposes a mechanism to perform cross-currency multi-hop payments. This protocol relies on the HTLC mechanism, aiming to ensure the atomicity of the payment across different hops.

However, apart from the already discussed issues associated with HTLC, the InterLedger protocol mandates that all cryptocurrencies implement HTLC contracts. This obviously hinders the deployment of this approach. Instead, it is possible to use the different PrivMuL constructions presented in this work on a single path, as described in Section V-E, therefore expanding the domain of cross-currency multi-hop payments.

VIII. Related Work

A recent work [24] shows a protocol to compute an ECDSA signature using multi-party computation. However, it is not as efficient as Lindell's approach [37].

There exists extensive literature proposing constructions for payment channels [22], [23], [35], [49]. These works focus on a single payment channel, and their extension to PCNs remain an open challenge. TumbleBit [31] and Bolt [30] support off-chain payments while achieving payment anonymity guarantees. However, the privacy guarantees of these approaches are restricted to single-hop payments and their extension to support multi-hop payments remains an open challenge.

State channels [25], [33], [43] and state channel networks [26] cannot work with prominent cryptocurrencies except from Ethereum. TeeChain [36] requires the availability of a trusted execution environment at each user. Instead, our proposal can be seamlessly deployed today in virtually all cryptocurrencies, including Ethereum.

The LN has emerged as the most promising approach for PCN in practice. Its current description [8] is being followed by several implementations [6], [10], [11]. However, these implementations suffer from the security and privacy issues with PCNs as described in this work. Instead, we provide several constructions for PrivMuLs that can be leveraged to have secure and privacy-preserving multi-hop payments.

Malavolta et al. [40] propose a protocol for secure and privacy-preserving multi-hop payments compatible with the current LN. Their approach, however, imposes an overhead of around 5 MB for the nodes in the network, therefore hindering its deployability. Here, we propose several efficient constructions that require only a few bytes of communication.

In the recent literature, we can find proposals for secure and privacy-preserving atomic swaps. Tesseract [15] leverages trusted hardware to perform real time cryptocurrency exchanges. The Merkleized Abstract Syntax Trees (MAST) protocol has been proposed as a privacy solution for atomic swaps [34]. However, MAST relies on scripts that are not available in the major cryptocurrencies today. Moreover, specific contracts for atomic swaps hinder the fungibility of the currency: An observer can easily differentiate between a regular payment and a payment resulting from an atomic swap.

IX. Conclusion

We rigorously study the cryptographic core functionality for security, privacy, and interoperability guarantees in PCNs, presenting a new attack on today's PCNs (the wormhole attack) and proposing a novel cryptographic construction (PrivMuLs). We instantiate PrivMuLs in two settings: script-based and scriptless. In the script-based setting, we demonstrate that PrivMuLs can be realized from any (partially) homomorphic operation. In the scriptless setting, we propose a construction based on ECDSA, thereby catering the vast majority of cryptocurrencies deployed today. Our performance evaluation shows that PrivMuLs are practical: All operations take less than 100 milliseconds to run and introduce a communication overhead of less than 500 bytes.

We show that PrivMuLs can be combined in a single path and are of interest in several applications apart from PCNs, such as atomic swaps and interoperable PCNs. In the future, we plan to devise cryptographic instantiations of PCNs for the few cryptocurrencies that are not yet covered, most notably Monero.

Acknowledgements. The authors would like to thank Elizabeth Stark (CEO of Lightning Network Labs) for insightful discussions on the writeup of this paper.

References

[1] "5 potential use cases for bitcoin's lightning network," https://tinyurl.com/y6u4tnda.
[2] "Blockchain explorer information," https://blockchain.info/.
[3] "Charm: A framework for rapidly prototyping cryptosystems," https://github.com/JHUISI/charm.
[4] ""ethereum website"," https://www.ethereum.org/.
[5] "Extended version of this paper," https://sites.google.com/site/multihoplock/home/main.pdf.
[6] "Lightning network daemon," https://github.com/lightningnetwork/lnd.
[7] "Lightning network developers mailing list," url omitted to maintain the anonymity of the authors.
[8] "Lightning network specifications," https://github.com/lightningnetwork/lightning-rfc.
[9] "Raiden network," http://raiden.network/.
[10] "A scala implementation of the lightning network," https://github.com/ACINQ/eclair.
[11] "c-lightning – a lightning network implementation in c," Acceses in May 2018, https://github.com/ElementsProject/lightning.
[12] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. D. Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolic, S. W. Cocco, and J. Yellick, "Hyperledger fabric: A distributed operating system for permissioned blockchains," *CoRR*, vol. abs/1801.10228, 2018.
[13] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi, "Anoa: A framework for analyzing anonymous communication protocols," in *CSF*, 2013, pp. 163–178.
[14] M. Bellare and P. Rogaway, "Random oracles are practical: A paradigm for designing efficient protocols," in *CCS*, 1993.
[15] I. Bentov, Y. Ji, F. Zhang, Y. Li, X. Zhao, L. Breidenbach, P. Daian, and A. Juels, "Tesseract: Real-time cryptocurrency exchange using trusted hardware," in *ePrint Archive*, 2017, p. 1153. [Online]. Available: http://eprint.iacr.org/2017/1153
[16] M. Blum, P. Feldman, and S. Micali, "Non-interactive zero-knowledge and its applications," in *Symposium on Theory of Computing*, 1988, pp. 103–112.
[17] S. Bowe and D. Hopwood, "Hashed time-locked contract transactions," Bitcoin Improvement Proposal, https://github.com/bitcoin/bips/blob/master/bip-0199.mediawiki.
[18] J. Camenisch and A. Lysyanskaya, "A formal treatment of onion routing," in *Annual International Cryptology Conference*, 2005.
[19] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *FOCS*, 2001, pp. 136–.
[20] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. Gün Sirer, D. Song, and R. Wattenhofer, "On Scaling Decentralized Blockchains," in *FC*, 2016, pp. 106–125.
[21] I. Damgård, "On σ-protocols," *Lecture Notes, University of Aarhus, Department for Computer Science*, 2002.
[22] C. Decker, R. Russel, and O. Osuntokun, "eltoo: A simple layer2 protocol for bitcoin," https://blockstream.com/eltoo.pdf.
[23] C. Decker and R. Wattenhofer, "A fast and scalable payment network with bitcoin duplex micropayment channels," in *Stabilization, Safety, and Security of Distributed Systems*, 2015.
[24] J. Doerner, Y. Kondi, E. Lee, and a. shelat, "Secure two-party threshold ecdsa from ecdsa assumptions," in *S&P*, 2018.
[25] S. Dziembowski, L. Eckey, S. Faust, and D. Malinowski, "Perun: Virtual payment hubs over cryptocurrencies," in *ePrint Archive*, 2017. [Online]. Available: https://eprint.iacr.org/2017/635
[26] S. Dziembowski, S. Faust, and K. Hostakova, "Foundations of state channel networks," in *ePrint Archive*, 2018. [Online]. Available: https://eprint.iacr.org/2018/320
[27] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Conference on the Theory and Application of Cryptographic Techniques*, 1986.
[28] C. Fromknecht, "Instantiating scriptless 2p-ecdsa: fungible 2-of-2 multisigs for bitcoin today," https://tokyo2018.scalingbitcoin.org/transcript/tokyo2018/scriptless-ecdsa, 2018, accessed: 2018-10-12.
[29] S. Goldwasser, S. Micali, and R. L. Rivest, "A digital signature scheme secure against adaptive chosen-message attacks," *SIAM Journal on Computing*, vol. 17, no. 2, pp. 281–308, 1988.
[30] M. Green and I. Miers, "Bolt: Anonymous payment channels for decentralized currencies," in *CCS*, 2017.
[31] E. Heilman, L. Alshenibr, F. Baldimtsi, A. Scafuro, and S. Goldberg, "TumbleBit: An untrusted bitcoin-compatible anonymous payment hub," in *NDSS*, 2017.
[32] M. Jakobsson and A. Juels, "Millimix: Mixing in small batches," DIMACS Technical report 99-33, Tech. Rep., 1999.
[33] R. Khalil and A. Gervais, "Revive: Rebalancing off-blockchain payment networks," in *CCS*, 2017, pp. 439–453.
[34] J. Lau, "Merkelized abstract syntax tree," Bitcoin Improvement Proposal, https://tinyurl.com/yc9jh6lv.
[35] J. Lind, I. Eyal, P. R. Pietzuch, and E. G. Sirer, "Teechan: Payment channels using trusted execution environments," 2016, http://arxiv.org/abs/1612.07766.
[36] J. Lind, O. Naor, I. Eyal, F. Kelbert, P. R. Pietzuch, and E. G. Sirer, "Teechain: Reducing storage costs on the blockchain with offline payment channels," in *Systems and Storage Conference*, 2018, p. 125.
[37] Y. Lindell, "Fast Secure Two-Party ECDSA Signing," in *CRYPTO*, 2017, pp. 613–644.
[38] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *CCS*, 2016, pp. 17–30.

[39] G. Malavolta, P. Moreno-Sanchez, A. Kate, and M. Maffei, "SilentWhispers: Enforcing security and privacy in credit networks," in *NDSS*, 2017.
[40] G. Malavolta, P. Moreno-Sanchez, A. Kate, M. Maffei, and S. Ravi, "Concurrency and privacy with payment-channel networks," in *CCS*, 2017.
[41] P. McCorry, M. Möser, S. F. Shahandashti, and F. Hao, "Towards bitcoin payment networks," in *ACISP*, 2016.
[42] S. Micali, K. Ohta, and L. Reyzin, "Accountable-subgroup multisignatures," in *CCS*, 2001, pp. 245–254.
[43] A. Miller, I. Bentov, R. Kumaresan, and P. McCorry, "Sprites: Payment channels that go faster than lightning," *CoRR*, vol. abs/1702.05812, 2017. [Online]. Available: http://arxiv.org/abs/1702.05812
[44] P. Moreno-Sanchez, N. Modi, R. Songhela, A. Kate, and S. Fahmy, "Mind your credit: Assessing the health of the ripple credit network," in *WWW*, 2018, pp. 329–338.
[45] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *International Conference on the Theory and Applications of Cryptographic Techniques*, 1999, pp. 223–238.
[46] A. Pfitzmann and M. Hansen, "A Terminology for Talking about Privacy by Data Minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management," https://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf, Aug. 2010, v0.34.
[47] A. Poelstra, "Lightning in scriptless scripts," Mailing list post, https://lists.launchpad.net/mimblewimble/msg00086.html.
[48] ——, "Scriptless scripts," Presentation slides, https://download.wpsoftware.net/bitcoin/wizardry/mw-slides/2017-05-milan-meetup/slides.pdf.
[49] J. Poon and T. Dryja, "The bitcoin lightning network: Scalable off-chain instant payments," Technical Report, https://lightning.network/lightning-network-paper.pdf.
[50] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *Journal of the ACM*, vol. 56, no. 6, p. 34, 2009.
[51] S. Roos, P. Moreno-Sanchez, A. Kate, and I. Goldberg, "Settling payments fast and private: Efficient decentralized routing for path-based transactions," in *NDSS*, 2018.
[52] C.-P. Schnorr, "Efficient signature generation by smart cards," *Journal of cryptology*, vol. 4, no. 3, pp. 161–174, 1991.
[53] E. S. Stefan Thomas, "A Protocol for Interledger Payments," Whitepaper, https://interledger.org/interledger.pdf.
[54] M. Trillo, "Stress test prepares visanet for the most wonderful time of the year," http://www.visa.com/blogarchives/us/2013/10/10/stress-test-prepares-visanet-for-the-most-wonderful-time-of-the-year/index.html, 2013, accessed: 2017-08-07.
[55] P. Wuille, "Schnorr Bitcoin Improvement Proposal," https://github.com/sipa/bips/blob/bip-schnorr/bip-schnorr.mediawiki.

