



Blitz: Secure Multi-Hop Payments Without Two-Phase Commits

Lukas Aumayr, *TU Wien*; Pedro Moreno-Sanchez, *IMDEA Software Institute*;
Aniket Kate, *Purdue University*; Matteo Maffei, *TU Wien*

<https://www.usenix.org/conference/usenixsecurity21/presentation/aumayr>

This paper is included in the Proceedings of the
30th USENIX Security Symposium.

August 11–13, 2021

978-1-939133-24-3

Open access to the Proceedings of the
30th USENIX Security Symposium
is sponsored by USENIX.

Blitz: Secure Multi-Hop Payments Without Two-Phase Commits

Lukas Aumayr
TU Wien
lukas.aumayr@tuwien.ac.at

Aniket Kate
Purdue University
aniket@purdue.edu

Pedro Moreno-Sanchez
IMDEA Software Institute
pedro.moreno@imdea.org

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

Abstract

Payment-channel networks (PCN) are the most prominent approach to tackle the scalability issues of current permissionless blockchains. A PCN reduces the load on-chain by allowing arbitrarily many off-chain multi-hop payments (MHPs) between any two users connected through a path of payment channels. Unfortunately, current MHP protocols are far from satisfactory. One-round MHPs (e.g., Interledger) are insecure as a malicious intermediary can steal the payment funds. Two-round MHPs (e.g., Lightning Network (LN)) follow the 2-phase-commit paradigm as in databases to overcome this issue. However, when tied with economical incentives, 2-phase-commit brings other security threats (i.e., wormhole attacks), staggered collateral (i.e., funds are locked for a time proportional to the payment path length) and dependency on specific scripting language functionality (e.g., Hash Time-Lock Contracts) that hinders a wider deployment in practice.

We present Blitz, a novel MHP protocol that demonstrates for the first time that we can achieve the best of the two worlds: a single round MHP where no malicious intermediary can steal coins. Moreover, Blitz provides the same privacy for sender and receiver as current MHP protocols do, is not prone to the wormhole attack and requires only constant collateral. Additionally, we construct MHPs using only digital signatures and a timelock functionality, both available at the core of virtually every cryptocurrency today. We provide the cryptographic details of Blitz and we formally prove its security. Furthermore, our experimental evaluation on a LN snapshot shows that (i) staggered collateral in LN leads to in between 4x and 33x more unsuccessful payments than the constant collateral in Blitz; (ii) Blitz reduces the size of the payment contract by 26%; and (iii) Blitz prevents up to 0.3 BTC (3397 USD in October 2020) in fees being stolen over a three day period as it avoids wormhole attacks by design.

1 Introduction

Permissionless cryptocurrencies such as Bitcoin enable secure payments in a decentralized, trustless environment. Transactions are verified through a consensus mechanism

and all valid transactions are recorded in a public, distributed ledger, often called blockchain. This approach has inherent scalability issues and fails to meet the growing user demands: In Bitcoin, the transaction throughput is technically limited to tens of transactions per second and the transaction confirmation time is around an hour. In contrast, more centralized payment networks such as the Visa credit card network, can handle peaks of 47,000 transaction per second.

This scalability issue is an open problem in industry and academia alike [15, 31]. Among the approaches proposed so far, payment channels (PC) have emerged as one of the most promising solutions; implementations thereof are already widely used in practice, e.g., the Lightning Network (LN) [22] in Bitcoin. A PC enables two users to securely perform an arbitrary amount of instantaneous transactions between each other, while burdening the blockchain with merely two transactions, (i) for opening and (ii) for closing. In particular, following the unspent transaction output (UTXO) model, two users open a PC by locking some coins in a shared multi-signature output. By exchanging signed transactions that spend from the shared output in a peer-to-peer fashion, they can capture and redistribute their balances off-chain. Either one of the two users can terminate the PC by publishing the latest of these signed transactions on the blockchain.

As creating PCs requires locking up some coins, it is economically infeasible to set up a PC with every user one wants to interact with. Instead, PCs can be linked together forming a graph known as payment channel network (PCN) [19, 22]. In a PCN, a payment of α coins from a sender U_0 to a receiver U_n can be performed via a path $\{U_i\}_{i \in [0, n]}$ of intermediaries.

1.1 State-of-the-art PCNs

A possible way of achieving such a multi-hop payment (MHP) is an optimistic 1-round approach, e.g., Interledger [27]. Here, U_0 starts paying to its neighbor on the path U_1 , who then pays to its neighbor U_2 and so on until U_n is reached. This protocol, however, relies on every intermediary behaving honestly, otherwise any intermediary can trivially steal coins by not forwarding the payment to its neighbor.

To achieve security in MHPs, most widely deployed PCNs (e.g., LN [22]) require an additional second round of communication (i.e., sequential, pair-wise communication between sender and receiver via intermediaries). Specifically, PCNs follow the principles of the 2-phase-commit protocol used to perform atomic updates in distributed databases. In the first communication round, the users on the payment path lock α coins of the PC with their right neighbor in a simple smart contract called Hash Time-Lock Contract (HTLC), which can be expressed even in restricted scripting languages such as the one used in Bitcoin. The money put into the HTLC by the left neighbor at each PC moves to the right neighbor, if this neighbor can present a secret chosen by U_n (i.e., the receiver of the payment); alternatively, it can be reclaimed by the left neighbor after some time has expired.

After HTLCs have been set up on the whole path, the users move to the second round, where they release the locks by passing the secret from U_n to U_0 via the intermediaries on the path before the time on the HTLCs has expired. Intermediaries are economically incentivized to assist in the 2-phase payment protocol. In the first round, when U_i receives α coins from the left neighbor U_{i-1} , it forwards only $\alpha - \text{fee}$ to the right neighbor U_{i+1} , charging fee coins for the forwarding service. In the second round, when U_{i+1} claims the $\alpha - \text{fee}$ coins from U_i , the latter is incentivized to recover the α coins from U_{i-1} .

1.2 Open problems in current PCNs

There are some fundamental problems with current PCNs that follow the 2-phase-commit paradigm. While 2-phase-commit has been successfully used for atomic updates in distributed databases, it is not well suited to applications where economic incentives are inherently involved. In particular, there exists a tradeoff between security, efficiency and number of rounds in the PCN setting that constitutes not only a challenging conceptual problem, but also one with strong practical impact, as we motivate below.

Staggered collateral After a user U_i has paid to U_{i+1} , it must have enough time to claim the coins put by U_{i-1} . If U_{i-1} is not cooperative, then this time is used to forcefully claim the funds with an on-chain transaction. The timing on the HTLCs (called collateral time in the blockchain folklore) grows therefore in a staggered manner from right to left, $t_i \geq t_{i+1} + \xi$. In practice, ξ has to be quite long: e.g., in the LN, it is set to one day (144 blocks). In the worst case, the funds are locked up for a time of $n \cdot \xi$. This means that a single payment of value α over n users can lock up a collateral of $\Theta(n^2 \cdot \alpha \cdot \xi)$. Reducing this locktime enables a faster release time of locked funds and directly improves the throughput of the network. Moreover, long locktimes are also problematic when looking at the high volatility of cryptocurrency prices, where prices can drop significantly within the same day.

Griefing attack A malicious user can start a MHP to itself, causing user U_i to lock up α coins for a time $(n - i) \cdot \xi$. The malicious user subsequently goes idle and lets the payment

fail with the intention of reducing the overall throughput of the network by causing users to lock up their funds. In a different scenario, an intermediary could do the same by accepting payments in the first round, but going idle in the second. It is interesting also to observe the amplification factor: with the relatively small amount of α coins, an attacker can lock $(n - 1) \cdot \alpha$ coins of the network. This attack is hard to detect and can even be used to target specific users in the PCN in order to lock up their funds.

Wormhole attack The wormhole attack [20] is an attack on PCNs where two colluding malicious users skip honest users in the open phase of the 2-phase-commit protocol and thereby cheat them out of their fees. This is problematic as now the payment does not happen atomically anymore: For some users the payment is successful and for others it is not, i.e., for the ones encased by the malicious users. The users for whom it is unsuccessful have to lock up some of their funds, but do not get any fees for offering their services, nor can they use their locked funds for other payments. These fees go instead to the attacker.

HTLC contracts PCNs built on top of 2-phase-commit payments depend largely on HTLCs and the underlying cryptocurrencies supporting them in their scripts. However, there are a number of cryptocurrencies that do not have this functionality or that do not provide scripting capabilities at all, such as Stellar or Ripple. Instead, these currencies provide only digital signature schemes and timelocks.

On a conceptual level, one could actually wonder whether or not it is required to add an agreement protocol (in the database literature, a protocol where if an honest party delivers a message m , then m is eventually delivered by every honest party), like the HTLC-based 2-phase-commit paradigm, on top of the blockchain-inherited consensus protocol.

The current state of affairs thus leads to the following question: *Is it possible to design a PCN protocol with a single round of communication (and thus without HTLCs) while preserving security and atomicity?*

1.3 Our contributions

We positively answer this question by presenting *Blitz*, a novel payment protocol built on top of the existing payment channel constructions, which combines the advantages of both the optimistic 1-round and the 2-phase-commit paradigms. Our contributions are as follows.

- With *Blitz*, we introduce for the first time a payment protocol that achieves a MHP in one round of communication while preserving security in the presence of malicious intermediaries (i.e., as in the LN). The *Blitz* protocol has constant collateral of only $\Theta(n \cdot \alpha \cdot \xi)$, allowing for PCNs that are far more robust against griefing attacks and provide a higher transaction throughput. Additionally, the *Blitz* protocol is immune to the wormhole attack and having only one communication round reduces the chance of unsuccessful payments due to network faults.

- We show that Blitz payments can be realized with only timelocks and signatures, without requiring, in particular, HTLCs. This allows for a more widespread deployment, i.e., in cryptocurrencies that do not feature hashlocks or scripting, but only signatures and timelocks, e.g., Stellar or Ripple. Since Blitz builds on standard payment channel constructions, it can be smoothly integrated as an (alternative or additional) multi-hop protocol into all popular PCNs, such as the LN.

- We formally analyze the security and privacy of Blitz in the Universal Composability (UC) framework. We provide an ideal functionality modeling the security and privacy notions of interest and show that Blitz is a UC realization thereof.

- We evaluate Blitz and show that while the computation and communication overhead is inline with that of the LN, the size of the contract used in Blitz is around 26% smaller than an HTLC in the LN, which in practice opens the door for a higher number of simultaneous payments within each channel. We have additionally evaluated the effect of the reduction of collateral from staggered in the LN to constant in Blitz and observed that it reduces the number of unsuccessful payments due to locked funds by a factor between 4x and 33x, depending on payment amount and percentage of disrupted payments. Finally, the avoidance of the wormhole attack by design in Blitz can save up to 0.3 BTC (3397 USD in October 2020) of fees in our setting (over a three day period).

2 Background and notation

The notation used in this work is adopted from [5]. We provide here an overview on the necessary background and for more details we refer the reader to [5, 19, 20].

2.1 Transactions in the UTXO model

Throughout this work, we consider cryptocurrencies that are built with the *unspent transaction output* (UTXO) model, as Bitcoin is for instance. In such a model, the units of cash, which we will call *coins*, exist in *outputs* of *transactions*. Let us define such an output θ as a tuple consisting of two values, $\theta := (\text{cash}, \phi)$, where $\theta.\text{cash}$ denotes the amount of coins held in this output and $\theta.\phi$ is the condition which must be fulfilled in order to spend this output. The condition is encoded in the scripting language used by the underlying cryptocurrency. We say that a user U owns the coins in an output θ , if $\theta.\phi$ contains a digital signature verification script w.r.t. U 's public key and the digital signature scheme of the underlying cryptocurrency. For this, we use the notation $\text{OneSig}(U)$. If multiple signatures are required, we write $\text{MultiSig}(U_1, \dots, U_n)$.

