

# Breaking and Fixing Virtual Channels: Domino Attack and Donner

Lukas Aumayr

TU Wien

lukas.aumayr@tuwien.ac.at

Pedro Moreno-Sanchez

IMDEA Software Institute

pedro.moreno@imdea.org

Aniket Kate

Purdue University / Supra

aniket@purdue.edu

Matteo Maffei

Christian Doppler Laboratory  
Blockchain Technologies for the  
Internet of Things / TU Wien  
matteo.maffei@tuwien.ac.at

**Abstract**—Payment channel networks (PCNs) mitigate the scalability issues of current decentralized cryptocurrencies. They allow for arbitrarily many payments between users connected through a path of intermediate payment channels, while requiring interacting with the blockchain only to open and close the channels. Unfortunately, PCNs are (i) tailored to payments, excluding more complex smart contract functionalities, such as the oracle-enabling Discreet Log Contracts and (ii) their need for active participation from intermediaries may make payments unreliable, slower, expensive, and privacy-invasive. Virtual channels are among the most promising techniques to mitigate these issues, allowing two endpoints of a path to create a direct channel over the intermediaries without any interaction with the blockchain. After such a virtual channel is constructed, (i) the endpoints can use this direct channel for applications other than payments and (ii) the intermediaries are no longer involved in updates.

In this work, we first introduce the Domino attack, a new DoS/griefing style attack that leverages virtual channels to destruct the PCN itself and is inherent to the design adopted by the existing Bitcoin-compatible virtual channels. We then demonstrate its severity by a quantitative analysis on a snapshot of the Lightning Network (LN), the most widely deployed PCN at present. We finally discuss other serious drawbacks of existing virtual channel designs, such as the support for only a single intermediary, a latency and blockchain overhead linear in the path length, or a non-constant storage overhead per user.

We then present Donner, the first virtual channel construction that overcomes the shortcomings above, by relying on a novel design paradigm. We formally define and prove security and privacy properties in the Universal Composability framework. Our evaluation shows that Donner is efficient, reduces the on-chain number of transactions for disputes from linear in the path length to a single one, which is the key to prevent Domino attacks, and reduces the storage overhead from logarithmic in the path length to constant. Donner is Bitcoin-compatible and can be easily integrated in the LN.

## I. INTRODUCTION

Payment channels (PCs) have emerged as one of the most promising solutions to the limited transaction throughput of permissionless blockchains, with the Lightning Network [32] being the most popular realization thereof in Bitcoin. A PC

enables arbitrarily many payments between two users while requiring to commit only two transactions to the ledger: one to open and another to close the channel. Aside from payments, several applications proposed so far benefit from the scalability gains of 2-party PCs [11], [12], [18]. Recent work [8] has further shown how to lift any operation supported by the underlying blockchain to the off-chain setting, thereby further expanding the class of supported off-chain applications.

Creating PCs between all pairs of users (i.e., a clique) is economically infeasible, as users must lock coins for each PC and funding occurs on-chain. On-demand creation of PCs with any potential partner is also infeasible due to the need for on-chain transactions for opening and closing each channel, which results in on-chain fees, long confirmation times (around 1h in Bitcoin) and again impacts the blockchain throughput. As a result, single PCs are instead linked together to form PCNs, using paths of PCs to connect two users instead of opening a PC between them. The interactions of PCN users can be classified into synchronization protocols and virtual channels.

**Synchronization protocols.** Synchronization protocols [9], [22], [28]–[30], [32] allow a sender to pay a receiver when they are connected by a path of PCs, atomically updating the balance of all PCs along the path. Although some of these synchronization protocols are deployed in practice (e.g., for multi-hop payments in the Lightning Network), there are several drawbacks: (i, online assumption) they require users in the path to be online; (ii, reliability) each intermediate user must participate, making the payment less reliable; (iii, cost) each intermediate charges a fee per synchronization round; (iv, latency) the latency of the application increases along with the number of intermediaries (e.g., in the Lightning Network up to one day latency per channel); (v, privacy) intermediaries are aware of every single operation; and (vi, efficiency) they can handle only a limited number of simultaneous payments (e.g., 483 in the Lightning Network) [4]. Finally, and perhaps more importantly, current synchronization protocols are tailored to payments. Supporting 2-party applications (as the ones mentioned before) would require thus to come up with a synchronization protocol for each application. Apart from being a burden, it is not trivial to design such protocols tailored to applications beyond payments, as exemplified by the recent quest in the Bitcoin community about the realization of Discreet Log Contracts across multiple hops [17].