APPENDIX

A. Wormhole Attack

In this section, we first describe generically the wormhole attack. We then formally prove that no two-round multi-hop payment in a payment-channel network is robust against this attack.

For describing the essence of the wormhole attack, we use an abstract view on payments in PCNs where we assume them to consist of a *commitment phase* and a *releasing phase*. In the commitment phase, pairwise locks (e.g. HTLCs) between the parties along the payment path are created in pairwise locking protocols between the neighboring nodes. In the releasing phase, starting from the payment's receiver, keys for 'opening' the locks are released (as the condition R in the lightning network). These keys should satisfy the property that each path node can – given a valid key for an outgoing lock – derive a key for it's incoming lock. Note that we assume communication to be restricted to nodes connected by a direct link in the PCN. This prevents that besides the specified messages in the releasing phase, keys can be sent to previous nodes in the path (e.g., via broadcast).

Figure 7 shows the payment from Alice to Edward in the abstract setting. Initially, Edwards gives a trapdoor t to Alice. Using this, Alice starts the commitment phase by creating the lock $\ell_{A,B}$ with Bob. To this end, Alice and Bob might use their secret local states s_A and s_B. In the same fashion all following pairwise locks are created in the commitment phase till reaching Edward. Edward then starts the releasing phase by using the trapdoor he initially sent to Alice for creating the key $k_{D,E}$ for opening lock $\ell_{D,E}$. From this key (and the information learned in the commitment phase), Dave can derive key $k_{C,D}$. In this fashion the whole lock chain can be released.

Note that in the setting of only two rounds of communication, the initial secret local states of the users involved in a payment are completely independent from the payment path and consequently from the local states of the other nodes in the path. This is as none of the users received any path-specific information upfront.

As a consequence, two nodes u_i and u_j (with $1 < i+1 < j$) on a payment path can exclude intermediate nodes u_k (with $i < k < j$) from taking part in the releasing phase as follows: After completing the commitment phase in an honest fashion, the releasing phase proceeds honestly till reaching u_j. At this point u_j can derive a key $k_{j-1,j}$ for releasing the lock $\ell_{j-1,j}$ with (honest) user u_{j-1}. Instead of releasing this lock, u_j forwards $k_{j-1,j}$ to u_i which again can use this key for producing a valid key for lock $\ell_{i-1,i}$ with its predecessor

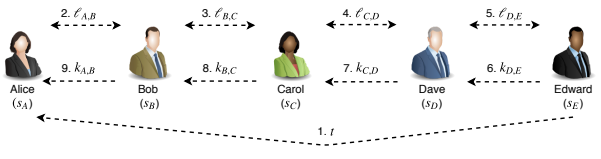

Fig. 7: Illustration of the abstract locking mechanism underlying payments in PCNs

u_{i-1}. This is possible as no secret information from the nodes $\{u_k\}_{i<k<j}$ is required for generating a valid key for $l_{i-1,i}$. Otherwise also opening the final lock would already require secret information from some intermediate nodes. As these pieces of secret information from the intermediate nodes are however completely unrelated to the path and consequently from the trapdoor t even the receiver could not earn the necessary knowledge from the trapdoor for opening the last lock. So finally, node u_i can release lock $\ell_{i-1,i}$ and consecutively all remaining locks can be released without contacting nodes $\{u_k\}_{i<k<j}$ at all. Together with the assumption that nodes $\{u_k\}_{i<k<j}$ cannot receive information through other channels than the direct communication with their immediate neighbors in the path and the fact that keys for locks can only be derived from the initial key, there is no way for nodes $\{u_k\}_{i<k<j}$ to open the locks with their successors in the path.

1) Inevitability of wormhole attacks in one-round payment protocols: In PCNS, payments that only encompass two rounds of communication are inevitably vulnerable to wormhole attacks. [2] More specifically, this means that when no path-specific information was communicated to the intermediate nodes of the payment path before performing the payment, nodes located between two corrupted users in the path can always be bypassed in the releasing phase. This situation occurs in cases where the path is not known upfront, but routing is performed dynamically (e.g., [51]).

We characterize this generic property of payments in PCNS in Theorem 1. In the following, we formally prove this theorem.

Proof. Assume a payment along the path $u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_n$ with u_1 being the sender and u_n being the receiver. Without loss of generality, assume two nodes u_i and u_j with $i < j$ being controlled by the attacker and all other nodes on the path being honest. We show that the view of honest nodes u_l with $l < i$ or $l > j$ in the scenario of u_i and u_j performing a wormhole attack on a successful payment don't differ from the view in the scenario of a successful payment. In addition, we show that the view of honest nodes u_m with $i < m < j$ in the scenario of the wormhole attack do not differ from their view in the scenario of an unsuccessful payment. To this end, we first show how an attacker can simulate the behavior of the nodes $u_{i+1}, \ldots, u_{j-1}$ without changing the view of the honest nodes u_l with $l < i$ or $l > j$.

In the commitment phase, the locks along the path $u_1 \rightarrow u_2 \rightarrow \cdots \rightarrow u_l$ have been created correctly. In the locking protocol between u_l and u_i, u_i behaves honestly and consequently u_l's view is the same as in the honest case. In parallel to starting the locking protocol between u_i and u_{i+1}, the attacker locally simulates the locking protocols for the user's $u_{i+1}, \ldots, u_j$ and creates simulated locks $\ell_{i+1}, \ldots, \ell_j$. This is possible as by sampling random local states for those nodes, the attacker can run the locking protocol locally. Finally, u_j can continue the commitment phase in an honest manner using as own local state the one resulting from the simulated commitments. This cannot be distinguished by node u_{j+1} as it's own local state is unrelated to the local states of the other intermediate nodes.

As we consider the case of a successful payment, the releasing phase will be performed honestly by nodes $u_n, \ldots u_{j+1}$. When u_{j+1} releases the lock between u_j and u_{j+1} with key k_j, then the attacker can simulate releasing the locks $\ell_{j-1}, \ldots, \ell_i$ locally without publishing the corresponding keys. This is possible as the attacker can use the local states of the intermediate nodes $u_{i+1}, \ldots, u_{j-1}$ from the simulated commitment phase for deriving keys $k_{j-1}, \ldots, k_{i-1}$. As k_{i-1} is hence also a valid key for the honestly created lock l_{i-1}, the releasing phase can from this point be concluded in an honest manner.

Finally, we can observe that nodes $u_{i+1}, \ldots, u_{j-1}$ are not contacted at all in the releasing phase of the payment which is the same as in the case that the payment was unsuccessful, i.e., the releasing phase was not initiated by the receiver at all. □

B. PrivMuLs Correctness

In this section, we define the notion of correctness for PrivMuLs.

Definition 3 (Correctness of PrivMuLs)**.** *Let $\mathbb{L}$ be a PrivMuL, $\lambda \in \mathbb{N}^+$ and $n \in \mathsf{poly}(\lambda)$. Let $(U_0, \ldots, U_n) \in \mathbb{U}^n$ be a vector of users, $(\mathsf{sk}_0, \ldots, \mathsf{sk}_{n-1})$ and $(\mathsf{sk}^*_1, \ldots, \mathsf{sk}^*_n)$ two vectors of private keys and $(\mathsf{pk}_0, \ldots, \mathsf{pk}_{n-1})$ a vector of shared public keys such that for all $0 \leq i < n$, it holds that*

$$\{(\mathsf{sk}_i, \mathsf{pk}_i), (\mathsf{sk}^*_{i+1}, \mathsf{pk}_i)\} \leftarrow \langle \mathsf{KGen}_{U_i}(1^\lambda), \mathsf{KGen}_{U_{i+1}}(1^\lambda) \rangle.$$

Let $(s^I_0, \ldots, s^I_n)$ be vector of initial states and k_n be a key such that for all $0 \leq i < n$

$$\{s^I_0, \ldots, (s^I_n, k_n)\} \leftarrow \left\langle \begin{array}{c} \mathsf{Setup}_{U_0}(1^\lambda, U_1, \ldots, U_n) \\ \cdots \\ \mathsf{Setup}_{U_n}(1^\lambda) \end{array} \right\rangle$$

[2] Note that we assume here PCNs of the previously described structure hence requiring that payments encompass a commitment and a revealing phase and communication to be restricted to direct neighbors.

Furthermore, let $(\ell_0, \ldots, \ell_{n-1})$ *be a vector of locks,* $(s_1^L, \ldots, s_n^L)$ *and* $(s_0^R, \ldots, s_{n-1}^R)$ *vectors of states, and* $(k_0, \ldots, k_{n-1})$ *a vector of keys such that for all* $0 \leq i < n$*, it holds that*

$$\{(\ell_i, s_i^R), (\ell_i, s_{i+1}^L)\} \leftarrow \left\langle \begin{array}{c} \mathsf{Lock}_{U_i}(s_i^I, \mathsf{sk}_i, \mathsf{pk}_i) \\ \mathsf{Lock}_{U_{i+1}}(s_{i+1}^I, \mathsf{sk}_{i+1}^*, \mathsf{pk}_i) \end{array} \right\rangle$$

and

$$k_i \leftarrow \mathsf{Rel}(k_{i+1}, (s_{i+1}^I, s_{i+1}^L, s_{i+1}^R))$$

where s_n^R *is* $\perp$*. We say that* $\mathbb{L}$ *is correct if there exists a negligible function* negl *such that for all* $0 \leq i < n$ *it holds that*

$$\Pr\left[\mathsf{Vf}(\ell_i, k_i) = 1\right] \geq 1 - \mathsf{negl}(\lambda).$$

C. Schnorr-based Scriptless Construction

In the following we cast the idea of Poelstra [47] in our framework.