Ownership of outputs can change via transactions. A transaction maps a non-empty list of existing outputs to a non-empty list of new outputs. For better distinction, we refer to these existing outputs as *transaction inputs*. We formally define a transaction body tx as an attribute tuple $\text{tx} := (\text{id}, \text{input}, \text{output})$. The identifier $\text{tx.id} \in \{0, 1\}^*$ is automatically assigned as the hash of the inputs and outputs, $\text{tx.id} := \mathcal{H}(\text{tx.input}, \text{tx.output})$, where \mathcal{H} is modelled as a




Figure 1: (Left) Transaction tx has two outputs, one of value x_1 that can be spent by B (indicated by the gray box) with a transaction signed w.r.t. pk_B at (or after) round t_1 , and one of value x_2 that can be spent by a transaction signed w.r.t. pk_A and pk_B but only if at least t_2 rounds passed since tx was accepted on the blockchain. (Right) Transaction tx' has one input, which is the second output of tx containing x_2 coins and has only one output, which is of value x_2 and can be spent by a transaction whose witness satisfies the output condition $\phi_1 \vee \phi_2 \vee (\phi_3 \wedge \phi_4)$. The input of tx is not shown.

random oracle. The attribute tx.input is a list of identifiers of the inputs of the transaction, while $\text{tx.output} := (\theta_1, \dots, \theta_n)$ is a list of new outputs. A full transaction $\bar{\text{tx}}$ contains additionally a list of witnesses, which fulfill the spending conditions of the inputs. We define $\bar{\text{tx}} := (\text{id}, \text{input}, \text{output}, \text{witness})$ or for convenience $\bar{\text{tx}} := (\text{tx}, \text{witness})$. Only a valid transaction can be published on the blockchain, i.e., one that has a valid witness for every input and has only inputs not used in other published transactions.

In fact, a transaction is not published on the blockchain immediately after it is submitted, but only after it is accepted through the consensus mechanism. We model that by defining a blockchain delay Δ , an upper bound on the time it takes for a transaction that is broadcast until it is added to the ledger.

For better readability we use charts to visualize transactions, their ordering and how they are used in protocols. The charts are expected to be read from left to right, i.e., the direction of the arrows. Every transaction is represented as a rectangle with rounded corners. Incoming arrows represent inputs. Every transaction has one or more output boxes inside it. Inside these boxes we write the amount of coins stored in the corresponding output. Every output box has one or more outgoing arrow. This arrow has the condition needed to spend the corresponding output written above and below it.

To present complex conditions in a compact way, we use the following notation. On a high level, we write the owner(s) of an output below the arrow and how they can spend it above. In a bit more detail, most output scripts require signature verification w.r.t. one or more public keys, a condition that we represent by writing the necessary public keys below a given arrow. Other conditions are written above the arrow. The conditions above can be any script supported by the underlying cryptocurrency, however in this paper we require only the following. We write “ $+t$ ” or $\text{RelTime}(t)$ to denote a relative timelock, i.e., the output with this condition can be spent, if and only if at least t rounds have passed since the transaction containing the output was published on the blockchain. Additionally, we consider absolute timelocks,

denoted as “ $\geq t$ ” or $\text{AbsTime}(t)$: this condition is satisfied if and only if the blockchain is at least t blocks long. If an output condition is a disjunction of several conditions, i.e., $\phi = \phi_1 \vee \dots \vee \phi_n$, we write a diamond shape in the output box and put each subcondition ϕ_i above/below its own arrow. For the conjunction of several conditions we write $\phi = \phi_1 \wedge \dots \wedge \phi_n$. We illustrate an example of our transaction charts in Figure 1.

2.2 Payment channels

A payment channel is used by two parties P and Q to perform several payments between them while requiring only two on-chain transactions (for opening and closing). The balances are kept and updated in what is called a state. For brevity and readability, we hereby abstract away from the implementation details of a payment channel and provide a more detailed description in Appendix C.

We assume that there is an off-chain transaction tx^{state} which holds the outputs representing the current state of the payment channel. We further assume that the current tx^{state} can always be published on the blockchain and if an old state is published by a dishonest user, the honest user gets the total channel balance through some punishment mechanism.

Formally, we define a channel $\bar{\gamma}$ as the following attribute tuple $\bar{\gamma} := (\text{id}, \text{users}, \text{cash}, \text{st})$. Here, $\bar{\gamma}.\text{id} \in \{0, 1\}^*$ is a unique identifier of the channel, $\bar{\gamma}.\text{users} \in \mathcal{P}^2$ denotes the two parties that participate in the channel out of the set of all parties \mathcal{P} . Further, $\bar{\gamma}.\text{cash} \in \mathbb{R}_{\geq 0}$ stores the total number of coins held in the channel and $\bar{\gamma}.\text{st} := (\theta_1, \dots, \theta_n)$ is the current state of the channel consisting of a list of outputs. For convenience, we also define a channel skeleton γ with respect to a channel $\bar{\gamma}$ as the tuple $\gamma := (\bar{\gamma}.\text{id}, \bar{\gamma}.\text{users})$. When the channel is used along a payment path as shown in the next section, we say the $\gamma.\text{left} \in \gamma.\text{users}$ accesses the user that is closer to the sender and $\gamma.\text{right} \in \gamma.\text{users}$ the one closer to the receiver. The balance of each user can be inferred from the state $\bar{\gamma}.\text{st}$, however for convenience we define a function $\bar{\gamma}_i.\text{balance}(U)$, that returns the coins of user $U \in \gamma_i.\text{users}$ in this channel.

2.3 Payment channel networks

Since maintaining a payment channel locks a certain amount of coins for a party, it is economically prohibitive to set up a payment channel with every party that one potentially wants to interact with. Instead, each party may open channels with a few other parties, creating thereby a network of channels. A payment channel network (PCN) [19] is thus a graph where vertices represent the users and edges represent channels between pairs of users. In a PCN, a user can pay any other user connected through a path of payment channels between them. Suppose user U_0 wants to pay some amount α to U_n , but does not have a payment channel directly with it. Now assume that instead, U_0 has a payment channel $\bar{\gamma}_0$ with U_1 , who in turn has a channel $\bar{\gamma}_1$ with U_2 and so on, until the receiver U_n . We say that U_0 and U_n are connected by a path and denote a payment using it as *multi-hop payment* (MHP).

Optimistic payment schemes In an MHP, the main challenge is to ensure that the payment happens atomically and for everyone, so that no (honest) user loses any money. In fact, there exists payment-channel network constructions where this security property does not hold. We call them *optimistic payment schemes* and give Interledger [27] as an example. In this scheme, the users on the path simply forward the payment without any guarantee of the payment reaching the receiver. The sender U_0 starts by performing an update for channel $\bar{\gamma}_0$, where $\bar{\gamma}_0.\text{balance}(U_1)$ is increased by α (and $\bar{\gamma}_0.\text{balance}(U_0)$ is decreased by α) compared to the previous state. U_1 does the same with U_2 and this step is repeated until the receiver U_n is reached. This scheme works if every user is honest. However, a malicious intermediary can easily steal the money by simply stopping the payment and keeping the money for itself.

Secure MHPs Since the assumption that every user is honest is infeasible in practice, most widely deployed systems instead ensure that no honest user loses coins. The Lightning Network (LN) [22] uses so called Hash Time-Lock Contracts (HTLCs). An HTLC works as follows. In a payment channel between Alice and Bob, party Alice locks some coins that belong to her in an output that is spendable in the following fashion: (i) After some pre-defined time t , Alice will get her money back. (ii) Bob can also claim the money at any time, if he knows a pre-image r_A for a certain hash value $\mathcal{H}(r_A)$, which is set by Alice.

For an MHP in the LN, suppose again that we have a sender U_0 who wants to pay α to a receiver U_n via some intermediaries U_i with $i \in [1, n-1]$, and that two users U_j and U_{j+1} for $j \in [0, n-1]$ have an opened payment channel. Now for the first step, U_n samples a random number r , computes the hash of it $y := \mathcal{H}(r)$ and sends y to U_0 . In the second step, the sender U_0 sets up an HTLC with U_1 by creating a new state with three outputs $\theta_1, \theta_2, \theta_3$ that correspondingly hold the amount of coins: α , U_0 's balance minus α and U_1 's balance. While θ_2 and θ_3 are spendable by their respective owners, θ_1 is the output used by the HTLC. The HTLC that is constructed spends the output containing α back to U_0 after n time, let us say n days, or to U_1 if it knows a value x such that $\mathcal{H}(x) = y$. Now U_1 repeats this step with its right neighbor, again using y but a different time, $(n-1)$ days, in the HTLC. This step is repeated until the receiver is reached, with a timeout of one day.

Now if constructed correctly, the receiver U_n can present r to its left neighbor U_{n-1} , which is the secret required in the HTLC for giving the money to U_n . We call this *opening the HTLC*. After doing that, the two parties can either agree to update their channel to a new state, where U_n has α coins more, or otherwise the receiver can publish the state and a transaction with witness r spending the money from the HTLC to itself on-chain. When a user U_i reveals the secret r to its left neighbor U_{i-1} , U_{i-1} can use r to continue this process. For this continuation, U_{i-1} needs to have enough time. Otherwise, U_i could claim the money of the HTLC it has with U_{i-1} by

spending the HTLC on-chain at the last possible moment. Because of the blockchain delay, user U_{i-1} will notice this too late and will not be able to claim the money of the HTLC with U_{i-2} anymore. This is the reason why the timelocks on the HTLCs are staggered, i.e., increasing from right to left.

The aforementioned process where each user presents r to the left neighbor is repeated until the sender U_0 is reached, at which point the payment is completed. We call this approach of performing MHPs *2-phase-commit*.

3 Solution overview

The goal of this work is to achieve the best of the two multi-hop payment (MHP) paradigms existing nowadays (optimistic and 2-phase-commit), that is, an MHP protocol with a single round of communication that overcomes the drawbacks of the current LN MHP protocol and yet maintains the security and privacy notions of interest.

For that, we propose a paradigm shift, which we call *pay-or-revoke*. The idea is to update the payment channels from sender to receiver in a single round of communication. The key technical challenge is thus to design a single channel update that can be used *simultaneously* for sending coins from the left neighbor to the right one if the payment is successful and for a refund of the coins to the left neighbor if the payment is unsuccessful (e.g., one intermediary is offline).

We present the pay-or-revoke paradigm in an incremental way, starting with a naive design, discussing the problems with it, and presenting a tentative solution. We iterate these steps until we finally reach our solution.

Naive approach Assume a setting with a sender U_0 who wants to pay α coins to a receiver U_n via a known path of some intermediaries U_i ($i \in [1, n-1]$), where each pair of consecutive users U_j and U_{j+1} for $j \in [0, n-1]$ has a payment channel $\bar{\gamma}_j$, where $\bar{\gamma}_j.\text{balance}(U_j) \geq \alpha$. We start out with an optimistic payment scheme, as presented in Section 2.3. We already explained that the success of such a payment relies on every intermediary behaving honestly and really forwarding the α coins. Should an intermediary not forward the payment, U_n will never receive anything. Additionally, a receiver could claim that it never received the money even though it actually did and it would be difficult for the sender to prove otherwise.

To solve these problems the sender faces when using this form of payment we introduce a possibility for the sender to step back from a payment, that is, refund itself and all subsequent users the α coins that they initially put, should the payment not reach U_n . With such a refund functionality, the sender can now check if a receiver is giving a confirmation that it got the payment. This confirmation is external to the system (e.g., a digital payment receipt) and serves additionally as a proof that the money was received. If such a confirmation is not received, the sender simply steps back from the payment and the payments in every channel are reverted.

Adding refund functionality Adding a refund functionality while avoiding additional security problems is challenging.

Two neighbors can no longer simply update their channel $\bar{\gamma}_i$ to a state where α coins are moved from the left to the right neighbor, as this only encodes the payment. Instead, we need to introduce an intermediate channel state tx^{state} , which encodes the possibility for both a refund and a payment.

We realize that as follows. This new state has an output holding α coins coming from $\bar{\gamma}_i.\text{left}$ ($= U_i$) while leaving the rest of the balance in the channel untouched. The output containing α coins becomes then the input for two mutually exclusive transactions: refund and payment. We denote the refund transaction as tx_i^r , which spends the money back to $\bar{\gamma}_i.\text{left}$ ($= U_i$). We denote the payment transaction as tx_i^p , which spends the money to $\bar{\gamma}_i.\text{right}$ ($= U_{i+1}$). The refund should only be possible until a certain time T . This gives the sender time to wait for the payment to reach the receiver and for the receiver to give a (signed) confirmation. Should something go wrong, the sender starts the refund procedure. After time T , if no refund happened, the payment is considered successful and the payment transaction becomes valid.

The latter condition can easily be expressed in the scripting language of virtually any cryptocurrency including Bitcoin, by making use of absolute timelocks, which in this work we defined as $\text{AbsTime}(T)$, meaning an output can be spent only after some time T . Unfortunately, the same cannot be done for expressing the condition that an output is spendable *only before* time T (e.g., see [13] for details).

We overcome this problem in a different way. Instead of making the refund transaction tx_i^r only valid before T , we allow both tx_i^r and the payment transaction tx_i^p to be valid after time T and encode a condition that, should both be posted after T , tx_i^p will always be accepted over tx_i^r . We can achieve this by adding a relative timelock on the input of tx_i^r of the blockchain delay Δ . In other words, should a user try to close the channel with tx^{state} appearing on the chain after time T , the other user will have enough time to react and post tx_i^p , which will get accepted before the relative timelock of tx_i^r expires. For the honest refund case nothing changes: If tx^{state} is on-chain and tx_i^r gets posted before $T - \Delta$, it will always be accepted over tx_i^p , since the latter transaction is only valid after time T .