**Virtual channels.** Virtual channels (VC) [7], [19]–[21], [25], [27], [31] allow two users connected by a path of PCs to establish a direct connection, bypassing intermediaries. Intuitively, a

VC is akin to a PC, but instead of being opened by an on-chain transaction, it is opened off-chain using funds from the path of PCs. Therefore, the opening phase involves all intermediaries, besides the endpoints. Once established, however, updates can proceed without the involvement of any intermediaries. In this manner, VCs overcome the aforementioned drawbacks of synchronization protocols: (i) intermediaries are no longer required to be online; (ii) the reliability of the channel does not depend on intermediaries; (iii) intermediaries do not charge a fee for each usage of the channel (perhaps only once to create and close the VC); (iv) the latency does not depend on intermediaries; (v) intermediaries do not learn each single VC update; (vi) a PC can host several VCs, each of which can be used to dispense up to 483 payments or potentially more VCs, bypassing the limitation on the number of payments in PCNs.

Since VCs can be used just as PCs, they constitute the most promising solution to perform repeated transactions as well as applications different from payments (e.g., [11], [12], [18]) between any pair of users connected by a path of PCs. In fact, applications built on top of PCs can be smoothly lifted to VCs, which constitutes a crucial improvement over synchronization protocols.<sup>1</sup> For instance, VCs support Discreet Log Contracts [18], an application that has received increased attention lately and that intuitively allows for bets based on attestations from an oracle on real world events. As compared to PCs, VCs offer the same advantages while requiring no on-chain transaction for their setup, thereby dispensing from the associated blockchain delays, on-chain fees, and on-chain footprints. This makes it possible to keep VCs short-lived, to frequently close, open, or extend them based on current needs. For a more detailed discussion see Appendix A.

VC constructions are difficult to design, since the balance of honest parties needs to be ensured even in the presence of malicious, and possibly colluding, intermediaries/endpoints. The first constructions have been proposed for blockchains supporting Turing-complete scripting languages based on the account model, like Ethereum [19]–[21]. In such blockchains, VC constructions are somewhat easier to design: For instance, stateful smart contracts can resolve conflicts on the current state of VCs by associating a different version number to each state update and, in case of conflict, by selecting the highest number as the valid state. Indeed, Ethereum-based constructions are based on this idea and do not suffer from the Domino attack presented in this paper. Unfortunately, this reliance on Turing-complete scripting languages makes these constructions incompatible with many of the cryptocurrencies available today, including Bitcoin itself.

It is not only of practically relevant, but also theoretically interesting to investigate what is the minimum scripting functionalities necessary to design secure VCs. Therefore, a bit later VC constructions have been proposed also for blockchains with a less expressive scripting language and based on the Unspent Transaction Output (UTXO) model (i.e., Bitcoin-compatible) [7], [25], [27]. Throughout the rest of this paper, we investigate VCs built on these blockchains if not specified otherwise. All of these VC constructions share one common design pattern: The VC is funded from all underlying PCs. We refer to this design pattern as rooted VCs and illustrate

TABLE I: Comparison to other multi-hop VC protocols.

|                              | LVPC [25]<br>Bitcoin     | Elmo [27]<br>Bitcoin +<br>ANYPREVOUT | Donner<br>Bitcoin |
|------------------------------|--------------------------|--------------------------------------|-------------------|
| Scripting requirements       |                          |                                      |                   |
| Multi-hop                    | ✓                        | ✓                                    | ✓                 |
| Secure against Domino attack | ✗                        | ✗                                    | ✓                 |
| Path privacy                 | ✗                        | ✗                                    | yes               |
| Time-based fee model         | ✓                        | ✗                                    | ✓                 |
| Unlimited lifetime           | ✗                        | ✓                                    | ✓                 |
| Storage Overhead per party   | $\Theta(n)$ <sup>‡</sup> | $\Theta(n^3)$                        | $\Theta(1)$       |
| Off-chain closing            | ✓                        | ✗                                    | ✓                 |
| Offload: txs on-chain        | $\Theta(n)$              | $\Theta(n)$                          | 1                 |
| Offload: time delay          | $\Theta(n)$ <sup>‡</sup> | $\Theta(n)$ <sup>‡</sup>             | 1                 |

<sup>‡</sup> by synchronizing all channels, this time can be only  $\Theta(\log(n))$ .

it on a high level in Figure 1(a.1). Because VCs are, unlike PCs, not funded on-chain, they rely on an operation called offloading, which transforms a VC to a PC. This is important for honest users so they can enforce their balance in case the other user misbehaves: first transforming the VC to a PC by putting the VC funding on-chain, and second using the means provided by the PC to enforce their balance. Rooted designs enable both endpoints to offload the VC, but because they are funded by all underlying PCs, every underlying PC has to be closed on-chain (see Figure 1(a.2)).