Schnorr Signatures. Let $\mathbb{G}$ be an elliptic curve group of order q with base point G and let $H : \{0,1\}^* \to \{0,1\}^{|q|}$ be a collision resistant hash function (modeled as a random oracle). The key generation algorithm $\mathsf{KGen}_{\mathsf{schnorr}}(1^\lambda)$ of a Schnorr signature [52] samples some $x \leftarrow_\$ \mathbb{Z}_q$ and sets the corresponding public key as $Q := x \cdot G$. To sign a message m, the signing algorithm $\mathsf{Sig}_{\mathsf{schnorr}}(\mathsf{sk}, m)$ samples some $k \leftarrow_\$ \mathbb{Z}_q$, computes $e := H(Q \| k \cdot G \| m)$, sets $s := k - xe$, and returns $\sigma := (R, s)$, where $R := k \cdot G$. The verification $\mathsf{Vf}_{\mathsf{schnorr}}(\mathsf{pk}, \sigma, m)$ returns 1 if and only if $s \cdot G = R + H(Q \| R \| m) \cdot Q$. Schnorr signatures are known to be strongly unforgeable against the discrete logarithm assumption [29]. We assume the existence of a 2-party protocol $\Pi_{\mathsf{KGen}}^{\mathsf{schnorr}}$ where the two players, on input x_0 and x_1, set a shared public key $Q := (x_0 + x_1) \cdot G$. Such a protocol can be realized using standard techniques [42].

Description. Let $\mathbb{G}$ be an elliptic curve group of order q with base point G and let $H : \{0,1\}^* \to \{0,1\}^{|q|}$ be a hash function. The Schnorr-based construction is formally described in Fig. 8. The key generation algorithm consists of an execution of the $\Pi_{\mathsf{KGen}}^{\mathsf{schnorr}}$ protocol. At the end of a successful run, U_i receives (x_i, pk) whereas U_j obtains (x_j, pk), where $\mathsf{pk} := (x_i + x_j) \cdot G$. The setup of a PrivMuL is identical to the ECDSA-based construction and can be found in Fig. 6.

Prior to the locking phase, two users U_i and U_{i+1} (implicitly) agree on the value Y_i and on a message m to be signed. Each message is assumed to be unique for each session (e.g., contains a transaction identifier). The locking algorithm consists of an "incomplete" distributed signing of m. First, the two parties agree on a randomly chosen element $R_0 + R_1$ using a standard coin tossing protocol, then they set the randomness of the signature to be $R := R_0 + R_1 + Y_i$. Note that at this point the parties cannot complete the signature since they do not know the discrete logarithm of Y_i. Instead, they jointly compute the value $s := r_0 + r_1 + e \cdot (x_0 + x_1)$ as if Y_i was not part of the randomness, where e is the hash of the transcript so far. The resulting (R, s) is *not* a valid signature on m, since the additive term y^* (where $y^* \cdot G = Y_i$) is missing from the computation of s. However, rearranging the terms, we have that $(R, s + y^*)$ is a valid signature on m. This implies that, once the discrete logarithm of Y_i is revealed, a valid signature m can be computed by U_{i+1}. Leveraging this observation, U_{i+1} can enforce an *atomic* opening: The subsequent locking (between U_{i+1} and U_{i+2}) is conditioned on some $Y_{i+1} = Y_i + y_{i+1} \cdot G$. This way, the opening of the right lock reveals the value $y^* + y_{i+1}$ and U_{i+1} can immediately extract y^* and open its left lock with a valid signature on m. The security of the construction is shown by the following theorem. We refer the reader to Appendix E for a full proof.

Theorem 5. *Let* COM *be a secure commitment scheme, let* NIZK *be a non-interactive zero knowledge proof, let* $\Pi_{\mathsf{KGen}}^{\mathit{schnorr}}$ *be a secure shared key generation protocol, and let* Π_{anon} *be an anonymous communication channel. If Schnorr signatures are strongly existentially unforgeable, then the construction in Fig. 8 UC-realizes the ideal functionality* $\mathcal{F}$*.*

D. Comparison of Privacy Notions and Guarantees

In this section we discuss our notion of relationship anonymity as the privacy notion of interest for PCNs and compare it with other possible privacy notions described in the literature related to PCN.

Our privacy model faithfully captures the reality of currently deployed PCN. In particular, Malavolta et al. [40] showed that it captures the well established notion of relationship anonymity. In a nutshell, relationship anonymity [46] requires that, given two simultaneous successful payment operations between $\text{sender}_{\{0,1\}}$ and $\text{receiver}_{\{0,1\}}$ that share the same path with at least one honest intermediate user, corrupted intermediate users cannot determine the correct pair (sender_b, receiver_b) for a given payment with probability better than $1/2$ (i.e., guessing). Note that this holds only for payments for the same value, since such an information it is trivially leaked to intermediate users, i.e., each users can monitor how adjacent links evolve and infer the amount that was transferred.

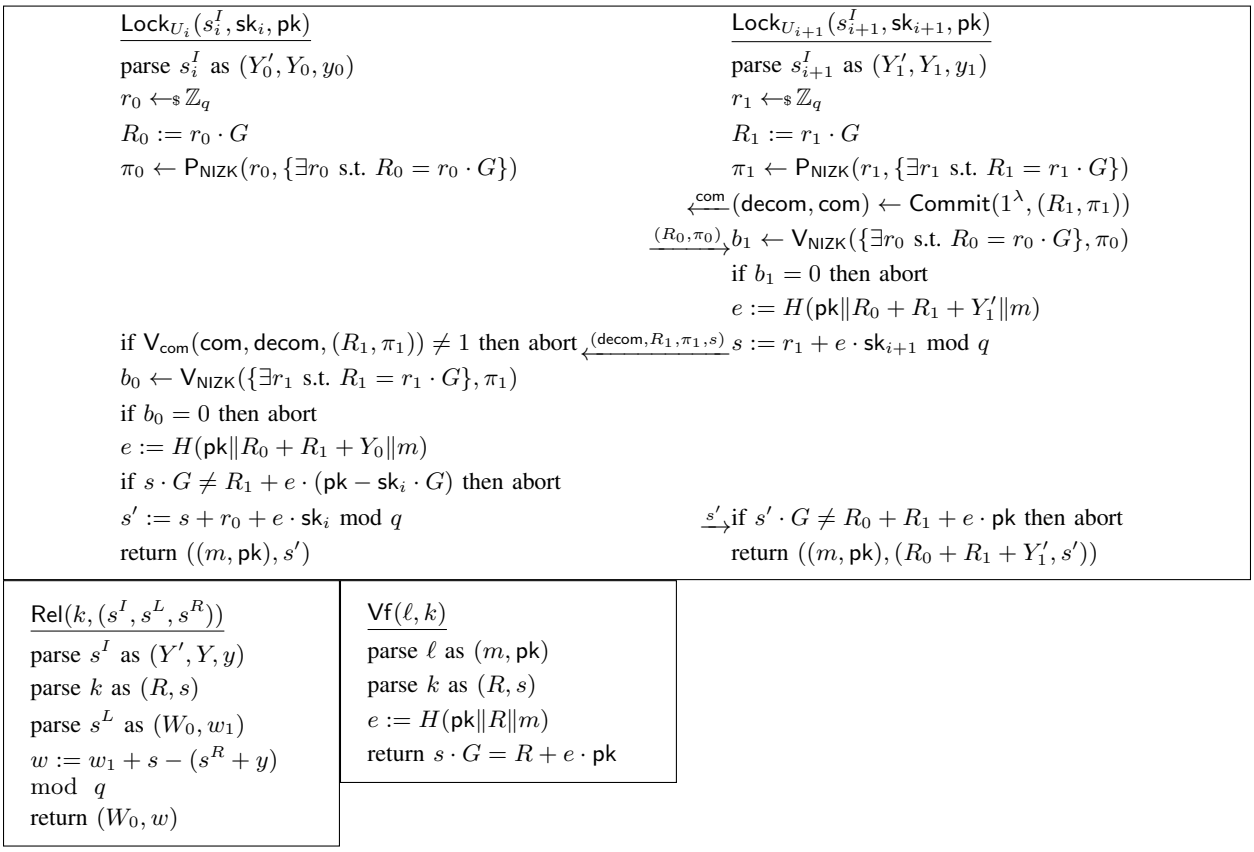

Fig. 8: Algorithms and protocols for the Schnorr-based construction. The Setup protocol is as defined in Fig. 6.

An alternative privacy notion is described in BOLT [30]. There, authors propose *payment anonymity*. Intuitively, payment anonymity requires that the merchant, even in collaboration with a set of malicious customers, learns nothing about a customer's spending pattern beyond the information available outside the payment protocol.

While this privacy notion additionally hides the value that is transacted, it is restricted to single-hop payments and does not consider the crucial aspect of conditional payment required when more than one intermediate user takes part in the payment. As discussed in Section III, many well-established networks use paths with multiple intermediaries and it is reasonable to expect long paths also in the LN. To obtain the best of both worlds, one could envision a protocol where private one-hop payments are performed "at the edges" (i.e., between sender and first hop as well as between last hop and the receiver) while the rest of intermediate users carry out a multi-hop payment á la LN.

However, this approach raises several questions. First, it is unclear whether the hypothetical resulting privacy guarantees are stronger or weaker than those presented in this work. It is possible that the naïve combination of the two systems would completely break down the guarantees of both schemes. Techniques presented in both works might be required to develop a new system. Second, BOLT requires a blockchain supporting a reach scripting language and it is therefore not compatible with prominent cryptocurrencies (such as Bitcoin). Thus, making this system Bitcoin-compatible would require fundamentally new techniques.

In summary, although it seems to be an interesting research direction, further work is required to study this approach and its privacy properties.

E. Security Analysis

Throughout the analysis we denote by $\mathsf{poly}(\lambda)$ any function that is bounded by a polynomial in λ. We denote any function that is negligible in the security parameter by $\mathsf{negl}(\lambda)$. We say that an algorithm is PPT if it is modelled as a probabilistic Turing machine whose

running time is bounded by some function $\mathsf{poly}(\lambda)$. In the following we elaborate on the security analysis of our constructions.

We shall point out that in this analysis we model a very strong variant of anonymous communication, which might not always be reasonable to assume. More realistic privacy guarantees are captured by onion routing [18] functionalities or mix networks [32]. For ease of exposition we stick to our simplistic model, noting that our proof is completely parametric and one can switch to a less idealized functionality in a modular manner.