Making the refund atomic So far, we added a refund functionality that is (i) not atomic and (ii) triggerable by every user on the path. An obvious attack on this scheme would be for any user on the path to commence the refund in a way that tx_i^r is accepted on the ledger just before T . Other users would not have enough time to react accordingly and lose their funds. Also, allowing intermediary users to start the refund opens up the door to grieving, where malicious users start a refund even though the payment reached the receiver. We therefore need a mechanism that (i) ensures the atomicity of the refund (or payment) and (ii) is triggerable only by the sender.

Following the LN protocol, one could add a condition $\mathcal{H}(r_A)$ on the refund transaction, such that the refund can only happen when a pre-image r_A chosen by the sender is

known. To prevent the sender from publishing at the last moment however, the timing for the refund in the next channel would have to be $T + \Delta$ to give U_1 enough time to react. In subsequent channels, this time would grow by Δ for every hop and we would then have an undesirable staggered time delay. Additionally, this approach would rely on the scripting language supporting hash-lock functionality.

To keep the time delay constant, we instead make the refund transactions dependent on a transaction being published by the sender. First, the sender creates a transaction that we name *enable-refund* and denote by tx^{er} . The unsigned transaction tx^{er} is then passed through the path and is used at each channel \bar{y}_j as an additional input for tx_j^{r} .

This makes the refund transaction at every channel dependent on tx^{er} and gives the sender and only the sender the possibility to abort the payment until time T in case something goes wrong along the path (e.g., a user is offline or the enable-refund transaction is tampered), and the receiver the guarantee to get the payment after time T otherwise.

In order to use the same tx^{er} for the refund transaction tx_j^{r} of every channel \bar{y}_j , we proceed as follows. For every user on the path (except for the receiver) there needs to exist an output in tx^{er} which belongs to it. Additionally, we observe that an intermediary U_i whose left neighbor U_{i-1} has used tx^{er} as input for its refund transaction $\text{tx}_{i-1}^{\text{r}}$ can safely construct a refund transaction tx_i^{r} dependent on the same tx^{er} , because it will know that if its left neighbor refunded, tx^{er} has to be on-chain, which means that it can refund itself. Also, since the appearance of tx^{er} on the ledger is a global event that is observable by everyone at the same time, the time T used for the refund can be the same for every channel, i.e., constant.

Putting everything together Our approach is depicted in Figure 2, tx^{er} is shown in Figure 3, and the transaction structure between two users is shown in Figure 4. Note that we change the payment value from α to α_i to embed a per-hop fee (see Appendix A for details). After the payment is set up from sender to receiver, the receiver sends a confirmation of tx^{er} back to U_0 , which acts both as verification that tx^{er} was not tampered and as a payment confirmation. Should the sender receive this in time, it will wait until time T , after which the payment will be successful. If no confirmation was received in time, or tx^{er} was tampered, the sender will publish tx^{er} in time to trigger the refund.

We remark that it is crucial that every intermediate user can safely construct tx_i^{r} only observing tx^{er} , but not the input funding it (or not even knowing whether it will be funded at all in the first place). Indeed, an intermediary U_i does not care if the transaction tx^{er} is spendable at all, it only cares that its left neighbor U_{i-1} uses an output of the same transaction tx^{er} as input for its refund transaction $\text{tx}_{i-1}^{\text{r}}$, as U_i does in tx_i^{r} .

In UTXO based cryptocurrencies, using the j^{th} output of a transaction tx as input of another transaction tx' means referencing the hash of the transaction body $\mathcal{H}(\text{tx})$, which we defined as tx.id , plus an index j . A transaction tx_i^{r} that was




Figure 2: Illustration of the pay-or-revoke paradigm.




Figure 3: Transaction tx^{er} , which enables the refunds and, here, spends the output of some other transaction tx^{in} .

created with an input referencing tx^{er} .id and some index j , can only be valid if tx^{er} is published. This means, in particular, that it is computationally infeasible to create a different transaction $\text{tx}^{\text{er}'} \neq \text{tx}^{\text{er}}$ and use one of $\text{tx}^{\text{er}'}$'s outputs as input of tx_i^{r} without finding a collision in \mathcal{H} . Further, as tx_i^{r} requires the signatures of both U_i and U_{i+1} , a malicious U_i on its own cannot create a different refund transaction $\text{tx}_i^{\text{r}'}$ that does not depend on tx^{er} .

A final timelock There is however still one subtle problem with the construction up to this point regarding the timing coming from the fact that the sender has the advantage of being the only one able to trigger the refund by publishing tx^{er} . In a bit more detail, as closing a channel takes some time, a malicious sender U_0 can forcefully close its channel with U_1 beforehand. Then, when $\text{tx}_0^{\text{state}}$ is on the ledger, the sender publishes tx^{er} so that it appears just before $T - \Delta$. The sender is able to publish $\text{tx}_0^{\text{state}}$ just in time before T . All other intermediaries however, who did not yet close their channel, with the result that $\text{tx}_i^{\text{state}}$ is not on the ledger, will not be able to do this and publish tx_i^{r} in time.

To solve this problem, we introduce a relative timelock on the outputs of tx^{er} of exactly $t_c + \Delta$, as shown in Figure 3 and

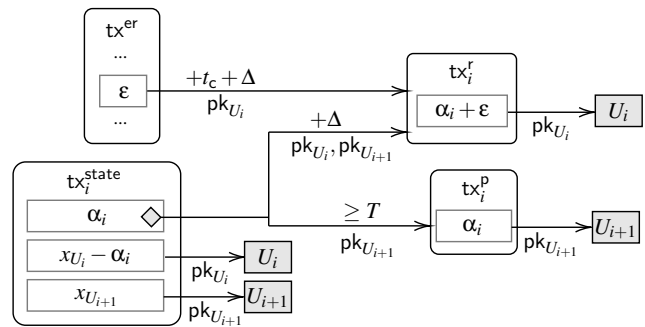


Figure 4: Payment setup in the channel \bar{y}_i of two neighboring users U_i and U_{i+1} with the new state tx^{state} . x_{U_i} and $x_{U_{i+1}}$ are the amounts that U_i and U_{i+1} own in the state prior to tx^{state} .

Figure 4. This relative time delay is an upper bound on the time it takes to (i) forcefully close the channel and (ii) wait for the time delay needed to publish tx_i^r . With this, we ensure that no user gains an advantage by closing its channel in advance, since this can be entirely done in this relative timelock on tx^{er} 's outputs. Honest intermediaries can easily check that this relative timelock is present in tx^{er} 's outputs and every user on the payment path has the same time.

A timeline of when the transactions have to appear on the ledger is given in Appendix E. Note that for the payment to be refunded, tx^{er} has to be posted to the ledger at the latest at time $T - t_c - 3\Delta$. Still, for better readability we sometimes refer to this case simply as tx^{er} being published before time T .

Improving anonymity of the path Until this point, we have shown a design of the pay-or-revoke paradigm, that, while ensuring that honest users do not lose coins, has an obvious drawback in terms of anonymity. In particular, the transaction outputs of tx^{er} contain the addresses of every user on the path in the clear (except for the receiver who does not need to refund and therefore needs no such output). This means that every intermediary (or any other user that sees tx^{er}) learns about the identity of every user on the payment path as soon as it sees tx^{er} . To prevent this leak, we use stealth addresses [29]. We overview our use of stealth addresses here and refer to Section 4.2 for technical details. On a high level, instead of spending to existing addresses, the sender uses fresh addresses for the outputs of tx^{er} . These addresses were never used before, but are under the control of the respective users. With this approach, if tx^{er} is leaked, the identities of all users on the path, especially the identity of the sender and the receiver, remain hidden. Note that we assume the input of tx^{er} to be an unused and unlinkable input of the sender.

Fast track payments The design considered so far has still a practical drawback compared to MHPs in the LN. In the LN, if every user is honest, the payment is carried out almost instantaneously, i.e. the channels are updated as soon as the HTLCs are opened. Obviously, users of a payment do not want to wait until some time T until the payment is carried out, even if all users are honest. To enable the same fast payments in Blitz, we extend the protocol design with an *optional* second communication round, called the fast track (we compare this second round to the one adopted in the LN below). Specifically, the users on the path can honestly update their channels from the sender to the receiver to a state where the α coins move from left to right.

For this, the sender does not go idle upon receiving the confirmation in time from the receiver. Instead, U_0 starts updating the channel \bar{y}_0 with its neighbor U_1 to a state where the α coins are paid to U_1 . Since U_0 is the only one able to publish tx^{er} , U_0 is safe when performing this update. After this update, U_1 does the same with U_2 . All users on the path repeat this step until the receiver is reached. If everyone is honest, the payment will be carried out as quick as in the LN honest case. If someone stops the update or some honest users

are skipped by colluding malicious users, honest users simply wait until time T , and claim their money (and fees) either by cooperatively updating the channel with their neighbor or forcefully on-chain. Intuitively, since intermediary users only update their right channel after updating their left channel, they cannot lose any money, even if tx^{er} is published.

Using the fast track seems to be a better choice for normal payments. However, there are applications, where the non fast track is more suitable, e.g., a service with a trial period or a subscription model, where a user might want to set up a payment, that gets confirmed after some time. Should the user decide against it, he/she can cancel the payment. The choice of fast track is up to the user. Having this second round is completely optional and for efficiency reasons only. A payment that is carried out in one round has the same security properties as one carried out in two rounds.

Fast revoke In the case that an intermediary is offline and the payment is unsuccessful, the refund can happen without necessarily publishing tx^{er} , saving the cost to put a transaction on-chain. Say U_{i+1} is offline and U_i has already set up the construction with U_{i-1} . As soon as an honest U_i notices that U_{i+1} is unresponsive, it can start asking U_{i-1} to update their channel to the state before the payment was set up. After doing this, U_{i-1} asks its left neighbor to do the same and so on until the sender is reached and the payment is reverted without tx^{er} being published. Should some intermediary refuse to honestly revoke, then tx^{er} can still be published. Apart from funds being locked for a shorter time, one could add additional incentives to the fast revocation (or fast track) by giving a small fee to the users that are willing to participate in it. Of course, users need a mechanism to find out whether others are offline. For that, we note that the LN protocol mandates users to periodically broadcast a heartbeat message. We consider such default messages orthogonal to payment protocols and do not count them in round complexity.

Honest update The transactions in Figure 4 between users are exchanged off-chain and used to guarantee that honest users do not lose any coins. However, should one of the users in a channel be able to convince the other that it is able to enforce either tx_i^r or tx_i^p on-chain (that is if tx^{er} is on-chain before time T or time T has already passed, respectively), two collaborating users can simply perform an honest update. For this, they update their channel to a state where both have their corresponding balance, with the benefit that no transaction has to be put on-chain and their channel remains open.

Blitz vs. ILP/LN/AMHL We claim that Blitz is a solution for the issues presented in Section 1 and allows for PCNs that have higher throughput, less communication complexity, additional security against certain attacks, and are implementable in cryptocurrencies without scripting capabilities. We highlight the differences between Blitz and other state-of-the-art payment methods such as Interledger Payments (ILP), the LN and the wormhole secure construction Anonymous Multi-Hop Locks (AMHL) [20] in Table 1.

Table 1: Features of different payment methods: Interledger (ILP), Lightning Network (LN), Anonymous Multi-Hop Locks (AMHL), Blitz and Blitz using the fast track payment (FT). We abbreviate timelocks as TL and signature functionality as σ . * The requirement of HTLC can be dropped from the LN using scriptless scripts when feasible.

	ILP	LN	AMHL	Blitz	Blitz FT
Bal. Security	No	Yes	Yes	Yes	Yes
Rounds	1	2	2	1	2
Atomicity	No	No (Wormhole)	Yes	Yes	Yes
Scripting	σ	σ , TL, HTLCs*	σ , TL	σ , TL	σ , TL
Collateral	n/a	linear	linear	constant	constant

Table 2: Collateral time for the LN, AMHL and Blitz for unsuccessful (refund) and successful payments (pay) as well as different threat models. We say *instant* when no one on the path stops the payment in either round. ξ denotes the time users need to claim their funds (e.g., in the LN 144 blocks).

	LN / AMHL		Blitz	
	refund	pay	refund	pay
anyone malicious	$n \cdot \xi$	$n \cdot \xi$	ξ	ξ
sender honest	$n \cdot \xi$	$n \cdot \xi$	Δ	ξ
everyone honest	instant	instant	instant	instant

First, Blitz offers balance security with only one round of communication, while ILP does not provide that and the LN requires two rounds. While the fast track optimization does involve a second round (from left to right, as opposed to right to left as in the LN), it is optional and affects only the efficiency (in the case everyone is honest) and not security: a payment that had a successful first round will be successful regardless of any network faults in the second round.