Conceptual advancements in this work. We show that rooted VCs are by design prone to severe drawbacks including the Domino attack (see Section III), a new DoS/griefing style attack in which (i) a malicious intermediary of a VC or (ii) an attacker establishing a VC with itself over a number of honest PCs can close the whole path of underlying PCs and bring them on-chain. Not only are all existing Bitcoin-compatible VC constructions affected by this attack, in fact the ideal functionalities against which they are proven secure do permit this attack, but also this attack is so severe that it can potentially shut down the underlying PCN, as we show in Section III-C. As a result, we argue that none of the existing Bitcoin-compatible VC constructions should be deployed in practice. Furthermore, the rooted design allows adversaries to learn the identity of participants other than their direct neighbors, thereby breaking what we call path privacy (see Section III-D). Given these security and privacy shortcomings, we introduce a paradigm shift towards the design of non-rooted VCs, based on two fundamental ingredients.

First, instead of being rooted, the VC is funded independently from the underlying PCs, by one of the VC endpoints. The underlying PCs are used to lock up some funds (or collateral) that are paid to the honest VC endpoint if the other VC endpoint misbehaves. We illustrate this concept on a high level in Figure 1(b.1). In contrast to rooted designs, VCs can be offloaded without having to close the underlying PCs, which is the key to prevent Domino attacks. Since the VC is only funded by one endpoint, only this funding endpoint has the means of transforming the VC to a PC (offload). Subsequently, the other one cannot get their money via offloading in case of misbehavior. This issue is solved by compensating the non-funding endpoint in case the funding endpoint has not transformed the VC to a PC within a channel lifetime  $T$ , see Figure 1(b.2) and (b.3).

This lifetime  $T$  is the second crucial aspect where we depart from the state of the art. Current solutions provide unlimited lifetime without guaranteeing however that the VC

<sup>1</sup> VCs expose all the functionalities of a PC and can be used interchangeably as a building block for off-chain applications, see Section V.




Fig. 1: Conceptual comparison of (a) state-of-the-art VCs (rooted) and (b) our protocol (non-rooted) on layers L1 (blockchain), L2 (PCs) and L3 (VCs). Note that the VC in (a.1) is funded by all the underlying channels. In (b.1), the VC is funded only by  $U_0$ , indirectly via a transaction  $tx^{vc}$ . Additionally, in (b.1), a payment is set up from  $U_0$  to  $U_3$ , whose outcome depends on whether the VC is offloaded. Offloading, i.e., the act of forcefully transforming a VC (L3) to a PC (L2) in (a.2), requires that all the underlying PCs (L2) are put on-chain (L1). In (b.2), offloading the VC keeps the PCs open, posting only  $tx^{vc}$  on-chain (L1). Since offloading enables  $U_3$  to receive their funds, the payment is refunded then. However, since in (b), only  $U_0$  can offload,  $U_3$  is compensated (b.3) after a timeout  $T$  via a payment that is executed iff  $U_0$  has not offloaded the VC (i.e., (b.2) did not happen).

will remain open, as an intermediary node could initiate the offloading. Instead, our design ensures that the VC is open until time  $T$ , which can be repeatedly prolonged if all involved parties agree. This allows intermediaries to charge fees based on the lifetime of the VC, which corresponds to the time they have to lock up their funds, something that is not possible in current VC solutions with unlimited lifetime [27]. The improvements over existing Bitcoin-compatible multi-hop VC constructions are summarized in Table I. We compare with single-hop constructions and with those relying on Turing-complete smart contracts in Table IV in Appendix B.

Our contributions can be summarized as follows:

- We introduce the Domino attack, which allows the adversary to close arbitrarily many PCs of honest users, thereby destructing the underlying PCN. We argue that any rooted construction, in particular, all existing Bitcoin-compatible VC constructions are prone to this attack. We show the severity of this attack in a quantitative analysis; given current BTC transaction fees, it suffices for an attacker to spend 1 BTC to close every channel in the current LN. Even though VC protocols are not yet used in practice, we find it crucial to show this attack before any construction gets implemented, offering instead a secure alternative.
- We present Donner, a new VC protocol that departs from the rooted paradigm by funding the VC from outside of the underlying PC path. In addition to being secure against the Domino attack, it significantly improves in terms of efficiency and interoperability over state-of-the-art VC protocols (see Table I).
- We introduce the notion of synchronized modification, a novel subroutine allowing parties to atomically change the

value or timeout of a synchronization protocol, a contribution of independent interest. Synchronized modification, non-rooted funding, and the pay-or-revoke paradigm [9] are the core building blocks of Donner.