1) Generic Construction: Here we elaborate the proof of Theorem 2.

Proof. We define the following sequence of hybrids, where we gradually modify the initial experiment.

$\mathcal{H}_0$: Is identical to the protocol as described in Section V-B.

$\mathcal{H}_1$: Instead of sending messages through the Π_{anon} channel, the parties communicate in interaction with the ideal functionality $\mathcal{F}_{\mathsf{anon}}$.

Anon(m, U_i)

Upon invocation U_j on input (m, U_i):

send (m, U_i) to U_i

$\mathcal{H}_2$: Consider the following ensemble of variables in the interaction with $\mathcal{A}$: A honest user U_i, a key pair $(\mathsf{sk}_i, \mathsf{pk})$, a state s^I, a tuple $(\ell_i, \ell_{i+1}, s^L, s^R)$ such that

$$\{\cdot, (\ell_i, s^L)\} \leftarrow \langle \cdot, \mathsf{Lock}_{U_i}(s^I, \mathsf{sk}_i, \mathsf{pk})\rangle$$

and

$$\{(\ell_{i+1}, s^R), \cdot\} \leftarrow \langle \mathsf{Lock}_{U_i}(s^I, \mathsf{sk}_i, \mathsf{pk}), \cdot\rangle.$$

If, for any set of these variables, the adversary returns some k such that $\mathsf{Vf}(\ell_{i+1}, k) = 1$ and $\mathsf{Vf}(\ell_i, \mathsf{Rel}(k, (s^I, s^L, s^R))) \neq 1$, then the experiment aborts.

$\mathcal{H}_3$: Consider the following ensemble of variables in the interaction with $\mathcal{A}$: A pair of honest users (U_0, U_i) a set of (possibly corrupted) users $(U_1, \ldots, U_n)$, a key pair $(\mathsf{sk}_i, \mathsf{pk})$, a set of initial states

$$(s_0^I \ldots, s_n^I) \leftarrow \left\langle \begin{array}{l} \mathsf{Setup}_{U_0}(1^\lambda, U_1, \ldots, U_n), \ldots, \\ \mathsf{Setup}_{U_n}(1^\lambda) \end{array} \right\rangle,$$

and a pair of locks (ℓ_{i-1}, ℓ_i) such that

$$\{\cdot, (\ell_{i-1}, \cdot)\} \leftarrow \langle \cdot, \mathsf{Lock}_{U_i}(s_i^I, \mathsf{sk}_i, \mathsf{pk})\rangle$$

and

$$\{(\ell_i, \cdot), \cdot\} \leftarrow \langle \mathsf{Lock}_{U_i}(s_i^I, \mathsf{sk}_i, \mathsf{pk}), \cdot\rangle.$$

If, for any set of these variables, the adversary returns some k_{i-1} such that $\mathsf{Vf}(\ell_{i-1}, k_{i-1}) = 1$ before the user U_i outputs a key k_i such that $\mathsf{Vf}(\ell_i, k_i) = 1$, then the experiment aborts.

$\mathcal{H}_4$: Let $S = (U_0, \ldots, U_m)$ be an ordered set of (possibly corrupted) users. We say that that an ordered subset $A = (U_1, \ldots, U_j)$ is *adversarial* if U_i is honest and $(U_{i+1}, \ldots, U_j)$ are corrupted. Note that every set of users can be expressed as a concatenation of adversarial subsets, that is $S = (A_1 || \ldots || A_{m'})$, for some $m' \leq m$. Whenever a honest user is requested to set up a lock for a certain set $S = (A_1 || \ldots || A_{m'})$, it initializes an independent lock for each subset (A_i, A_{i+1}^0), where A_{i+1}^0 is the first element of the $(i+1)$-th set, if present. Whenever some A_{i+1}^0 is requested to release the key for the corresponding lock (recall that all A_{i+1}^0 are honest nodes) it releases the key for the fresh lock (A_i, A_{i+1}^0) instead.

$\mathcal{S}$: The interaction of the simulator is identical to $\mathcal{H}_4$ except that the actions of $\mathcal{S}$ are dictated by the interaction with $\mathcal{F}$. The simulator reads the communication of $\mathcal{A}$ with the honest users via $\mathcal{F}_{\mathsf{anon}}$ and is queried by $\mathcal{F}$ on the following set of inputs.

1) $(\cdot, \cdot, \cdot, \cdot, \mathsf{Init})$: The simulator reconstruct the adversarial set (defined above) from the ids and sets up a fresh lock chain.
2) $(\cdot, \mathsf{Lock})$: The simulator initiates the locking procedure with the adversary and replies with $\perp$ if the execution is not successful.
3) $(\cdot, \mathsf{Rel})$ The simulator releases the key of the corresponding lock and publishes it.

If $\mathcal{A}$ interacts with a honest user (e.g., by releasing a lock) the simulator queries the corresponding interface of $\mathcal{F}$.

Note that the simulator is efficient and interacts as the adversary with the ideal world. Furthermore, the simulation is always consistent with the ideal world, i.e., if the adversary's action is not supported by the interfaces of $\mathcal{F}$ the simulation aborts. What is left to be shown is that the neighboring hybrids are indistinguishable to the eyes of the environment $\mathcal{E}$.

Lemma 1. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_0, \mathcal{A}, \mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_1, \mathcal{A}, \mathcal{E}}.$$

Proof. Follows directly from the security of Π_{anon}. □

Lemma 2. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_1, \mathcal{A}, \mathcal{E}} \equiv \mathsf{EXEC}_{\mathcal{H}_2, \mathcal{A}, \mathcal{E}}.$$

Proof. Follows from the homomorphic property of the function g: Recall that a key-lock pair (k, ℓ) is valid if and only if $g(k) = \ell$. Let (k_i, ℓ_i) be the output of $\mathcal{A}$, by construction we have that $\ell_i = \ell_{i-1} + g(y_i)$, for some (ℓ_{i-1}, y_i), which is part of the state of the honest node. Since the release algorithm computes $k_i - y_i$ we have that

$$\begin{aligned} g(k_i - y_i) &= g(k_i) - g(y_i) \\ &= \ell_i - g(y_i) \\ &= \ell_{i-1} + g(y_i) - g(y_i) \\ &= \ell_{i-1} \end{aligned}$$

with probability 1, by the homomorphic property of g. □

Lemma 3. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_2,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}}.$$

Proof. Let $q \in \mathsf{poly}(\lambda)$ be a bound on the number of interactions. Recall that $\mathcal{H}_2$ and $\mathcal{H}_3$ differ only for the case where the adversary outputs a key for a honestly generated lock before the trapdoor is released. Assuming towards contradiction that the probability that this event happens is non-negligible, we can construct the following reduction against the one-wayness of g: On input some $Y^* \in \mathcal{R}$, the reduction guesses a session $j \in [1, q]$ and some index $i \in [1, n]$. The setup algorithm of the j-th session is modified as follows: Y_i is set to be Y^*. Then, for all $\iota \in [i-1, 0]$, the setup samples some $y_\iota \in \mathcal{D}$ and returns $(Y_\iota = Y_{\iota+1} - g(y_\iota), Y_{\iota+1}, y_\iota)$. The setup samples a random $y_i \in \mathcal{D}$ and sets $Y_{i+1} = g(y_i)$. Then, for $\iota \in [i+1, n-1]$, the setup samples $y_\iota \in \mathcal{D}$ returns $(Y_\iota, Y_\iota + g(y_\iota), y_\iota)$. The nodes $(U_1, \ldots, U_{n-1})$ are given the corresponding output (except for U_i) and U_n is given $(Y_{n-1}, \sum_{j=i}^{n-1} y_j)$. If the node U_i is requested to release the lock, the reduction aborts. At some point of the execution the adversary $\mathcal{A}$ outputs some $y*$, and the reduction returns $y^* + y_{i-1}$.

The reduction is clearly efficient and, whenever j and i are guessed correctly, the reduction does not abort. Since the group defined by g is abelian, the distribution induced by the modified setup algorithm is identical to the original (except for the initial state of U_1). Also note that, whenever j and i are guessed correctly, the user U_i is honest and therefore the adversary does not not see the corresponding internal state. It follows that the reduction is identical to $\mathcal{H}_2$, to the eyes of the adversary. Finally, whenever the adversary outputs some valid k_{i-1} for ℓ_{i-1}, then it holds that $g(k_{i-1}) = \ell_{i-1}$. Substituting we have that

$$\begin{aligned} g(k_{i-1}) &= \ell_{i-1} \\ g(y^*) &= Y_{i-1} \\ g(y^*) &= Y^* - g(y_{i-1}) \\ g(y^*) + g(y_{i-1}) &= Y^* \\ g(y^* + y_{i-1}) &= Y^*. \end{aligned}$$

It follows that the reduction is successful with probability at least $\frac{1}{q \cdot n \cdot \mathsf{poly}(\lambda)}$. This proves our statement. □

Lemma 4. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}} \equiv \mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}}.$$

Proof. Recall that adversarial sets are always interleaved by a honest node. Therefore in $\mathcal{H}_3$ for each adversarial set starting at index i there exists a y such that $Y_i = Y_{i-1} + g(y)$ and $\mathcal{A}$ is not given y. Since y is randomly sampled from $\mathcal{D}$ we have that $Y_{i-1} + g(y) \equiv Y'$, for some Y' sampled uniformly from $\mathcal{R}$, which corresponds to the view of $\mathcal{A}$ in $\mathcal{H}_4$. □

Lemma 5. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}} \equiv \mathsf{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{E}}.$$

Proof. The changes between the two experiments are only conceptual and the equivalence of the views follows. □

This concludes our analysis. □

2) Schnorr-based Construction: Here we prove Theorem 5.

Proof. We define the following sequence of hybrids, where we gradually modify the initial experiment.

$\mathcal{H}_0$: Is identical to the protocol as described in Appendix C.

$\mathcal{H}_1$: Instead of sending messages through the Π_{anon} channel, the parties communicate in interaction with the ideal functionality $\mathcal{F}_{\mathsf{anon}}$.

Anon(m, U_i)

Upon invocation U_j on input (m, U_i):

send (m, U_i) to U_i

$\mathcal{H}_2$: All the calls to the commitment scheme are replaced with interactions with the ideal functionality $\mathcal{F}_{\mathsf{com}}$, defined in the following.