Indeed, the same holds true for the wormhole attack: Once a user has successfully set up a Blitz payment, it cannot be skipped anymore in the second round, even with the fast track. The payment is successful for everyone or no one, achieving thus the atomicity property missing in ILP and the LN, and honest intermediaries are not cheated out of their fees.

Secondly, Blitz reduces the collateral from linear (in the size of the path) to constant in the case some of the parties are malicious, while offering comparable performance in the optimistic case, as shown in Section 6. For a corner case where the sender is honest, the collateral can even be unlocked almost instantaneously. We show in which cases Blitz outperforms the LN in Table 2. Finally, in terms of interoperability, we require only signatures and timelocks from the underlying blockchain, with the LN additionally requiring HTLCs and ILP only signatures.

Concurrent payments In Blitz, multiple payments can be carried out in parallel, analogous to concurrent HTLC-based payments in the LN (see Appendix A for further discussion and an illustrative example).

4 Our construction

4.1 Security and privacy goals

We informally review the security and privacy goals of a PCN, deferring the formal definitions to the full version [6].

Balance security Honest intermediaries do not lose money [19].

Sender/Receiver privacy In the case of a successful payment, malicious intermediaries cannot determine if the left neighbor along the path is the actual sender or just an honest user connected to the sender through a path of non-compromised users. Similarly, malicious intermediaries cannot determine if the right neighbor is the actual receiver or an honest user connected to the receiver through a path of non-compromised users.

Path privacy In the case of a successful payment, malicious intermediaries cannot determine which users participated in the payment aside from their direct neighbors.

4.2 Assumptions and building blocks

System assumptions We assume that every party has a publicly known pair of public keys (A, B) as required for stealth address creation (see below). We further assume that honest parties are required to stay online for the duration of the protocol. Finally, we consider the route finding algorithm an orthogonal problem and assume that every user (U_0) has access to a function $\text{pathList} \leftarrow \text{GenPath}(U_0, U_n)$, which generates a valid path from U_0 to U_n over some intermediaries. We refer the reader to [24, 25] for more details on recent routing algorithms for PCNs. We now introduce the cryptographic building blocks that we require in our protocol.

Ledger and payment channels We rely, as a blackbox, on a public ledger to keep track of all balances and transactions and a PCN that supports the creation, update, and closure of channels (see Section 2). We further assume that payment channels between users that want to conduct payments are already opened. We denote the standard operations to interact with the blockchain and the channels as follows:

$\text{publishTx}(\bar{\text{tx}})$: If $\bar{\text{tx}}$ is a valid transaction (Section 2), it will be accepted on the ledger after at most time Δ .

$\text{updateChannel}(\bar{y}_i, \text{tx}_i^{\text{state}})$: When called by a user $\in \bar{y}_i$.users, initiates an update in \bar{y}_i to the state $\text{tx}_i^{\text{state}}$. If the update is successful, (update-ok) is returned to both users of the channel, else (update-fail) is returned to them. We define t_u as an upper bound on the time it takes for a channel update after this procedure is called.

$\text{closeChannel}(\bar{y}_i)$: When called by a user $\in \bar{y}_i$.users, closes the channel, such that the latest state transaction $\text{tx}_i^{\text{state}}$ will appear on the ledger. We define t_c as an upper bound on the time it takes for $\text{tx}_i^{\text{state}}$ to appear on the ledger after this procedure is called.

Digital signatures A digital signature scheme is a tuple of algorithms $\Sigma := (\text{KeyGen}, \text{Sign}, \text{Vrfy})$ defined as follows:

$(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\lambda)$ is a PPT algorithm that on input

the security parameter λ , outputs a pair of public and private keys (pk, sk) .

$\sigma \leftarrow \text{Sign}(sk, m)$ is a PPT algorithm that on input the private key sk and a message m outputs a signature σ .

$\{0, 1\} \leftarrow \text{Vrfy}(pk, \sigma, m)$ is a DPT algorithm that on input the public key pk , an authentication tag σ and a message m , outputs 1 if σ is a valid authentication for m .

We require that the digital signature scheme is correct, that is, $\forall (pk, sk) \leftarrow \text{KeyGen}(\lambda)$ it must hold that $1 \leftarrow \text{Vrfy}(pk, \text{Sign}(sk, m), m)$. We additionally require a digital signature scheme that is strongly unforgeable against message-chosen attacks (EUF-CMA) [14].

Stealth addresses [29] On a high level, this scheme allows a user (say Alice) to derive a fresh public key in a digital signature scheme Σ controlled by another user (say Bob) on input two of Bob's public keys. In a bit more detail, a *stealth addresses* scheme is a tuple of algorithms $\Phi := (\text{GenPk}, \text{GenSk})$ defined as follows:

$(P, R) \leftarrow \text{GenPk}(A, B)$ is a PPT algorithm that on input two public keys A, B controlled by some user U , creates a new public key P under U 's control. This is done by first sampling some randomness $r \leftarrow_{\$} [0, l - 1]$, where l is the prime order of the group used in the underlying signature scheme Σ , and computing $P := g^{\mathcal{H}(A^r)} \cdot B$, where \mathcal{H} is a hash function modelled as a random oracle. Then, the value $R := g^r$ is calculated. P is the public key under U 's control and R is the information required to construct the private key.

$p \leftarrow \text{GenSk}(a, b, P, R)$ is a DPT algorithm that on input two secret keys a, b corresponding to the two public keys A, B and a pair (P, R) that was generated as $P \leftarrow \text{GenPk}(A, B)$, creates the secret key p corresponding to P . This is done by computing $p := \mathcal{H}(R^a) + b$.

We see that correctness follows directly: $g^p = g^{\mathcal{H}(R^a) + b} = g^{\mathcal{H}(g^{ra})} \cdot g^b = g^{\mathcal{H}(A^r)} \cdot B = P$. In [29] it is argued that this new one-time public key P is *unlinkable* for a spectator even when observing R , meaning on a high level that P for some user U cannot be linked to any existing public key of U . For simplicity, we denote $\tilde{U}_i, pk_{\tilde{U}_i}$ when referring to the stealth identity or the stealth public key under the control of user U_i .

Anonymous communication network (ACN) An ACN allows users to communicate anonymously with each other. One such ACN is based on onion routing, whose ideal functionality is defined in [8]. Sphinx [10] is a realization of this and (extended with a per-hop payload) is used in the Lightning Network (LN). We use this functionality here as well in a blackbox way. On a high level, routing information and a per-hop payload is encrypted and layered for every user along a path, in what is called an *onion*. Every user on the path can then, when it is its turn, “peel off” such a layer, revealing: (i) the next neighbor; (ii) the payload meant for it; and (iii) the rest of the data, which is again an onion that can only be opened by the next neighbor. This rest of data is then forwarded to the next user and so on until the receiver is reached.

For readability, we use two algorithms, where $\text{onion} \leftarrow \text{CreateRoutingInfo}(\{U_i\}_{i \in [1, n]}, \{\text{msg}_i\}_{i \in [1, n]})$ creates such a routing object (an onion) using (publicly known) public encryption keys of the corresponding users on the path. Moreover, when called by correct user U_i , the algorithm $\text{GetRoutingInfo}(\text{onion}_i, U_i)$ returns $(U_{i+1}, \text{msg}_i, \text{onion}_{i+1})$, that is, the next user on the path, a message and a new onion or returns msg_n if called by the recipient. A wrong user $U \neq U_i$ calling $\text{GetRoutingInfo}(\text{onion}_i, U_i)$ will result in an error \perp .

4.3 2-party protocol for channel update

In this section, we show the necessary steps to update a single channel $\bar{\gamma}_i$ between two consecutive users U_i and U_{i+1} on a payment path to a state encoding our payment functionality as shown in Figure 4. We will describe later in Section 4.4 the complete multi-hop payment (MHP) protocol.

As overviewed in Section 3, a channel update requires to create a series of transactions to realize the “pay-or-revoke” semantics at a given channel. In particular, for readability, we define the following transaction creation methods and in Figure 7 some macros to be used hereby in the paper:

$\text{tx}_i^p := \text{GenPay}(\text{tx}_i^{\text{state}})$ This transaction takes $\text{tx}_i^{\text{state}}.\text{output}[0]$ as input and creates a single output $:= (\alpha_i, \text{OneSig}(U_{i+1}))$.

$\text{tx}_i^r := \text{GenRef}(\text{tx}_i^{\text{state}}, \text{tx}^{\text{er}}, \theta_{\epsilon_i})$ This transaction takes as input $\text{tx}_i^{\text{state}}.\text{output}[0]$ and $\theta_{\epsilon_i} \in \text{tx}^{\text{er}}.\text{output}$. The calling user U_i makes sure that this output belongs to a stealth address under U_i 's control. It creates a single output $\text{tx}_i^r.\text{output} := (\alpha_i + \epsilon, \text{OneSig}(U_i))$, where α_i, U_i, U_{i+1} are taken from $\text{tx}_i^{\text{state}}$.

We now explain in detailed order, how these transactions have to be created, signed and exchanged. A full description in pseudocode is given in Figure 5. This two party update procedure, which we call pcSetup , is called by a user U_i giving as parameters the channel $\bar{\gamma}_i$ with its right neighbor U_{i+1} , the transaction tx^{er} , a list containing the values R_i for the stealth addresses of each user on the path, onion_{i+1} containing some routing information for the next user, the output $\theta_{\epsilon_i} \in \text{tx}^{\text{er}}.\text{output}$ that belongs to a stealth address of U_i , the amount to be paid α_i and the time T . The user U_i knows these values either from performing pcSetup with its left neighbor U_{i-1} or because U_i is the sender.

The first step for U_i is to create the new channel state from the channel $\bar{\gamma}_i$ and the amount α_i by calling $\text{tx}_i^{\text{state}} := \text{genState}(\bar{\gamma}_i, \alpha)$. In the second step, U_i creates the transaction tx_i^r from $\text{tx}_i^{\text{state}}.\text{output}[0]$ and θ_{ϵ_i} . Then, U_i sends $\text{tx}^{\text{er}}, \text{tx}_i^{\text{state}}, \text{tx}_i^r, \text{rList}$ and onion_{i+1} to its right neighbor U_{i+1} .

Now U_{i+1} checks if tx^{er} is well-formed and, if it is not the receiver, has an output $\theta_{\epsilon_{i+1}}$, which belongs to its stealth address (using its stealth address private keys a, b) under some $R_i \in \text{rList}$. Moreover, it checks that onion_{i+1} contains the correct routing information and a message indicating that the tx^{er} was not tampered, for instance a hash of it. All this is done using the macro (see Figure 7) $(\text{sk}_{\tilde{U}_{i+1}}, \theta_{\epsilon_{i+1}}, R_{i+1}, U_{i+2}, \text{onion}_{i+2}) :=$

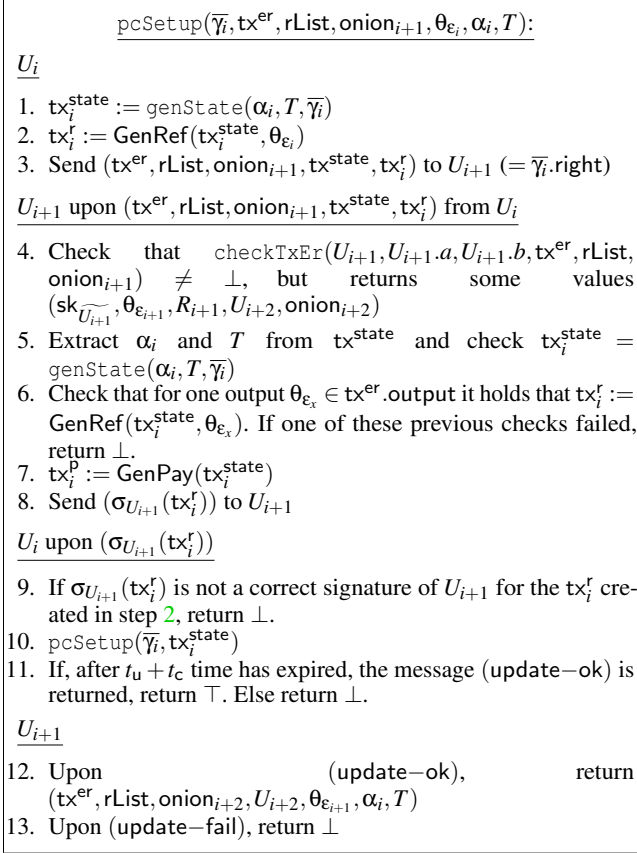


Figure 5: Protocol for 2-party channel update

$\text{checkTxEr}(U_{i+1}, U_{i+1}.a, U_{i+1}.b, U_{i+1}.\text{tx}^{\text{er}}, \text{rList}, \text{onion}_{i+1})$, which returns \perp if any of the checks fail.