Commit(*sid*, m)

Upon invocation by U_i (for $i \in \{0, 1\}$):

record (sid, i, m) and send (com, sid) to U_{1-i}

if some $(sid, \cdot, \cdot)$ is already stored ignore the message

Decommit(*sid*)

Upon invocation by U_i (for $i \in \{0,1\}$):
if (sid, i, m) is recorded then send (decom, sid, m) to U_{1-i}

Instead of calling the Commit algorithm on some message m, the parties sent a message of the form **Commit(sid, m)** to the ideal functionality, and the decommitment algorithm is replaced with a call to **Decommit(sid)**. The verifying party simply records messages from $\mathcal{F}_{\mathsf{com}}$.

$\mathcal{H}_3$: All the calls to the NIZK scheme are replaced with interactions with the ideal functionality $\mathcal{F}_{\mathsf{NIZK}}$:

Prove(*sid*, x, w)

Upon invocation by U_i (for $i \in \{0,1\}$):
if $R(x,w) = 1$ then send (proof, sid, x) to U_{1-i}

Instead of running the proving algorithm in input (x, w), the proving party queries the functionality on **Prove(sid, x, w)**. The verifier records the messages from $\mathcal{F}_{\mathsf{NIZK}}$.

$\mathcal{H}_4, \mathcal{H}_5, \mathcal{H}_6, \mathcal{S}$: The subsequent hybrids are defined as $\mathcal{H}_2, \mathcal{H}_3, \mathcal{H}_4, \mathcal{S}$, respectively, in Theorem 2.

As argued before, the simulator is efficient and the interaction is consistent with the inputs of the ideal functionality. In the following we prove the indistinguishability of the neighboring experiments.

Lemma 6. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_0,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_1,\mathcal{A},\mathcal{E}}.$$

Proof. Follows directly from the security of Π_{anon}. □

Lemma 7. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_1,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_2,\mathcal{A},\mathcal{E}}.$$

Proof. Follows directly from the security of the commitments scheme COM. □

Lemma 8. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_2,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}}.$$

Proof. Follows directly from the security of the non-interactive zero-knowledge scheme NIZK. □

Lemma 9. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}}.$$

Proof. In order to show this claim, we introduce an intermediate experiment.

$\mathcal{H}_3^*$: The key generation and locking algorithms are substituted with the interaction with the functionality $\mathcal{F}_{\mathsf{schnorr}}$, which provides any two users with the interfaces specified below. Note that the signing interface is called by both parties on input m and $y = \sum_{j=0}^{i} y_j$, where i is the position of the lock in the chains and the y_j are defined as in the original protocol.

KeyGen($\mathbb{G}$, G, q)

Upon invocation by both U_0 and U_1 on input ($\mathbb{G}$, G, q):
sample $x \leftarrow\!\!\$\, \mathbb{Z}_q$ and compute $Q = x \cdot G$
set $\mathsf{sk}_{U_0,U_1} = x$
sample x_0 and x_1 randomly
sample a hash function $H : \{0,1\}^* \to \{0,1\}^{|q|}$
send (x_0, Q) to U_0 and (x_1, Q) to U_1
ignore future calls by (U_0, U_1)

Sign(m, y)

Upon invocation by both U_0 and U_1 on input (m, y):
compute $(R, s) = \mathsf{Sig}_{\mathsf{schnorr}}(\mathsf{sk}_{U_0,U_1}, m)$
return $(R, s - y)$

We defer the indistinguishability proof to lemma 10. Let cheat by the event that triggers an abort of the experiment in $\mathcal{H}_4$, that is, the adversary returns some k such that $\mathsf{Vf}(\ell_{i+1}, k) = 1$ and that $\mathsf{Vf}(\ell_i, \mathsf{Rel}(k, (s^I, s^L, s^R))) \neq 1$. Assume towards contradiction that $\Pr[\mathsf{cheat} \mid \mathcal{H}_3^*] \geq \frac{1}{\mathsf{poly}(\lambda)}$, then we can construct the following reduction against the strong-existential unforgeability of Schnorr signatures: The reduction receives as input a public key pk and samples an index $j \in [1, q]$, where $q \in \mathsf{poly}(\lambda)$ is a bound on the total amount of interactions. Let Q be the key generated in the j-th interaction, the reduction sets $Q = \mathsf{pk}$. All the calls to the signing algorithm are redirected to the signing oracle. If the event cheat happens, the reduction returns corresponding $(k^*, \ell^*) = (\sigma^*, (m^*, \mathsf{pk}^*))$, otherwise it aborts.

The reduction is clearly efficient. Assume for the moment that j is the index of the interaction where cheat happens, and let $i+1$ be the index that identifies the lock ℓ^* in the corresponding chain. Note that in case the guess of the reduction is correct we have that $\mathsf{pk}^* = \mathsf{pk}$. Since cheat happens we have that $\mathsf{Vf}_{\mathsf{schnorr}}(\mathsf{pk}^*, m^*, \sigma^*) = 1$ and the release fails, i.e., $\mathsf{Vf}(\ell_i, \mathsf{Rel}(k, (s_i^I, s_i^L, s_i^R)) \neq 1$ (where ℓ_i is the lock in the previous position as ℓ^* in the same chain). Recall that the release algorithm parses s_i^L as $(W_{i,0}, w_{i,1})$ and σ^* as (R^*, s^*) and returns $(W_{i,0}, w_{i,1} + s^* - (s_i^R + y_i))$. Substituting with the corresponding values

$$\begin{aligned}
&\left(W_{i,0}, w_{i,1} + s^* - (s_i^R + y_i)\right) \\
&= \left(R_i, \left(s_i - \sum_{j=0}^{i-1} y_j\right) + s^* - \left(s_j - \sum_{j=0}^{i} y_j\right) - y_i\right) \\
&= \left(R_i, s_i + s^* - s_j\right),
\end{aligned}$$

where s_j is the answer of the oracle on the j-th session on input m_j. This implies that $s^* \neq s_j$, otherwise (R_i, s_i) would be a valid signature since it is an output of the signing oracle. Since each message uniquely identifies a session (the same message is never queried twice to the interface **Sign**(m,y)) this implies that $(\sigma^*, (m^*, \mathsf{pk}^*))$ is a valid forgery. By assumption this happens with probability at least $\frac{1}{q \cdot \mathsf{poly}(\lambda)}$, which is a contradiction and proves that $\mathsf{Pr}\left[\mathsf{cheat} \mid \mathcal{H}_3^*\right] \leq \mathsf{negl}(\lambda)$. Since the experiments $\mathcal{H}_3$ and $\mathcal{H}_4$ differ only when cheat happens (and $\mathcal{H}_4$ aborts), we are only left with showing the indistinguishability of $\mathcal{H}_3$ and $\mathcal{H}_3^*$.

Lemma 10. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_3^*,\mathcal{A},\mathcal{E}}.$$

Proof. The proof consists of the description of the simulator for the interactive lock algorithm. The simulator for the key generation phase is trivial and therefore it is omitted. We describe two simulators depending on whether the honest adversary is playing the role of the "left" or "right" party. For each proof, both the simulators implicitly check that the given witness is valid and abort if this is not the case.

1) Left corrupted: Prior to the interaction the simulator is sent $(Y, y, (\mathsf{prove}, \{\exists y^* \text{ s.t } y^* \cdot G = Y\}, y^*))$, which is the state corresponding to the execution of the lock. After agreeing on a message m, the simulator sends $(\mathsf{com}, \mathit{sid})$ to $\mathcal{A}$, for a random *sid*. The simulator also queries the interface **Sign** on input m, y^* and receives a signature $\sigma = (R, s)$. At some point of the execution $\mathcal{A}$ sends $(R_0, (\mathsf{prove}, \{\exists r_0 \text{ s.t } r_0 \cdot G = R_0\}, r_0))$. The simulator replies with
$$\left(\begin{array}{c} R^* = R - (R_0 + Y), \\ \mathsf{decom}, \mathit{sid}, \left(\begin{array}{l} \mathsf{proof}, \mathit{sid}, \\ \{\exists r^* \text{ s.t } r^* \cdot G = R^*\} \end{array}\right), \\ R^*, (s - r_0 - e \cdot x_0) \end{array}\right)$$
where $e = H(\mathsf{pk} \| R^* \| m)$ and x_0 is the value returned by the key generation to $\mathcal{A}$. The rest of the execution is unchanged.
2) Right corrupted: Prior to the interaction the simulator is sent $(Y, y, (\mathsf{prove}, \{\exists y^* \text{ s.t } y^* \cdot G = Y\}, y^*))$, which is the state corresponding to the execution of the lock. After agreeing on a message m, the simulator is given
$$\left(\mathsf{com}, \mathit{sid}, \left(\begin{array}{l} R_1, \ \mathsf{prove}, \mathit{sid}, \\ \{\exists r_1 \text{ s.t } r_1 \cdot G = R_1\}, r_1 \end{array}\right)\right)$$
by $\mathcal{A}$. The simulator then queries the interface **Sign** on input m, y^* and receives a signature $\sigma = (R, s)$. The simulator sends $(R^* = R - (R_1 + Y), (\mathsf{proof}, \mathit{sid}, \{\exists r^* \text{ s.t } r^* \cdot G = R^*\}))$ to $\mathcal{A}$ and receives $((\mathsf{decom}, \mathit{sid}), s^*)$ in response. The simulator checks whether $s^* = r_1 + e \cdot x_1$, where $e = H(\mathsf{pk} \| R^* \| m)$, and returns s if this is the case.

Both simulators are obviously efficient and the distributions induced by the simulated views are identical to the ones of the original protocol. □

This concludes the proof of lemma 9. □

Lemma 11. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_5,\mathcal{A},\mathcal{E}}.$$

Proof. Let $q \in \mathsf{poly}(\lambda)$ be a bound on the number of interactions. Let cheat denote the events that triggers an abort in $\mathcal{H}_5$ but not in $\mathcal{H}_4$. In the following we are going to show that $\mathsf{Pr}\left[\mathsf{cheat} \mid \mathcal{H}_4\right] \leq \mathsf{negl}(\lambda)$, thus proving the indistinguishability of $\mathcal{H}_4$ and $\mathcal{H}_5$. Assume that the converse is true, then we can construct the following reduction against the discrete logarithm problem (which is implied by the sEUF of Schnorr): On input some $Y^* \in \mathbb{G}$, the reduction guesses a session $j \in [1, q]$ and some index $i \in [1, n]$. The setup algorithm of the j-th session is modified as follows: Y_i is set to be Y^*. Then, for all $\iota \in [i-1, 0]$, the setup samples some $y_\iota \in \mathbb{Z}_q$ and returns $(Y_\iota = Y_{\iota+1} - y_\iota \cdot G, Y_{\iota+1}, y_\iota)$. The setup samples a random $y_i \in \mathbb{Z}_q$ and sets $Y_{i+1} = y_i \cdot G$. Then, for $\iota \in [i+1, n-1]$, the setup samples $y_\iota \in \mathbb{Z}_q$ returns $(Y_\iota, Y_\iota + y_\iota \cdot G, y_\iota)$. The nodes $(U_1, \ldots, U_{n-1})$ are given the corresponding output (except for U_i) and U_n is given $(Y_{n-1}, \sum_{j=i}^{n-1} y_j)$. If the node U_i is requested to release the lock, the reduction aborts. At some point of the execution the adversary $\mathcal{A}$ outputs some $k^* = (R^*, s^*)$. The reduction parses s^R as the updated state of U_i and returns $s^* + y_{i-1} - s^R$.