Then, U_{i+1} checks if $\text{tx}_i^{\text{state}}$ and tx_i^f were well-constructed and in particular, that tx_i^f uses an output of tx^{er} as input. If everything is ok, then U_{i+1} can independently create tx_i^p , since it requires only its own signature. Next, U_{i+1} pre-signs tx_i^f and sends this signature to U_i . U_i checks if this signature is correct and then invokes a channel update with U_{i+1} to $\text{tx}_i^{\text{state}}$.

After this step, the pcSetup function is finished and returns either $(\text{tx}^{\text{er}}, \text{rList}, \text{onion}_{i+2}, U_{i+2}, \theta_{\varepsilon_{i+1}}, \alpha_i, T)$ to U_{i+1} and \top to U_i if successful or \perp otherwise to the users \overline{y}_i .users. If U_{i+1} is not the receiver, it will continue this process with its own neighbor as shown in the next section.

4.4 Multi-hop payment description

In this section we describe the MHP protocol. The pseudocode for carrying out MHPs in Blitz is shown in Figure 6, the macros used in are listed in Figure 7. For the full description of the macros, see Appendix H.

Setup Say the sender wants to pay α coins to U_n via a path channellist and for some timeout T . In the setup phase, the sender derives a new stealth address $\text{pk}_{\overline{U}_i}$ and some R_i for every user except the receiver. Then, the sender creates a list rList of entries R_i and onions encoding the right neighbor U_{i+1} for every user U_i . Moreover, the sender constructs tx^{er} .

Then, it adds the sum of all per-hop fees to the initial amount α : $\alpha_i := \alpha + (n - 1) \cdot \text{fee}$ where fee is the fee charged by every user (see Appendix A). The setup ends when the sender starts the open phase with its right neighbor U_1 .

Open After successfully setting up the payment with its left user U_{i-1} , U_i knows tx^{er} , rList , onion_{i+1} , α_{i-1} , T and its stealth output for $\theta_{\varepsilon_i} \in \text{tx}^{\text{er}}.\text{output}$. Using these values and reducing α_{i-1} by fee, U_i carries out the 2-party channel update with U_{i+1} . The right neighbor continues this step with its right neighbor until the receiver is reached.

Finalize Once the receiver has finished the open phase with its left neighbor, it sends back a signature of tx^{er} as a confirmation to the sender, who will then check if that transaction was tampered with. If yes, or if the sender did not receive such a confirmation in time, the sender publishes tx^{er} on the blockchain. Otherwise the sender goes idle.

Respond At any given time after opening a payment construction, users need to check if tx^{er} was published. If it was, they need to refund themselves via tx_i^f . Also, if some user's left neighbor tries to publish tx_i^f after time T , the user publishes tx_i^p . This ensures, that if the refund did not happen before time T , the users have a way to enforce the payment. Note that due to the relative timelock on both tx^{er} and $\text{tx}_i^{\text{state}}$, tx_i^p will always be possible if tx^{er} is published after T (or if the left neighbor tries to refund after T by closing the channel).

The protocol is shown in Figure 6. Note that we simplified the protocol for readability purposes, (e.g., by omitting the payment ids that are required for multiple concurrent payments). The full protocol modelled in the Universal Composability framework can be seen in the full version [6].

5 Security analysis

5.1 Security model

The security model we use closely follows [5, 11, 12]. We model the security of Blitz in the synchronous, global universal composability (GUC) framework [9]. We use a global ledger \mathcal{L} to capture any transfer of coins. The ledger is parameterized by a signature scheme Σ and a blockchain delay Δ , which is an upper bound on the number of rounds it takes between when a transaction is posted to \mathcal{L} and when said transaction is added to \mathcal{L} . Our security analysis is fully presented in the full version [6] and briefly outlined here.

Firstly, we provide an ideal functionality \mathcal{F}_{Pay} , which is an idealized description of the behavior we expect of our pay-or-revoke payment paradigm. This description stipulates any input/output behavior and the impact on the ledger of a payment protocol, as well as how adversaries can influence the execution. In this idealized setting, all parties communicate only with \mathcal{F}_{Pay} , which acts as a trusted third party.

We then provide our protocol Π formally defined in the UC framework and show that Π emulates \mathcal{F}_{Pay} . On a high level, we show that any attack that can be performed on Π can also be simulated on \mathcal{F}_{Pay} or in other words that Π is at least as

Setup
U_0 upon receiving $(\text{setup}, \text{channelList}, \text{tx}^{\text{in}}, \alpha, T)$
<ol style="list-style-type: none"> 1. If $\text{checkChannels}(\text{channelList}, U_0) = \perp$, abort. 2. Let $n := \text{channelList}$. If $\text{checkT}(n, T) = \perp$, abort. 3. If $\text{checkTxIn}(\text{tx}^{\text{in}}, n, U_0) = \perp$, abort. 4. $(\text{tx}^{\text{er}}, \text{rList}, \text{onion}) := \text{genTxEr}(U_0, \text{channelList}, \text{tx}^{\text{in}})$ 5. $\alpha_0 := \alpha + \text{fee} \cdot (n - 1)$ 6. $(\text{sk}_{\tilde{U}_0}, \theta_{\epsilon_0}, R_0, U_1, \text{onion}_1)$:= <li style="padding-left: 20px;">$\text{checkTxEr}(U_0, U_0.a, U_0.b, \text{tx}^{\text{er}}, \text{rList}, \text{onion})$ 7. $\text{pcSetup}(\overline{\gamma_0}, \text{tx}^{\text{er}}, \text{rList}, \text{onion}_1, U_1, \theta_{\epsilon_0}, \alpha_0, T)$
Open
U_{i+1} upon receiving $(\text{tx}^{\text{er}}, \text{rList}, \text{onion}_{i+2}, U_{i+2}, \theta_{\epsilon_{i+1}}, \alpha_i, T)$
<ol style="list-style-type: none"> 1. If U_{i+1} is the receiver U_n, send $(\text{confirm}, \sigma_{U_n}(\text{tx}^{\text{er}})) \leftrightarrow U_0$ and go idle. 2. $\text{pcSetup}(\overline{\gamma_{i+1}}, \text{tx}^{\text{er}}, \text{rList}, \text{onion}_{i+2}, U_{i+2}, \theta_{\epsilon_{i+1}}, \alpha_i - \text{fee}, T)$
Finalize
U_0 : Upon $(\text{confirm}, \sigma_{U_n}(\text{tx}^{\text{er}})) \leftarrow U_n$, check that $\sigma_{U_n}(\text{tx}^{\text{er}})$ is U_n 's valid signature for the transaction tx^{er} created in the Setup phase. If not, or if tx^{er} was changed, or no such confirmation was received until $T - t_c - 3\Delta$, publish $\text{Tx}(\text{tx}^{\text{er}}, \sigma_{U_0}(\text{tx}^{\text{er}}))$.
Respond (Executed in every round τ_x)
<ol style="list-style-type: none"> 1. If $\tau_x < T - t_c - 2\Delta$ and tx^{er} on the blockchain, closeChannel($\overline{\gamma_i}$) and, after $\text{tx}_i^{\text{state}}$ is accepted on the blockchain within at most t_c rounds, wait Δ rounds. Let $\sigma_{\tilde{U}_i}(\text{tx}_i^{\text{r}})$ be a signature using the secret key $\text{sk}_{\tilde{U}_i}$. publish $\text{Tx}(\text{tx}_i^{\text{r}}, (\sigma_{\tilde{U}_i}(\text{tx}_i^{\text{r}}), \sigma_{U_i}(\text{tx}_i^{\text{r}}), \sigma_{U_{i+1}}(\text{tx}_i^{\text{r}})))$. 2. If $\tau_x > T$, $\overline{\gamma_i}$ is closed and tx^{er} and $\text{tx}_i^{\text{state}}$ is on the blockchain, but not tx_i^{r}, publish $\text{Tx}(\text{tx}_{i-1}^{\text{p}}, (\sigma_{U_i}(\text{tx}_{i-1}^{\text{p}})))$.

Figure 6: The Blitz payment protocol

secure as \mathcal{F}_{Pay} . To prove this, we design a simulator \mathcal{S} , which translates any attack on the protocol into an attack on the ideal functionality. Then, we show that no PPT *environment* can distinguish between interacting with the real world and interacting with the ideal world. In the real world, the environment sends instructions to a real attacker \mathcal{A} and interacts with Π . In the ideal world, the environment sends attack instructions to \mathcal{S} and interacts with \mathcal{F}_{Pay} .

We need to show that the same messages are output in the same rounds and the same transactions are posted on the ledger in the same rounds in both the real and the ideal world, regardless of adversarial presence. To achieve this, the simulator needs to instruct the ideal functionality to output a message whenever one is output in the real protocol and the simulator needs to post the same transactions on the ledger. By achieving this, the environment cannot trivially distinguish between the real and the ideal world anymore just by looking at the messages and transactions as well as their respective timing. Formally, in the full version [6] we prove Theorem 1.

Theorem 1. (informal) *Let Σ be a EUF-CMA secure signature scheme. Then, for any ledger delay $\Delta \in \mathbb{N}$, the protocol Π UC-realizes the ideal functionality \mathcal{F}_{Pay} .*

Macros (see Appendix H)
<p>checkTxIn($\text{tx}^{\text{in}}, n, U_0$): If tx^{in} is well-formed and has enough coins, returns \top. checkChannels($\text{channelList}, U_0$): If channelList forms a valid path, returns the receiver U_n, else \perp. checkT(n, T): If T is sufficiently large, return \top. Otherwise, return \perp. genTxEr($U_0, \text{channelList}, \text{tx}^{\text{in}}$): Generates tx^{er} from tx^{in} along with a list of values rList to redeem their stealth addresses and an onion containing the routing information. genState($\alpha_i, T, \overline{\gamma_i}$): Generates and returns a new channel state carrying transaction $\text{tx}_i^{\text{state}}$ from the given parameters, shown in Figure 4. checkTxEr($U_i, a, b, \text{tx}^{\text{er}}, \text{rList}, \text{onion}_i$): Checks if tx^{er} is correct, U_i has a stealth address in it and onion_i holds routing information. If unsuccessful, returns \perp. If U_i is the receiver, returns $(\top, \top, \top, \top, \top)$. Else, returns $(\text{sk}_{\tilde{U}_i}, \theta_{\epsilon_i}, R_i, U_{i+1}, \text{onion}_{i+1})$ containing the output belonging to U_i; θ_{ϵ_i}, the secret key to spend it $\text{sk}_{\tilde{U}_i}$, the next user and the next onion.</p>

Figure 7: Subprocedures used in the protocol

5.2 Informal security discussion

Due to space constraints, we only argue informally here why Blitz achieves security and privacy (see Section 4.1). We give a more formal discussion in the full version [6] and consider the security against some concrete attacks in Appendix D.

Balance security An honest intermediary will forward a payment to its right neighbor only if first invoked by its left neighbor. If constructed correctly, the refund transactions in both channels depend on tx^{er} being published and the timing is identical. Also, the payment transactions have identical conditions in both channels. The only possible way for an intermediary to lose money is, if it were to pay its money to the right neighbor, while the left neighbor refunded. However, if the left neighbor is able to refund, this means that also the intermediary itself can refund. Similarly, if the right neighbor is able to claim the money, the intermediary can also claim it.

Honest sender A sender that does not receive a confirmation of the receiver that it received the money in time, can trigger a refund by publishing tx^{er} . In the setup phase of the protocol, the sender ensures that there is enough time for this.

Honest receiver The receiver gets the money in exchange for some service. It will wait until being certain that the money will be received before shipping the product. The transaction tx^{er} on the blockchain is a proof that a refund has occurred.

Privacy Blitz requires to share with intermediaries tx^{er} , routing information and the value that is being paid. The transaction tx^{er} uses stealth addresses for its outputs and an unlinkable input, thereby granting sender, receiver and path privacy in the honest case, as defined in Section 4.1. As in the LN however, the stronger notion of relationship anonymity [19] does not hold; the payment can be linked by comparing (i) in Blitz, tx^{er} and (ii) in the LN, the hash value. In the pessimistic case, the balance is claimed on-chain. In both Blitz and the LN, this breaks sender, path and receiver privacy. We defer the reader to Appendix A for a more detailed discussion on all privacy properties mentioned in this paragraph.

6 Evaluation

In this section, we evaluate the benefits that Blitz offers over the LN. The source code for our simulation is at [1].