The reduction is clearly efficient and, whenever j and i are guessed correctly, the reduction does not abort. Since the group $\mathbb{G}$ is abelian and the U_i is honest, the distribution induced by the modified setup algorithm is identical to the original to the eyes of the adversary. Recall that cheat happens only in the case where k^* is a valid opening for ℓ_i and the release algorithm is successful on input k^* (if the last condition is not satisfied both $\mathcal{H}_4$ and $\mathcal{H}_5$ abort). Substituting, we have that s^R is of the form $r_0 + r_1 + e \cdot (x0 + x_1) - y = s' - y$, for some $y \in \mathbb{Z}_q$. Since the release is successful, then it must be the case that $(R' = (r_0 + r_1) \cdot G + Y_{i-1}, s')$ is a valid Schnorr signature on the message m_{i-1} (agreed by the two parties in the locking algorithm for ℓ_{i-1}), which implies that $y \cdot G = Y_{i-1}$. As argued in the

proof of lemma 9, if $s^* \neq s'$, then we have an attacker against the strong unforgeability of the signature scheme. It follows that $s^* = s'$ with all but negligible probability. Substituting we have

$$\begin{aligned}(s^* + y_{i-1} - s^R) \cdot G &= (s^* + y_{i-1} - s' + y) \cdot G \\ &= (y_{i-1} + y) \cdot G \\ &= y_{i-1} \cdot G + y \cdot G \\ &= y_{i-1} \cdot G + Y_{i-1} \\ &= y_{i-1} \cdot G + (Y^* - y_{i-1} \cdot G) \\ &= Y^*\end{aligned}$$

as expected. Since, by assumption, this happens with probability at least $\frac{1}{q \cdot n \cdot \mathsf{poly}(\lambda)}$ we have a successful attacker against the discrete logarithm problem. This proves our statement.

□

Lemma 12. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_5,\mathcal{A},\mathcal{E}} \equiv \mathsf{EXEC}_{\mathcal{H}_6,\mathcal{A},\mathcal{E}}.$$

Proof. Recall that adversarial sets are always interleaved by a honest node. Therefore in $\mathcal{H}_5$ for each adversarial set starting at index i there exists a y such that $Y_i = Y_{i-1} + y \cdot G$ and $\mathcal{A}$ is not given y. Since y is randomly sampled from $\mathbb{Z}_q$ we have that $Y + i - 1 + y \cdot G \equiv Y'$, for some Y' sampled uniformly from $\mathbb{G}$, which corresponds to the view of $\mathcal{A}$ in $\mathcal{H}_6$. □

Lemma 13. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_6,\mathcal{A},\mathcal{E}} \equiv \mathsf{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{E}}.$$

Proof. The change is only syntactical and the indistinguishability follows. □

This concludes our analysis. □

3) ECDSA-based Construction: In the following we prove Theorem 3.

Proof. The sequence of hybrids that we define is identical to the one described in the proof of Theorem 5. In the following we prove the indistinguishability of neighboring experiments only for the cases where the argument needs to be modified. If the argument is identical, the proof is omitted.

Lemma 14. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_3,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}}.$$

Proof. In order to show this claim, we introduce an intermediate experiment.

$\mathcal{H}_3^*$: The key generation and locking algorithms are substituted with the interaction with the functionality $\mathcal{F}_{\mathsf{ECDSA}}$, which provides any pair of users with the interfaces specified below. Note that the locking algorithm is called by both parties on input m and $y = \sum_{j=0}^{i} y_j$, where i is the position of the lock in the chains and the y_j are defined as in the original protocol.

KeyGen($\mathbb{G}$, G, q)

Upon invocation by both U_0 and U_1 on input ($\mathbb{G}$, G, q):
sample $x \leftarrow\!\!\$\, \mathbb{Z}_q$ and compute $Q = x \cdot G$
sample x_0 and x_1 randomly
sample a hash function $H : \{0,1\}^* \to \{0,1\}^{|q|}$
sample a key pair $(\mathsf{sk}_{U_0,U_1}, \mathsf{pk}_{U_0,U_1}) \leftarrow \mathsf{KGen}_{\mathsf{HE}}(1^\lambda)$
compute $c \leftarrow \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, \tilde{r})$ for a random $\tilde{r}$
send (x_0, Q, H, sk) to U_0 and (x_1, Q, H, c) to U_1
ignore future calls by (U_0, U_1)

Sign(m, y)

Upon invocation by both U_0 and U_1 on input (m, y):
compute $(r, s) = \mathsf{Sig}_{\mathsf{ECDSA}}(\mathsf{sk}_{U_0,U_1}, m)$
return $(r, \mathsf{min}(s \cdot y, -s \cdot y))$

The indistinguishability proof of $\mathcal{H}_3$ and $\mathcal{H}_3^*$ is formally shown in lemma 15. Let cheat by the event that triggers an abort of the experiment in $\mathcal{H}_4$, that is, the adversary returns some k such that $\mathsf{Vf}(\ell_{i+1}, k) = 1$ and $\mathsf{Vf}(\ell_i, \mathsf{Rel}(k, (s^I, s^L, s^R))) \neq 1$. Assume towards contradiction that $\Pr[\mathsf{cheat} \mid \mathcal{H}_3^*] \geq \frac{1}{\mathsf{poly}(\lambda)}$, then we can construct the following reduction against the strong-existential unforgeability of ECDSA signatures: The reduction receives as input a public key pk and samples an index $j \in [1, q]$, where $q \in \mathsf{poly}(\lambda)$ is a bound on the total amount of interactions. Let Q be the key generated in the j-th interaction, the reduction sets $Q = \mathsf{pk}$. All the calls to the signing algorithm are redirected to the signing oracle. If the event cheat happens, the reduction returns corresponding $(k^*, \ell^*) = (\sigma^*, (m^*, \mathsf{pk}^*))$, otherwise it aborts.

The reduction runs in polynomial time. Assume for the moment that j is the index of the interaction where cheat happens, and let $i+1$ be the index that identifies the lock ℓ^* in the corresponding chain. Note that in case the guess of the reduction is correct we have that $\mathsf{pk}^* = \mathsf{pk}$. Since cheat happens we have that $\mathsf{Vf}_{\mathsf{ECDSA}}(\mathsf{pk}^*, m^*, \sigma^*) = 1$ and the release fails, i.e., $\mathsf{Vf}(\ell_i, \mathsf{Rel}(k^*, k^*, (s_i^I, s_i^L, s_i^R))) \neq 1$ (where ℓ_i is the lock in the previous position as ℓ^* in the same chain). Recall that the release algorithm parses s_i^L as $(w_{i,0}, w_{i,1})$, σ^* as (r^*, s^*), and s_i^R as (s', m, pk) and computes $t = w_1 \cdot (\frac{s'}{s^*} - y)^{-1}$ and $t' = w_1 \cdot (-\frac{s'}{s^*} - y)^{-1}$. Then it returns either $(w_{i,0}, \mathsf{min}(t, -t))$ or $(w_{i,0}, \mathsf{min}(t', -t'))$

depending on which verifies as a valid signature on m under pk. Substituting with the corresponding values (for the case t is the lower term)

$$(w_{i,0}, t) = \left(r_i, w_{i,1} \cdot \left(\frac{s'}{s^*} - y \right)^{-1} \right)$$
$$= \left(r_i, s_i \cdot \sum_{j=0}^{i-1} y_j \cdot \left(\frac{s_j \cdot \sum_{j=0}^{i} y_j}{s^*} - y_i \right)^{-1} \right)$$

where s_j is the answer of the oracle on the j-th session on input the corresponding message m_j. If we set $s^* = s_j$ then we have

$$(w_{i,0}, t) = \left(r_i, s_i \cdot \sum_{j=0}^{i-1} y_j \cdot \left(\sum_{j=0}^{i} y_j - y_i \right)^{-1} \right)$$
$$= (r_i, s_i)$$

which is a valid signature on m_i (since it is the output of the signing oracle) and the release would be successful. So this cannot happen and we can assume that $s^* \neq s_j$. A similar argument (substituting t with t') can be used to show that it must be the case that $s^* \neq -s_j$. Since each message uniquely identifies a session (the same message is never queried twice to the interface **Sign**(m,y)) this implies that $(\sigma^*, (m^*, \mathsf{pk}^*))$ is a valid forgery. By assumption this happens with probability at least $\frac{1}{q \cdot \mathsf{poly}(\lambda)}$, which is a contradiction and proves that $\Pr[\mathsf{cheat} \mid \mathcal{H}_3^*] \leq \mathsf{negl}(\lambda)$. Since the experiments $\mathcal{H}_3$ and $\mathcal{H}_4$ differ only when cheat happens (and $\mathcal{H}_4$ aborts), we are only left with showing the indistinguishability of $\mathcal{H}_3$ and $\mathcal{H}_3^*$.

Lemma 15. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_3, \mathcal{A}, \mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_3^*, \mathcal{A}, \mathcal{E}}.$$

Proof. The proof consists of the description of the simulator for the interactive lock algorithm. The simulator for the key generation phase is identical as the one described in the work of Lindell [37]. In the following we describe the two simulators for the locking protocol depending on whether the honest adversary is playing the role of the "left" or "right" party. For each zero-knowledge proof, both the simulators implicitly check that the given witness is valid and abort if this is not the case.