Testbed We took a snapshot of the LN graph (October 2020) from <https://ln.bigsun.xyz/> containing 11.6k nodes, 6.5k of which have 30.9k active channels with a total capacity of 1166.7 BTC, which account for around 13.2M USD in October 2020. We ignore the nodes without active channels. The initial distribution of the channel balance is unknown. We assume that initially the balance at each channel is available to both users. It is assigned to a user as required by payments in a first come, first serve basis. Naturally, the balance that has already been used and thus assigned to one user in the channel, is not reassigned to the other user. Since we use this strategy consistently throughout all our experiments, this assignment does not introduce any bias in the results.

Simulation setup We discretize the time in rounds and each round represents the collateral time per hop (i.e., 1 day or 144 blocks as in the LN). In such a setting, we simulate payments in batches as follows. Assume that we want to simulate $N\text{Pay}$ payments for an amount of Amt and with a failure rate of $F\text{Rate}$. For that, in a first batch we simulate the $F\text{Rate}$ % of $N\text{Pay}$ payments, where each payment is between two nodes s and r (such that $s \neq r$) selected at random in the graph and routed through the cheapest path according to fees. Moreover, each payment in this batch is disrupted at an intermediary node chosen at random in the path between s and r . Finally, for each payment, some balance is marked to be locked at the channels for a certain number of rounds during the second batch, depending on whether we are evaluating the LN (i.e., staggered rounds) or Blitz (i.e., single round). We model thereby a setting where the network contains locked collateral due to disrupted payments.

After the first batch, we simulate a second one of $N\text{Pay}$ payments over 3 rounds as before, assuming that they are not disrupted (e.g., go over paths of honest nodes). We remark that here each payment may still be unsuccessful because there are not enough unlocked funds in the path between s and r . We focus thus on the effect that staggered vs. constant collateral has in the number of successful payments.

Setting parameters Due to the off-chain nature of the LN, there is no ground truth for payment data, a common limitation in PCN related work. We try to make reasonable assumptions for these unknown parameters in our simulation. We sample the payment amount Amt for each payment from the range $[1000, ub]$. We use a lower bound of 1000, as technically the minimum is 546 satoshis (=1 dust) and we additionally account for fees. We select an upper bound (ub) out of $\{3000, 6000, 9000\}$, which is around 0.1%, 0.2% and 0.3% of the average channel capacity. We consider two different number of payments $N\text{Pay}$, 78k and 978k. The former corresponds to four payments per active node and per round (ppnpr) modeling a setting with sporadic payments (e.g., a




Figure 8: Ratio $\text{fail}_{\text{LN}}/\text{fail}_{\text{Blitz}}$. (Left) we fix the number of disrupted payments at 0.5% and vary ub . (Right) we fix ub at 3000 and vary the number of disrupted payments.

banking system), whereas the latter corresponds to 50 ppnpr, modeling a higher payment frequency (e.g., micropayments).

Finally, we vary the amount of disrupted payments $F\text{Rate}$ as $\{0.5, 1, 2.5\}$ % of the total payments $N\text{Pay}$. We divide these disrupted payments into two groups of equal size. In the first half, the payment is stopped during the setup phase (from s to r). In the LN, the channels before the faulty/malicious node are locked with a staggered collateral lock time. In Blitz, due to the sender publishing tx^{er} , the funds are immediately unlocked. In the second half, the payment fault occurs in the second phase, which in the LN is the unlocking and in Blitz the fast track. This models the case where a node is offline or an attacker delays the completion of the payment until the last possible moment. In the LN, the collateral left of the malicious node is again staggered, whereas in Blitz the channels right of that node are locked for one simulation round. Finally, we note that distributing the disrupted payments differently into these groups will alter the results accordingly (see Appendix G).

Collateral effect We calculate the number of unsuccessful payments in a baseline case (i.e., omitting the first batch of disrupted payments), in Blitz as well as in the LN and we say that $\text{fail}_{\text{Blitz}}$ (correspondingly fail_{LN}) is the number of payments that fail in Blitz (correspondingly the LN) when subtracting those failing also in the baseline case. We carry out every experiment for a given setting eight times and calculate the average. In Figure 8 we show the ratio $\text{fail}_{\text{LN}}/\text{fail}_{\text{Blitz}}$. For all choices of parameters, there are more unsuccessful payments in the LN than in Blitz, showing thus the practical advantage of Blitz by requiring only constant collateral. We also observe that difference grows in favor of Blitz with the number of payments, showing that the advantage in terms of collateral is higher in use cases for which initially the LN was designed such as micropayments. Finally, we observe that Blitz offers higher transaction throughput even with an arguably small ratio of disrupted payments (i.e., a reduced adversarial effect).

Wormhole attack We measure an upper bound on the amount of fees potentially at risk in the LN, due to it being prone to the wormhole attack. We observe that the amount of coins at risk grows with the number of payments and their amount. In particular, with 50 ppnpr and an upper bound of

3000 (modeling e.g., a micropayment setting), we observe that the LN put at risk 0.25 BTC (2831 USD in October 2020). Increasing the upper bound to 9000 while keeping 65 ppnr, we observe that the LN put at risk 0.30 BTC. Blitz prevents the wormhole attack and the stealing of these fees by design.

Computation overhead The Blitz protocol does not require any costly cryptography. In particular, it requires that each user verifies locally the signatures for the involved transactions. Moreover, each user must compute three signatures (see Figure 4) independently on the number of channels involved in the payment. In the LN, each user requires to compute only two signatures, one per each commitment transaction representing the new state. We remark, however, that these are all simple computations that can be executed in negligible time even with commodity hardware.

Communication overhead We find that the contract size in Blitz is 26% smaller than the size of the HTLCs in the LN. This advantage is crucial in practice as current LN payment channels cannot hold more than 483 HTLC (and thus 483 in-flight payments) simultaneously, because otherwise, the size of the off-chain state would be higher than a valid Bitcoin transaction [23, 28]. The reduced communication overhead in Blitz implies then that it allows for more simultaneous in-flight payments per channel than in the LN.

In the pessimistic case, the LN requires to include on-chain one transaction per channel (158 Bytes for refund, 192 Bytes for payment), while Blitz requires not only one on-chain transaction per channel (307 Bytes for refund, 158 Bytes for payment), but also that the sender includes the transaction tx^{er} to ensure that the refund is atomic. In this sense, the LN requires a smaller overhead than Blitz for the pessimistic case. We remark that there exist incentives in PCNs for the nodes to follow the optimistic case and reduce entering the pessimistic case because it requires to close the channels and cannot be used for further off-chain payments without re-opening them, with the consequent cost in time and fees. We give detailed results about communication overhead in Appendix F.

7 Related work

PCNs have attracted plenty of attention from academia [13, 18, 20, 21, 27] and have been deployed in practice [22]. These PCNs, with the exception of Interledger [27], follow the 2-phase-commit paradigm and suffer from (some of) the drawbacks we have discussed in this work, namely, prone to the wormhole attack, griefing attacks, staggered collateral or rely on scripting functionality not widely available. Interledger is a 1-phase protocol that however does not provide security.

Sprites [21] is the first multi-hop payment (MHP) that achieves constant collateral. It, however, relies on Turing complete smart contracts (available in, e.g., Ethereum) thereby reducing its applicability in practice. Other constructions that require Turing complete smart contracts, e.g., State channels [11], achieve constant collateral, but have similar privacy issues as the LN when used for MHPs. AMCU [13] achieves

constant collateral and is compatible with Bitcoin. AMCU, however, reveals every participant to each other, a privacy leakage undesirable in the MHP setting.

To improve privacy, [19] introduced MHTLCs. In [30], CHTLCs based on Chameleon hash functions were introduced, a functionality that is again not supported in most cryptocurrencies (e.g., in Bitcoin). AMHL [20] replaces the HTLC contract with novel cryptographic locks to avoid the wormhole attack. MHTLC, CHTLC or AMHL based MHPs all follow the 2-phase-commit paradigm and require staggered collateral. We defer to Appendix B for works on 1-phase commits in the context of distributed databases.

8 Conclusion

Payment-channel networks (PCNs) are the most prominent solution to the scalability problem of cryptocurrencies with practical adoption (e.g., the LN). While optimistic 1-round payments (e.g., Interledger) are prone to theft by malicious intermediaries, virtually all PCNs today follow the 2-phase-commit paradigm and are thus prone to a combination of: (i) security issues such as wormhole attacks; (ii) staggered collateral; and (iii) limited deployability as they rely on either HTLC or Turing complete smart contracts.

We find a redundancy implementing a 2-phase-commit protocol on top of the consensus provided by the blockchain and instead design Blitz, a multi-hop payment protocol that demonstrates for the first time that it is possible to have a 1-round payment protocol that is secure, resistant to wormhole attacks by design, has constant collateral, and builds upon digital signatures and timelock functionality from the underlying blockchain's scripting language. Our experimental evaluation shows that Blitz reduces the number of unsuccessful payments by a factor of between 4x and 33x, reduces the size of the payment contract by a 26% and saves up to 0.3 BTC (3397 USD in October 2020) in fees over a three day period as it avoids wormhole attacks by design.

Blitz can be seamlessly deployed as a (additional or alternative) payment protocol in the current LN. We believe that Blitz opens possibilities of performing more efficient and secure payments across multiple different cryptocurrencies and other applications built on top, research directions which we intend to pursue in the near future.

Acknowledgements This work has been supported by the European Research Council (ERC) under the Horizon 2020 research (grant 771527-BROWSEC); by the Austrian Science Fund (FWF) through the projects PROFET (grant P31621), the Meitner program (grant M-2608) and the project W1255-N23; by the Austrian Research Promotion Agency (FFG) through the Bridge-1 project PR4DLT (grant 13808694) and the COMET K1 SBA; by the Vienna Business Agency through the project Vienna Cybersecurity and Privacy Research Center (VISP); by CoBloX Labs; by the National Science Foundation (NSF) under grant CNS-1846316.

References

- [1] Blitz simulation: Github repository, 2020. <https://github.com/blitz-payments/simulation>.
- [2] M. Abdallah, R. Guerraoui, and P. Pucheral. One-phase commit: does it make sense? *Conference on Parallel and Distributed Systems*, 1998.
- [3] Yousef J. Al-houmaily and Panos K. Chrysanthis. Two-Phase Commit in Gigabit-Networked Distributed Databases. In *Parallel and Distributed Computing Systems*, 1995.
- [4] Yousef J Al-Houmaily and Panos K Chrysanthis. 1-2PC: the one-two phase atomic commit protocol. In *Symposium on Applied Computing*, 2004.
- [5] Lukas Aumayr, Oguzhan Ersoy, Andreas Erwig, Sebastian Faust, Kristina Hostakova, Matteo Maffei, Pedro Moreno-Sanchez, and Siavash Riahi. Generalized Bitcoin-Compatible Channels. Cryptology ePrint Archive. <https://eprint.iacr.org/2020/476>.
- [6] Lukas Aumayr, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. Blitz: Secure Multi-Hop Payments Without Two-Phase Commits. Cryptology ePrint Archive, Report 2021/176, 2021. <https://eprint.iacr.org/2021/176>.
- [7] Vivek Bagaria, Joachim Neu, and David Tse. Boomerang: Redundancy Improves Latency and Throughput in Payment-Channel Networks. In *FC*, 2020.
- [8] Jan Camenisch and Anna Lysyanskaya. A Formal Treatment of Onion Routing. In *CRYPTO*, 2005.
- [9] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally Composable Security with Global Setup. In *TCC*, 2007.
- [10] G. Danezis and I. Goldberg. Sphinx: A Compact and Provably Secure Mix Format. In *IEEE S&P*, 2009.
- [11] Stefan Dziembowski, Lisa Eckey, Sebastian Faust, Julia Hesse, and Kristina Hostáková. Multi-party Virtual State Channels. In *EUROCRYPT*, 2019.
- [12] Stefan Dziembowski, Lisa Eckey, Sebastian Faust, and Daniel Malinowski. PERUN: Virtual Payment Channels over Cryptographic Currencies. In *IEEE S&P*, 2019.
- [13] Christoph Egger, Pedro Moreno-Sanchez, and Matteo Maffei. Atomic Multi-Channel Updates with Constant Collateral in Bitcoin-Compatible Payment-Channel Networks. In *CCS*, 2019.
- [14] Shafi Goldwasser, Silvio Micali, and Ronald L Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on computing*, 1988.
- [15] Lewis Gudgeon, Pedro Moreno-Sanchez, Stefanie Roos, Patrick McCorry, and Arthur Gervais. SoK: Off The Chain Transactions. In *FC*, 2020.
- [16] Rachid Guerraoui and Jingjing Wang. How Fast can a Distributed Transaction Commit? In *PODS*, 2017.
- [17] Maurice Herlihy, Liuba Shrira, and Barbara Liskov. Cross-chain Deals and Adversarial Commerce. *VLDB*, 2019.
- [18] Aggelos Kiayias and Orfeas Stefanos Thyfronitis Litos. A Composable Security Treatment of the Lightning Network. In *CSF*, 2019.
- [19] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Srivatsan Ravi. Concurrency and Privacy with Payment-Channel Networks. In *CCS*, 2017.
- [20] Giulio Malavolta, Pedro Moreno-Sanchez, Clara Schneidewind, Aniket Kate, and Matteo Maffei. Anonymous Multi-Hop Locks for Blockchain Scalability and Interoperability. In *NDSS*, 2019.
- [21] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. Sprites: Payment Channels that Go Faster than Lightning. In *FC*, 2019.
- [22] Joseph Poon and Thaddeus Dryja. The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments, 2016. <https://lightning.network/lightning-network-paper.pdf>.
- [23] EmelyanenkoK (pseudonym). Payment channel congestion via spam-attack. <https://github.com/lightningnetwork/lightning-rfc/issues/182>.
- [24] Stefanie Roos, Pedro Moreno-Sanchez, Aniket Kate, and Ian Goldberg. Settling Payments Fast and Private: Efficient Decentralized Routing for Path-Based Transactions. In *NDSS*, 2018.
- [25] Vibhaalakshmi Sivaraman, Shaileshh Bojja Venkatarishnan, Kathleen Ruan, Parimarjan Negi, Lei Yang, Radhika Mittal, Giulia C. Fanti, and Mohammad Alizadeh. High Throughput Cryptocurrency Routing in Payment Channel Networks. In *NSDI*, 2020.
- [26] James W Stamos and Flavio Cristian. Coordinator log transaction execution protocol. *Distributed and Parallel Databases*, 1993.
- [27] Stefan Thomas and Evan Schwartz. A protocol for interledger payments, 2015. <https://interledger.org/interledger.pdf>.
- [28] Sergei Tikhomirov, Pedro Moreno-Sanchez, and Matteo Maffei. A Quantitative Analysis of Security, Anonymity and Scalability for the Lightning Network. In *IEEE S&B Workshop*, 2020.
- [29] Nicolas Van Saberhagen. Cryptonote v 2.0, 2018. <https://cryptonote.org/whitepaper>.
- [30] Bin Yu, Shabnam Kasra Kermanshahi, Amin Sakzad, and Surya Nepal. Chameleon Hash Time-lock Contract for Privacy Preserving Payment Channel Networks. In *Conference on Provable Security*, 2019.
- [31] Alexei Zamyatin, Mustafa Al-Bassam, Dionysis Zindros, Eleftherios Kokoris-Kogias, Pedro Moreno-Sanchez, Aggelos Kiayias, and William J. Knottenbelt. SoK: Communication Across Distributed Ledgers. In *FC*, 2021.

A Discussion on practical deployment

Payment fees We encode a fee mechanism in our construction. For simplicity, we assume that every intermediary charges the same fee amount: fee . However, it is trivial to extend this mechanism to allow for different fees. The sender initially puts an amount $\alpha_0 := \alpha + \text{fee} \cdot (n - 1)$ in the output $\theta_{i,0}$. Every intermediary now deducts fee from this amount when opening the construction with its own right neighbor. Specifically, an intermediary U_i receives α_{i-1} and forwards only $\alpha_i := \alpha_{i-1} - \text{fee}$. Thereby, every intermediary effectively gains fee coins in the case of a successful payment.

Refund tradeoff In the case of a refund, where a fast refund (see Section 3) is not possible, the sender has to publish tx^{er} . Doing this will have the cost of publishing this transaction (and possibly the transaction containing its input) plus the $(n - 1) \cdot \varepsilon$ that go to the intermediaries. The amount ε can be the smallest possible amount of cash, since it is just used to enable the payment. In other words, for Bitcoin we can say $\varepsilon := 1 \text{ satoshi}$ ¹, which is currently around 0.00011 USD. However, the refund of Blitz payments has a fundamental advantage over the one in the Lightning Network (LN). The refund time is only constant in the worst case and if the sender is honest, is only the time it takes to publish tx^{er} (i.e., Δ) instead of $n \cdot \xi$. We presented this advantage in Section 3.

So the tradeoff is a more expensive, but much faster refund. This immensely reduces the effect of grieving attacks and increases the overall transaction throughput.

Race We already mentioned that only the sender can publish tx^{er} and because of the time delays, the timing is the same for every user on the path. We claimed that the latest possible time to safely publish tx^{er} and still be able to claim the refund is $T - t_c - 3\Delta$. However, there is a time frame after $T - t_c - 3\Delta$ up until $T - t_c - 2\Delta$, where the sender could publish tx^{er} and still, tx_i^{f} would be sent to the ledger before time T . However now, everyone is at risk, because we said that accepting a transaction takes at most Δ time and at time T , already tx_i^{p} might be sent to the ledger and there might be a race over which of these two transactions is accepted first. We argue, that a sender will not do this, as this puts himself at the same risk as every other intermediary. For a way of preventing this race entirely, we defer the reader to the full version [6].

Obfuscate the length of the path By adding additional dummy outputs (that belong to fresh addresses of the sender) to tx^{er} , a sender can obfuscate the path length. Note that the rList has to include some random values as well, so that it has the same number of elements as tx^{er} has outputs. Note that by looking at the time lock in the LN, the path length or at least one position within the path is leaked to some degree.

¹In practice, Bitcoin transactions need to carry a total amount of one dust, which is 546 satoshis. Having individual outputs of one satoshi is not a problem, as the sender can include an additional output to a stealth address under its control, such that the sum is greater than one dust. In tx_i^{f} , the output of tx^{er} holding one satoshi is combined with the first output of the state tx^{state} , resulting in a sum larger than one dust.

Extended privacy discussion As mentioned in Section 4.1, Blitz achieves sender, receiver and path privacy, which provide a measure of privacy in the case of a successful payment. To hide the path from users observing tx^{er} , we use stealth addresses for the outputs of tx^{er} . This allows to have path privacy as defined in Section 4.1, where malicious intermediaries cannot determine the participants of the payment other than their direct neighbors. We stress that as in the LN, the stronger notion of relationship anonymity [19] does not hold. Two users can link a payment by comparing the transaction tx^{er} in Blitz, or the hash value in the LN.

To make an on-chain linking of the sender impossible, we require the input of tx^{er} to be fresh and unlinkable to the sender. In practice, this can be achieved as follows. The sender creates off-chain an intermediary transaction tx^{in} that spends from an output under the sender's control tx^{sd} to a newly generated address of the sender, never used before. Then, tx^{er} uses this output with the new address of tx^{in} as input. Since tx^{in} is off-chain, users observing tx^{er} are unable to link the payment to an on-chain identity. Again, this is due to inputs referring to a transaction hash plus an id of the output.

In the pessimistic case, these properties do not hold anymore. If the transactions go on-chain, they can be linked together by observing a shared transaction tx^{er} or time T . The same holds true in the LN, where transactions that spend from an HTLC with the same hash value, can be linked.

Redundancy for improving throughput and latency Routing a payment through a path can fail or be delayed due to unknown channel balances, offline or malicious users or other reasons. Following Boomerang [7], a sender can construct several redundant payments across several paths, that differ in one or more users. For this, the sender creates a transaction tx^{er} for each of these redundant payments and forwards them. Intermediary users have to open a payment construction (build tx_i^{f} and tx_i^{p}) for every tx^{er} that they receive.

Should an intermediary user have a choice of forwarding a payment to several different neighbors, it can choose one and start a fast refund (Section 3) for the other payments. Should several different payments reach the receiver, it can start the fast refund for all but one of them. In the worst case, if the sender sees that after some time more than one payment is active, it can start the refund by publishing the according transaction tx^{er} . With this, the sender can ensure that at most one of the redundant payments is carried out. This technique is useful to improve transaction throughput and latency and we achieve it without any additional cryptography.

Concurrent payments Two parties of a payment channel can achieve concurrent payments as follows. They agree to update their current channel state $\text{tx}_i^{\text{state}}$ to a new state $\text{tx}_i^{\text{state}'}$, where any unresolved in-flight Blitz payments are carried over. More concretely, for every unresolved payment the transactions tx_i^{f} and tx_i^{p} are recreated, but the input for these transactions is changed from using an output of $\text{tx}_i^{\text{state}}$ to using an output of $\text{tx}_i^{\text{state}'}$. Afterwards, the right user's signature

for tx_i^f is given to the left user and only then, the old state $\text{tx}_i^{\text{state}}$ is revoked using the revocation technique in the LN (outlined in Appendix C). In other words, the same channel state-management of the LN is reused in Blitz, but changing the HTLC contract for the Blitz contract. We show an illustrative example of concurrent payments in Figure 9.

B 1-phase commits in distributed databases

The concepts of 1-phase commits [2, 3, 26] and one-two commit [4] have been studied for distributed databases in general. These protocols introduce recovery mechanisms such as coordinator Log [26], implicit Yes-Vote [3] or logical logging [2] towards avoiding the voting/commit/prepare phase of 2-phase commits. However, extending observation by Herlihy, Liskov, and Shrira [17], traditional 1-phase commit ideas are not directly applicable to PCNs: while PCNs (with blockchain-based conflict resolution) are structurally similar to transactions over distributed database, they are fundamentally different in terms of the ACID properties and the adversarial assumptions. Nevertheless, analyses such as [16] can still be interesting to understand lower-bounds for PCNs.

C Payment channels in more detail

In this section, we give a more detailed account on payment channels. A payment channel is used by two parties P and Q to perform several payments between them while requiring only two on-chain transactions. It is set up by two parties spending some coins to a shared multisig output (i.e., an output θ with $\theta.\phi := \text{MultiSig}(P, Q)$). Before signing and publishing this transaction however, they create transactions (so called *commitment transactions* tx^c) that spend this shared output in some way, e.g., giving each party some balance. We also refer to this as the (current) *state* of the channel. Now after publishing this tx^f on-chain, they can update their balances by creating new commitment transactions tx^c , rebalancing the funds of the channel and thereby carrying out payments. We note that there are implementations that use two commitment transactions per state (in other words, one per party) such as the Lightning Network (LN) [22] whereas a more recent construction called generalized channels [5] requires one commitment transaction per state. In this work, we leverage the latter construction, although other ledger channel protocols such as the one of the LN would work as well.

After a channel has been updated several times, there exist several tx^c that can be published. In order to prevent misbehavior, where one party publishes an older state of the channel, which perhaps is financially more advantageous to it, we employ a punishment mechanism. If an old state is published, the other, honest user can carry out this punishment to gain all funds of the channel. For this to work, both parties exchange revocation secrets every time a state is succeeded by a new one. This secret, together with the outdated tx^c that is published by the misbehaving user is enough to claim all funds of the channel. The latest state can always be safely

published as the corresponding revocation secret was not yet revealed. This mechanism provides an economical incentive not to publish an old tx^c .

To close a payment channel, the parties can merely publish the latest tx^c to the ledger, which terminates the channel. In summary, two parties can use a payment channel to carry out arbitrary many off-chain payments that rebalance some funds, but only need to publish two transactions on the blockchain, one to open the channel and one to close it, saving both fees and increasing the cryptocurrency's transaction throughput.

D Concrete attack scenarios (informal)

In this section, we consider some attacks against Blitz and argue informally, why balance security still holds.

tx^{er} is tampered If tx^{er} is tampered by some intermediary, the next intermediary will see that the message embedded in the routing information is not $\mathcal{H}(\text{tx}^{\text{er}})$ anymore. Assuming that a malicious intermediary does not know the routing information especially not the receiver, changing the routing information will result in the receiver not being reached.

Also, note that balance security holds even in the case where tx^{er} is tampered, as long as every intermediary U_i makes sure, that its refund tx_i^f depends on the same tx^{er} as the refund of its neighbor tx_{i-1}^f . Also note, that intermediaries have to ensure the same for the time T , in order to have the same time as their neighbor. Should an intermediary change the time T to a smaller value, it potentially only hurts itself by not being able to refund in time, while its left neighbor actually is. If the time T is changed to a larger value, this may delay the execution of the payment, however it is detectable, if the receiver sends this time T back to the sender, who can check if it was tampered.

Some users are skipped (wormhole) Users cannot be skipped, as the routing information can only be opened by the next user. A malicious user would not know the receiver and would not be able to forge the sender's signature of $\mathcal{H}(\text{tx}^{\text{er}})$ that is embedded as a message to the receiver in this onion. The only thing for the malicious user is to stop forwarding the payment (griefing attack). Users that are skipped in the fast-track payment will not be cheated out of their fees or funds, rather this money will be locked until at most until T instead of being accessible immediately (see Section 3).