1) Left corrupted: Prior to the interaction the simulator is sent $(Y, y, (\mathsf{prove}, \{\exists y^* \text{ s.t } y^* \cdot G = Y\}, y^*))$, which is the state corresponding to the execution of the lock. After agreeing on a message m, the simulator sends $(\mathsf{com}, \mathit{sid})$ to $\mathcal{A}$, for a random *sid*. The simulator also queries the interface **Sign** on input m, y^* and receives a signature $\sigma = (r, s)$. The simulator sets $R = \frac{H(m)}{s} \cdot G + \frac{r}{s} \cdot \mathsf{pk}$. At some point of the execution $\mathcal{A}$ sends $(R_0, R_0', (\mathsf{prove}, \{\exists r_0 \text{ s.t } r_0 \cdot G = R_0 \text{ and } r_0 \cdot Y = R_0'\}, r_0))$. Then the simulator samples a $\rho \leftarrow \mathbb{Z}_{q^2}$ and computes $c' \leftarrow \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, s \cdot r_0 + \rho q)$. Then it provides the attacker with

$$\left(\begin{array}{ll} & R^* = (r_0)^{-1} \cdot R, R_1 = y^{-1} \cdot R^*, \\ \mathsf{decom}, & \mathsf{proof}, \mathit{sid}, \\ \mathit{sid}, & \left\{ \begin{array}{l} \exists r^* \text{ s.t } r^* \cdot G = R_1 \\ \quad \text{and } r^* \cdot Y = R^* \end{array} \right\} \\ R_1, R^*, c' & \end{array} \right),$$

and the rest of the execution is unchanged.

The executions are identical except for the way c' is computed. In order to show the statistical proximity we invoke a the following helping lemma.

Lemma 16. *[37] For all $(r, s, p) \in \mathbb{Z}_q$ and for a random $\rho \in \mathbb{Z}_{q^2}$, the distributions $\mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, r \cdot s \bmod q + pq + \rho q)$ and $\mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, r \cdot s \bmod q + \rho q)$ are statistically close.*

In the real world c' is computed as $\mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, r \cdot s \bmod q + pq + \rho q)$, for some p which is bounded by q since the only operation performed without modular reduction are one multiplication and one addition, which cannot increase the result by more than q^2. Since the distribution $\mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, r \cdot s \bmod q + \rho q)$ is identical to the simulation, the indistinguishability follows.

2) Right corrupted: Prior to the interaction the simulator is sent $(Y, y, (\mathsf{prove}, \{\exists y^* \text{ s.t } y^* \cdot G = Y\}, y^*))$, which is the state corresponding to the execution of the lock. After agreeing on a message m, the simulator is given

$$\left(\begin{array}{l} \mathsf{com}, \mathit{sid}, R_1, R_1', \\ \mathsf{prove}, \mathit{sid}, \\ \left\{ \begin{array}{l} \exists r_1 \text{ s.t } r_1 \cdot G = R_1 \text{ and} \\ \quad r_1 \cdot Y = R_1' \end{array} \right\}, \\ r_1 \end{array} \right)$$

by $\mathcal{A}$. The simulator then queries the interface **Sign** on input m, y^* and receives a signature $\sigma = (r, s)$. Then it sets $R = \frac{H(m)}{s} \cdot G + \frac{r}{s} \cdot \mathsf{pk}$ and $R^* = R - (R_1 + Y)$ and sends $(R_0 = y^{-1} \cdot R^*, R^*, (\mathsf{proof}, \mathit{sid}, \{\exists r^* \text{ s.t } r^* \cdot G = R_0 \text{ and } r^* \cdot Y = R^*\}))$ to $\mathcal{A}$. The attacker sends $((\mathsf{decom}, \mathit{sid}), c')$ in response. The simulator checks

$$\mathsf{Dec}_{\mathsf{HE}}(\mathsf{sk}, c') = \tilde{r} \cdot r \cdot (r_1)^{-1} + H(m) \cdot r_1^{-1} \bmod q,$$

where $\tilde{r}$ was sampled in the key generation algorithm. If the check holds true, the simulator sends s to $\mathcal{A}$.

The distributions induced by the simulator is identical to the real experiment except for the way c is computed. Towards showing indistinguishability, consider the following modified simulator, that is given the oracle $\mathcal{O}(c', a, b)$ as defined in the following security experiment of the Paillier encryption scheme.

$\mathsf{Exp} - \mathsf{ecCPA}^{\mathcal{A}}_{\mathsf{HE}}(\lambda)$:

$(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}_{\mathsf{HE}}(1^\lambda)$

$(w_0, w_1) \leftarrow_\$ \mathbb{Z}_q$

$Q = w_0 \cdot G$

$b \leftarrow_\$ \{0, 1\}$

$c \leftarrow \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, w_b)$

$b' \leftarrow \mathcal{A}(\mathsf{pk}, c, Q)^{\mathcal{O}(\cdot,\cdot,\cdot)}$

where $\mathcal{O}(c', a, b)$ returns 1 iff $\mathsf{Dec}_{\mathsf{HE}}(\mathsf{sk}, c') = a + b \cdot w_b$

return 1 iff $b = b'$

Instead of performing the last check, the simulator queries the oracle on input $(c', a = H(m) \cdot r_1^{-1}, b = r \cdot (r_1)^{-1})$. It is clear that the modified simulator accepts if and only if the simulator described above accepts. Assume towards contradiction that the modified simulator can be efficiently distinguished from the real world experiment. Then we can reduce to the security of Paillier as follows: On input (pk, c, Q), the reduction simulates the inputs of $\mathcal{A}$ as described in the modified simulator using the input pk, Q, and c as the corresponding variables. It is easy to see that the reduction is efficient. Note that if $b = 0$ then we have that $c = \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, w_0)$ and $Q = w_0 \cdot G$, which is identical to the real world execution. On the other hand if $b = 1$ then it holds that $c = \mathsf{Enc}_{\mathsf{HE}}(\mathsf{pk}, w_1)$ and $Q = w_0 \cdot G$, where w_1 is uniformly distributed in $\mathbb{Z} - q$, which is identical to the (modified) simulated experiment. This implies that the modified simulation is computationally indistinguishable from the real world experiment. Since the modified simulation and the simulation (as described above) are identical to the eyes of the adversary, the validity of the lemma follows. □

This concludes the proof of lemma 14. □

Lemma 17. *For all PPT distinguishers $\mathcal{E}$ it holds that*

$$\mathsf{EXEC}_{\mathcal{H}_4,\mathcal{A},\mathcal{E}} \approx \mathsf{EXEC}_{\mathcal{H}_5,\mathcal{A},\mathcal{E}}.$$

Proof. Let $q \in \mathsf{poly}(\lambda)$ be a bound on the number of interactions. Let cheat denote the event that triggers an abort in $\mathcal{H}_5$ but not in $\mathcal{H}_4$. In the following we are going to show that $\Pr[\mathsf{cheat} \mid \mathcal{H}_4] \leq \mathsf{negl}(\lambda)$, thus proving the indistinguishability of $\mathcal{H}_4$ and $\mathcal{H}_5$. Assume that the converse is true, then we can construct the following reduction against the discrete logarithm problem (which is implied by the sEUF of ECDSA): On input some $Y^* \in \mathbb{G}$, the reduction guesses a session $j \in [1, q]$ and some index $i \in [1, n]$. The setup algorithm of the j-th session is modified as follows: Y_i is set to be Y^*. Then, for all $\iota \in [i-1, 0]$, the setup samples some $y_\iota \in \mathbb{Z}_q$ and returns $(Y_\iota = Y_{\iota+1} - (y_\iota) \cdot G, Y_{\iota+1}, y_\iota)$. The setup samples a random $y_i \in \mathbb{Z}_q$ and sets $Y_{i+1} = y_i \cdot G$. Then, for $\iota \in [i+1, n-1]$, the setup samples $y_\iota \in \mathbb{Z}_q$ and returns $(Y_\iota, Y_\iota + y_\iota \cdot G, y_\iota)$. The nodes $(U_1, \ldots, U_{n-1})$ are given the corresponding output (except for U_i) and U_n is given $(Y_{n-1}, \sum_{j=i}^{n-1} y_j)$. If the node U_i is requested to release the lock, the reduction aborts. At some point of the execution the adversary $\mathcal{A}$ outputs some $k^* = (r^*, s^*)$. The reduction parses $s^R = (s', m, \mathsf{pk})$ as the updated state of U_i then checks the following:

1) $\left(\frac{s}{s^*} + y_{i-1}\right) \cdot G = Y^*$

2) $-\left(\frac{s'}{s^*} + y_{i-1}\right) \cdot G = Y^*$

and returns the LHS term of the equation that satisfies the relation.

The reduction is clearly efficient and, whenever j and i are guessed correctly, the reduction does not abort. Since the $\mathbb{G}$ is abelian and the U_i is honest, the distribution induced by the modified setup algorithm is identical to the original to the eyes of the adversary. Recall that cheat happens only in the case where k^* is a valid opening for ℓ_i and the release algorithm is successful on input k^* (if the last condition is not satisfied both $\mathcal{H}_4$ and $\mathcal{H}_5$ abort). Substituting, we have that s' is of the form $\frac{x_0 \cdot x_1 \cdot r_x + H(m)}{r_0 \cdot r_1} = \tilde{s} \cdot y$, where $R' = r_0 \cdot r_1 \cdot Y_{i-1} = (r_x, r_y)$, for some $y \in \mathbb{Z}_q$. Since the release is successful, then it must be the case that $(r_x, \tilde{s})$ is a valid ECDSA signature on the message m_{i-1} (agreed by the two parties in the locking algorithm for ℓ_{i-1}). This implies that $y \cdot G = Y_{i-1}$. As argued in the proof of lemma 14, if $s^* \neq \tilde{s}$ and $s^* \neq \tilde{s}$, then we have an attacker against the strong unforgeability of the signature scheme. It follows that $s^* = \tilde{s}$ or $s^* = -\tilde{s}$

with all but negligible probability. Substituting we have

$$\begin{aligned}\left(\frac{s'}{s^*} + y_{i-1}\right) \cdot G &= \left(\frac{\tilde{s} \cdot y}{s^*} + y_{i-1}\right) \cdot G \\ &= \frac{\tilde{s} \cdot y}{s^*} \cdot G + y_{i-1} \cdot G \\ &= y \cdot G + y_{i-1} \cdot G \\ &= Y_{i-1} + y_{i-1} \cdot G \\ &= (Y^* - y_{i-1} \cdot G) + y_{i-1} \cdot G \\ &= Y^*\end{aligned}$$

which implies that condition (1) holds if $s^* = \tilde{s}$. For the other case

$$\begin{aligned}-\left(\frac{s'}{s^*} + y_{i-1}\right) \cdot G &= -\left(\frac{\tilde{s} \cdot y}{s^*} + y_{i-1}\right) \cdot G \\ &= -\frac{\tilde{s} \cdot y}{s^*} \cdot G + y_{i-1} \cdot G \\ &= y \cdot G + y_{i-1} \cdot G \\ &= Y_{i-1} + y_{i-1} \cdot G \\ &= (Y^* - y_{i-1} \cdot G) + y_{i-1} \cdot G \\ &= Y^*\end{aligned}$$

which means that condition (2) is satisfied if $s^* = -\tilde{s}$. Since, by assumption, this happens with probability at least $\frac{1}{q \cdot n \cdot \mathsf{poly}(\lambda)}$ we have a successful attacker against the discrete logarithm problem. This proves our statement. □

This concludes our proof. □

F. PCNs from Multi-Hop Locks

In this section we show that PrivMuLs are sufficient to construct a full-fledged PCN that satisfy the standard security definition from Malavolta et al. [40].

1) Ideal Functionalities: We assume an ideal realization of PrivMuLs in the form of an ideal functionality $\mathcal{F}_{\mathbb{L}}$ as described in Fig. 4. That is, all parties have oracle access to $\mathcal{F}_{\mathbb{L}}$ through the specified interfaces.

Furthermore, we assume the existence of a blockchain B that we model as a trusted append-only bulletin board: The corresponding ideal functionality $\mathcal{F}_{\mathsf{B}}$ maintains B locally and updates it according to the transactions between users. At any point in the execution, anyone can send a distinguished message read to $\mathcal{F}_{\mathsf{B}}$, who sends the whole transcript of B to U. We denote the number of entries of B by $|\mathsf{B}|$. We assume that users can specify arbitrary *contracts*, i.e., transactions in B may be associated with arbitrary conditions which require to be me in order to make the transaction effective. $\mathcal{F}_{\mathsf{B}}$ is entrusted to enforce that a contract is fulfilled before the corresponding transaction is executed.

We model time as the number of entries of the blockchain B, i.e., time t is whenever $|\mathsf{B}| = t$. Note that we can artificially elapse time by adding dummy entries to B and that the current time is available to all parties by simply reading B and counting the number of entries.

2) System Assumptions: We assume that every user in the PCN is aware of the complete network topology, that is, the set of all users and the existence of a payment channel between every pair of users. We further assume that the sender of a payment chooses a payment path to the receiver according to her own criteria.

The current value on each payment channel is not published but instead kept locally by the users sharing a payment channel. The two users U_0 and U_1 are assumed to maintain locally the capacity of their channel, denoted by $\mathsf{cap}(U_0, U_1)$. We further assume that every user is aware of the payment fees charged by each other user in the PCN. For ease of exposition we define the predicate $\mathsf{fee}(U_i)$ to return the fee charged by the user U_i. We assume that pairs of users sharing a payment channel communicate through secure and authenticated channels (such as TLS), which is easy to implement given that every user is uniquely identified by a public key.

3) Our System: In the following we describe the three operations (open channel, close channel, and payment) that constitute the core of our system. For the sake of simplicity we restrict each pair of user to at most one channel, however our construction can be easily extended to support multiple channels per pair.

OPEN CHANNEL. The open channel protocol generates a new payment channel between users U_1 and U_2. The user U_1 invokes $\mathcal{F}_{\mathbb{L}}$ on input (U_2, L), depending on the direction of the channel, which returns the users identifiers (U_1, U_2) if the operation was successful. Then the users create an initial blockchain deposit that includes the following information: Their addresses, the initial capacity of the channel, the channel timeout, and the fee charged to use the channel agreed beforehand between both users. After the deposit has been successfully added to the blockchain, the operation returns 1. If any of the previous steps is not carried out as defined, the operation returns 0.

CLOSE CHANNEL. The close channel protocol is run by two users U_1 and U_2 sharing an open payment channel to close it at the state defined by v and accordingly update their bitcoin balances in the Bitcoin blockchain. From this point on, U_1 and U_2 ignore all the requests from $\mathcal{F}_{\mathbb{L}}$ relative to their link.

contract(Alice, Bob, *lid*, x, t)
1) If GetStatus(*lid*) = Rel before t days, then Alice pays Bob x coins.
2) If t elapse, then Alice gets back x coins.

PAYMENT. A payment operation transfers a value v from a sender (U_0) to a receiver (U_{n+1}) through a path of open payment channels between them ($U_0, \ldots, U_{n+1}$). The sender (prot. 1) first computes the cost of sending v coins to the receiver as $v_1 := v + \sum_{i=1}^{n} \mathsf{fee}(U_i)$, and the corresponding cost at each of the intermediate hops in the payment path. Then it setups up a PrivMuL by calling the ideal functionality $\mathcal{F}_{\mathbb{L}}$ on the set of identifiers of the intermediate users. Finally, it sends each user the corresponding value to be transferred and a timeout information t_i.

Each intermediate user (prot. 3) checks whether the capacity of the channel is high enough to support the transfer of the coins and whether the timeouts give by the sender are consistent, i.e., $t_{i+1} = t_i - \Delta$ for some fixed Δ. Starting from (U_0, U_1), each pair of users query the ideal functionality $\mathcal{F}_{\mathbb{L}}$ on the **Lock** interface using the *lid* received in the previous phase. If the ideal functionality signals to proceed, then the two users establish a contract specified in the following.

The contract is authenticated by both users and can be uploaded to B by either of them at any time. If every user in the path locks the corresponding *lid*, eventually the receiver (prot. 2) is reached. U_{n+1} checks whether the transacted value is what it expects, and whether the latest timeout t_{n+1} is well-formed. If both conditions hold, the receiver releases the lock lid_n by querying the ideal functionality. This triggers a cascade of release calls in the path from the sender to the receiver, thereby enabling the left user in the link to pull the payment (using the previously established contract). If for some reason one of the intermediate links is not released, then all of the previous contracts are voided after the corresponding timeout.

4) Analysis: In the following we argue that the system as described above ideally realizes the functionality $\mathcal{F}_{PCN}$ as defined in [40], assuming oracle access to $\mathcal{F}_{\mathbb{L}}$ and $\mathcal{F}_{\mathsf{B}}$.

Theorem 6. *The system described above UC-realizes $\mathcal{F}_{PCN}$ in the $(\mathcal{F}_{\mathbb{L}}, \mathcal{F}_{\mathsf{B}})$-hybrid model.*

Proof. The proof consists of the observation that the ideal functionality $\mathcal{F}_{\mathbb{L}}$ enforces balance security and satisfies relationship anonymity (as defined in [40]). A subtlety is that now all users have access to a GetStatus interface and they might be able to query $\mathcal{F}_{\mathbb{L}}$ on a certain *lid* and learn its status even when they are not involved in the generation of such a lock. However one can easily show that this happen only with negligible probability since it require guessing *lid*, which is a string sampled uniformly at random. It is also easy to see that $\mathcal{F}_{\mathbb{L}}$ does not allow one to perform wormhole attacks, by construction. What is left to be shown is that the rest of the information exchanged by the machines does not break any of these properties. Note that the only information that is sent outside $\mathcal{F}_{\mathbb{L}}$ consists of user identifiers, timeouts, and values to lock. The first are already known by the intermediate users, whereas the

Algorithm 1: Payment routine for the sender

Input : $(U_0, \ldots, U_{n+1}, v)$
$v_1 := v + \sum_{i=1}^{n} \mathsf{fee}(U_i)$
if $v_1 \leq \mathsf{cap}(U_0, U_1)$ **then**
  query $\mathcal{F}_{\mathbb{L}}$ **on Setup(**$U_0, \ldots, U_{n+1}$**)**
  $\mathcal{F}_{\mathbb{L}}$ **returns** $(\bot, lid_0, \bot, U_1, \mathsf{Init})$
  $\mathsf{cap}(U_0, U_1) := \mathsf{cap}(U_0, U_1) - v_1$
  $t_0 := t_{\text{now}} + \Delta \cdot n$
  forall $i \in \{1, \ldots, n\}$
    $v_i := v_1 - \sum_{j=1}^{i-1} \mathsf{fee}(U_j)$
    $t_i := t_{i-1} - \Delta$
    send $((U_{i-1}, U_{i+1}, v_{i+1}, t_i, t_{i+1}), \mathsf{fwd})$ **to** U_i
  end for
  send (U_n, v_{n+1}, t_{n+1}) **to** U_{n+1}
  query $\mathcal{F}_{\mathbb{L}}$ **on Lock(**lid_0**)**
  if $\mathcal{F}_{\mathbb{L}}$ **returns** (lid_0, Lock)
    $\mathsf{contract}(U_0, U_1, lid_0, v_1, t_1)$
  else
    abort
  end if else
  abort
end if

Algorithm 2: Payment routine for the receiver

Input : $(U_n, v_{n+1}, t_{n+1}, v)$
$\mathcal{F}_{\mathbb{L}}$ **returns** $(lid_n, \bot, U_n, \bot, \mathsf{Init})$
if $(t_{n+1} > t_{\text{now}} + \Delta) \wedge (v_{n+1} = v) \wedge$ (**GetStatus(**lid_n**)** $= \mathsf{Lock}$) **then**
  query $\mathcal{F}_{\mathbb{L}}$ **on Release(**lid_n**)**
  send ok **to** U_n
else
  send $\bot$ **to** U_n
end if

Algorithm 3: Payment routine for the i-th intermediate user

Input : $(m, decision)$

if $decision = \mathsf{fwd}$ **then**
  parse m **as** $(U_{i-1}, U_{i+1}, v_{i+1}, t_i, t_{i+1})$
  $\mathcal{F}_{\mathbb{L}}$ **returns** $(lid_{i-1}, lid_i, U_{i-1}, U_{i+1}, \mathsf{Init})$
  if $(v_{i+1} \leq \mathsf{cap}(U_i, U_{i+1})) \wedge (t_{i+1} = t_i - \Delta) \wedge$ (**GetStatus(**lid_{i-1}**)** $= \mathsf{Lock}$) **then**
    $\mathsf{cap}(U_i, U_{i+1}) := \mathsf{cap}(U_i, U_{i+1}) - v_{i+1}$
    query $\mathcal{F}_{\mathbb{L}}$ **on Lock(**lid_i**)**
    if $\mathcal{F}_{\mathbb{L}}$ **returns** (lid_i, Lock)
      $\mathsf{contract}(U_i, U_{i+1}, lid_i, v_{i+1}, t_{i+1})$
    else
      send $\perp$ **to** U_{i-1}
    end if
  else
    send $\perp$ **to** U_{i-1}
else if $decision = \perp$ **then**
  $\mathsf{cap}(U_i, U_{i+1}) := \mathsf{cap}(U_i, U_{i+1}) + v_{i+1}$
  send $\perp$ **to** U_{i-1}
else if $(decision = \mathsf{ok}) \wedge$ **GetStatus(**lid_i**)** $= \mathsf{Rel}$ **then**
  query $\mathcal{F}_{\mathbb{L}}$ **on Release(**lid_{i-1}**)**
  send ok **to** U_{i-1}
else
  send $\perp$ **to** U_{i-1}
end if

rest of the items are chosen exactly as described in $\mathcal{F}_{PCN}$. This concludes our argument. $\square$