Sender publishes tx^{er} after starting fast track Assume a malicious sender started the fast track with its neighbor, but the fast track updates have not yet reached the receiver. Should the sender now publish tx^{er} , the intermediaries that did not yet perform the fast track will refund. The receiver will say that it did not receive the money and will not ship the promised product. The sender cannot prove that the receiver got the money, even though it has the payment confirmation in form of the receiver's signature of tx^{er} . The transaction tx^{er} on the blockchain is a proof of revocation, and the sender will have lost its money without getting anything in return. The sender should thus not publish tx^{er} after starting the fast track.

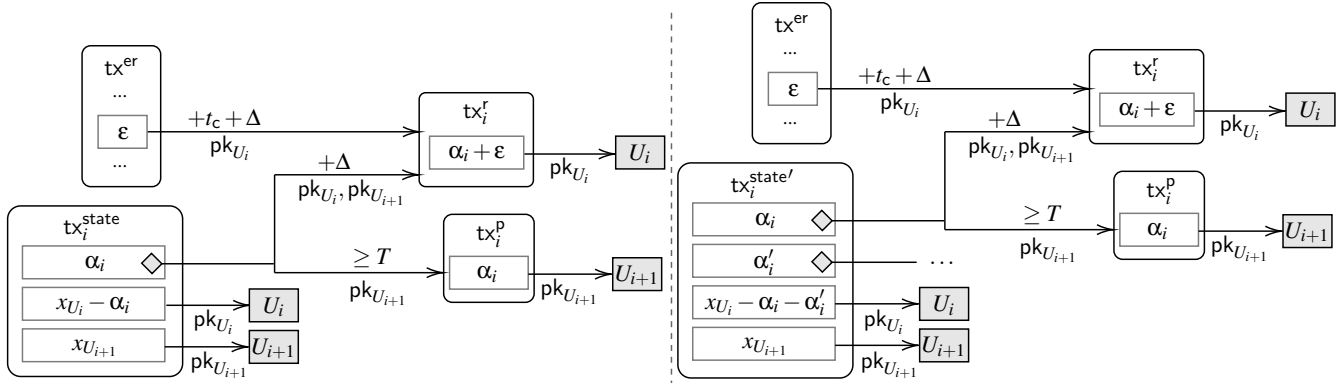


Figure 9: Concurrent payments between users U_i and U_{i+1} : (left) a Blitz channel with a single payment; (right) an updated channel that has this payment and a second concurrent one. To add a second payment of value α'_i to the channel, the transactions for the in-flight payment of value α_i are recreated with the new state $tx_i^{state'}$ as input, the channel is updated to $tx_i^{state'}$ and finally, the old state tx_i^{state} is revoked. In the LN, this process is the same, except that the HTLC contract and transactions are recreated, instead of the Blitz ones.

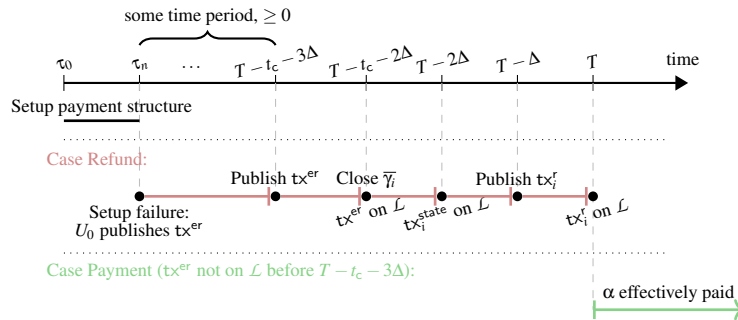


Figure 10: Timeline of when transactions appear on the ledger \mathcal{L} in the case payment and refund. $\tau_n - \tau_0$ denotes the time needed for the setup of the whole payment.

E Timeline

We show a timeline of posting the transaction of the Blitz payment construction between two users in Figure 10. Red shows the refund case, green the payment case.

F Communication overhead

To evaluate our payment scheme, we created an implementation that creates the transactions necessary for setting up the payment. The source code is publicly available at <https://github.com/blitz-payments/overhead>. We tested the compatibility by deploying the transactions on the Bitcoin testnet and checking if the transactions achieve our intended functionalities. Furthermore, we measured the transaction sizes in Bytes and compare them to multi-hop payments in the Lightning Network (LN) in a case-by-case analysis.

We present the number of transactions and their sizes for the different sizes in Table 3. Note that the size of the contract in our construction is only 88 Bytes compared to the 119 of the HTLC, a difference mostly due to the part of the script that verifies the hash pre-image. This means, that state transactions

Table 3: Communication overhead of the LN and Blitz. The pessimistic transactions are on-chain, the rest off-chain.

Cases	LN		Blitz	
	# txs	size	# txs	size
Pay (pessimistic)	1	192	1	158
Refund (pessimistic) per channel	1	158	1	307
Additional pess. refund cost for sender	0	0	1	$157 + 34 \cdot n$
Cost of p in-flight payments	1	$225 + 119 \cdot p$	1	$225 + 88 \cdot p$

holding several different in-flight payments, which directly implement the contract in their outputs, can hold around 26% more Blitz payments than LN payments. For one payment, this difference results in a state of size 311 Bytes for Blitz and a state of 345 Bytes for the LN. In Blitz, additionally to the state we require the refund transaction to be exchanged, which is 307 Bytes, resulting in 618 Bytes for a 2-party setup.

For the rest of the cases, the Blitz payments and the LN payments are similar. In the pessimistic case, both Blitz and the LN require to publish one transaction (after closing the channel) per disputed channel. In the pessimistic refund case, the it is 158 Bytes in the LN and 307 Bytes in Blitz, due to the

Table 4: Extended results of our simulation.

ub	FRate	ppnpr	fail _{Blitz}	fail _{LN}	ratio
25% disrupted type 1, 75% type 2					
3000	0.5%	4	4	33	8.25
3000	0.5%	50	13	4343	334.08
3000	1%	4	15	56	3.73
3000	1%	50	751	32807	43.68
3000	2.5%	4	28	182	6.50
3000	2.5%	50	1076	77213	71.76
75% disrupted type 1, 25% type 2					
3000	0.5%	4	18	31	1.72
3000	0.5%	50	505	4422	8.76
3000	1%	4	19	61	3.21
3000	1%	50	1458	33386	22.90
3000	2.5%	4	78	195	2.50
3000	2.5%	50	15427	77574	5.03

additional signature of the input spending from tx^{er} . In the pay case it is 192 Bytes in the LN and 158 Bytes in Blitz, due to the additional hash in the LN. The most notable difference in comparing the transaction overhead comes from the fact that in the Blitz payment, the sender has to publish tx^{er} in the pessimistic refund case, which is a total of $157 + 34 \cdot (n)$ Bytes, for a payment path of length $n + 1$. However, in the LN there is an additional communication overhead of sending the hash pre-image of 32 Bytes per channel back in the open phase.

G Extended simulation results

In this section, we include results for the simulation when we do not distribute the disrupted payments equally between the two types. As expected, letting 75% of the disrupted payments be of the second type is more favorable for Blitz, while having 25% is less favoring than dividing equally. We show the results in Table 4.

H Extended macros

In this section, we give concrete pseudo-code for the used subprocedures.

Subprocedures
<p><u>checkTxIn($\text{tx}^{\text{in}}, n, U_0$):</u></p> <ol style="list-style-type: none"> 1. Check that tx^{in} is a transaction on the ledger \mathcal{L}. 2. If $\text{tx}^{\text{in}}.\text{output}[0].\text{cash} \geq n \cdot \epsilon$ and $\text{tx}^{\text{in}}.\text{output}[0].\phi = \text{OneSig}(U'_0)$, that is spendable by an unused address of U_0, return \top. Otherwise, return \perp. When using this transaction (to fund tx^{er}), the sender will pay any superfluous coins back to a fresh address of itself. <p><u>checkChannels($\text{channelList}, U_0$):</u></p> <p>Check that channelList forms a valid path from U_0 via some intermediaries to a receiver U_n and that no users are in the path twice. If not, return \perp. Else, return U_n.</p>

checkT(n, T):

Let τ be the current round. If $T \geq \tau + n(2 + t_u) + 3\Delta + t_c + 1$, return \top . Otherwise, return \perp .

genTxEr($U_0, \text{channelList}, \text{tx}^{\text{in}}$):

1. Let $\text{outputList} := \emptyset$ and $\text{rList} := \emptyset$
2. For every channel γ_i in channelList :
 - $(\text{pk}_{\tilde{U}_i}, R_i) \leftarrow \text{GenPk}(\gamma_i.\text{left}.A, \gamma_i.\text{left}.B)$
 - $\text{outputList} := \text{outputList} \cup (\epsilon, \text{OneSig}(\text{pk}_{\tilde{U}_i}) \wedge \text{RelTime}(t_c + \Delta))$
 - $\text{rList} := \text{rList} \cup R_i$
3. Let $\mathcal{P} := \{\gamma_i.\text{left}, \gamma_i.\text{right}\}_{\gamma_i \in \text{channelList}}$ and let nodeList be a list, where \mathcal{P} is sorted from sender to receiver. Let $n := |\mathcal{P}|$.
4. Shuffle outputList and rList .
5. Let $\text{tx}^{\text{er}} := (\text{tx}^{\text{in}}.\text{output}[0], \text{outputList})$
6. Create a list $[\text{msg}_i]_{i \in [0, n]}$, where $\text{msg}_i := \mathcal{H}(\text{tx}^{\text{er}})$
7. $\text{onion} \leftarrow \text{CreateRoutingInfo}(\text{nodeList}, [\text{msg}_i]_{i \in [0, n]})$
8. Return $(\text{tx}^{\text{er}}, \text{rList}, \text{onion})$

genState($\alpha_i, T, \tilde{\gamma}_i$):

1. For the users $U_i := \tilde{\gamma}_i.\text{left}$ and $U_{i+1} := \tilde{\gamma}_i.\text{right}$, create the output vector $\tilde{\theta}_i := (\theta_0, \theta_1, \theta_2)$, where
 - $\theta_0 := (\alpha_i, (\text{MultiSig}(U_i, U_{i+1}) \wedge \text{RelTime}(T)) \vee (\text{OneSig}(U_{i+1}) \wedge \text{AbsTime}(T)))$
 - $\theta_1 := (x_{U_i} - \alpha_i, \text{OneSig}(U_i))$
 - $\theta_2 := (x_{U_{i+1}}, \text{OneSig}(U_{i+1}))$

where x_{U_i} and $x_{U_{i+1}}$ is the amount held by U_i and U_{i+1} in the channel, respectively.
2. Let $\text{tx}_i^{\text{state}}$ be a channel transaction carrying the state with $\text{tx}_i^{\text{state}}.\text{output} = \tilde{\theta}_i$. Return $\text{tx}_i^{\text{state}}$.

checkTxEr($U_i, a, b, \text{tx}^{\text{er}}, \text{rList}, \text{onion}_i$):

1. $x := \text{GetRoutingInfo}(\text{onion}_i, U_i)$. If $x = \perp$, return \perp . If U_i is the receiver and $x = \mathcal{H}(\text{tx}^{\text{er}})$, return $(\top, \top, \top, \top, \top)$. Else, if $x \neq (U_{i+1}, \mathcal{H}(\text{tx}^{\text{er}}), \text{onion}_{i+1})$, return \perp .
2. For all outputs $(\text{cash}, \phi) \in \text{tx}^{\text{er}}.\text{output}$ it must hold that:
 - $\text{cash} = \epsilon$
 - $\phi = \text{OneSig}(\text{pk}_x) \wedge \text{RelTime}(t_c + \Delta)$ for some identity pk_x
3. For exactly one output $\theta_{\epsilon_i} := (\epsilon, \text{OneSig}(\tilde{U}_i) \wedge \text{RelTime}(t_c + \Delta)) \in \text{tx}^{\text{er}}.\text{output}$ and one element $R_i \in \text{rList}$ it must hold that
 - Let $\text{pk}_{\tilde{U}_i}$ be the corresponding public key of $\text{OneSig}(\tilde{U}_i)$
 - $\text{sk}_{\tilde{U}_i} := \text{GenSk}(a, b, \text{pk}_{\tilde{U}_i}, R_i)$ must be the corresponding secret key of $\text{pk}_{\tilde{U}_i}$
4. If the checks in 2 or 3 do not hold, return \perp
5. Return $(\text{sk}_{\tilde{U}_i}, \theta_{\epsilon_i}, R_i, U_{i+1}, \text{onion}_{i+1})$