- We conduct a formal security and privacy analysis of Donner in the Universal Composability framework.
- We conduct experimental evaluations to quantify the severity of the Domino attack and demonstrate that Donner requires significantly less transactions than state-of-the-art VCs; Donner decreases the on-chain costs for offloading VCs from linear in the path length to a single one and the storage overhead per PC from linear or logarithmic in LVPC [25] (depending on how the VC is constructed) or cubic in Elmo [27] to constant.

## II. BACKGROUND AND NOTATION

### A. UTXO based blockchains

We adopt the notation for UTXO-based blockchains from [8], which we shortly review next. In UTXO-based blockchains, the units of currency, i.e., the coins, exist in outputs of transactions. We define such an output as a tuple  $\theta := (\text{cash}, \varphi)$ ;  $\theta.\text{cash}$  contains the amount of coins stored in this output and  $\theta.\varphi$  defines the condition under which the coins can be spent. The latter is done by encoding such a condition in the scripting language of the underlying blockchain. This can range from simple ownership, specifying which public key can spend the output, to more complex conditions (e.g., timelocks, multi-signatures, or logical boolean functions).

Coins can be spent with transactions, resulting in the change of ownership of the coins. A transaction maps a list of outputs to a list of new outputs. For better readability, we denote the former outputs as transaction inputs. Formally, we define a transaction body as a tuple  $tx := (\text{id}, \text{input}, \text{output})$ . The identifier  $tx.\text{id} \in \{0, 1\}^{\mathbb{D}}$  is assigned as the hash of the other attributes,  $tx.\text{id} := H(tx.\text{input}, tx.\text{output})$ . We model  $H$  as a random oracle. The attribute  $tx.\text{input}$  is a non-empty list of the identifiers of the transaction's inputs and  $tx.\text{output} := (\theta_1, \dots, \theta_n)$  a non-empty list of new outputs. To prove that the spending conditions of the inputs are known, we introduce full transactions, which contain in addition to the transaction body also a witness list. We define a full transaction  $\bar{tx} := (\text{id}, \text{input}, \text{output}, \text{witness})$  or for convenience also  $\bar{tx} := (tx, \text{witness})$ . Valid transactions can be recorded on the public ledger  $L$  called blockchain, with a delay of  $\Delta$ . A transaction is valid if and only if (i) all its inputs exist and are not spent by other transaction on  $L$ ; (ii) it provides a valid witness for the spending condition  $\varphi$  of every input; and (iii) the sum of coins in the outputs is equal (or smaller) than the sum of coins in the inputs.

There are several conditions under which coins can be spent. Usually they consist of a signature that verifies w.r.t. one or more public keys, which we denote as  $\text{OneSig}(pk)$  or  $\text{MultiSig}(pk_1, pk_2, \dots)$ . Additional conditions could be any script supported by the scripting language of the underlying blockchain, but in this paper we only use relative and absolute time-locks. For the former, we write  $\text{RelTime}(t)$  or simply  $+t$ , which signifies that the output can be spent only if at least  $t$  rounds have passed since the transaction holding this output was accepted on  $L$ . Similarly, we write  $\text{AbsTime}(t)$

or simply  $\geq t$  for absolute time-locks, which means that the transaction can be spent only if the blockchain is at least  $t$  blocks long. A condition can be a disjunction of subconditions  $\varphi = \varphi_1 \vee \dots \vee \varphi_n$ . A conjunction of subconditions is simply written as  $\varphi = \varphi_1 \wedge \dots \wedge \varphi_n$ .

To visualize how transactions are used in a protocol, we use transaction charts. The charts are to be read from left to right. Rounded rectangles represent transactions, with incoming arrows being their inputs. The boxes within the transactions are the outputs and the value in them represents the amount of output coins. Outgoing arrows show how outputs can be spent. Transactions that are on-chain have a double border (see, e.g., Figure 9 in Appendix C.1).

## B. Payment channels

Two users can utilize a payment channel (PC) in order to perform arbitrarily many payments, while putting only two transactions on the ledger. On a high level, there are three operations in a PC operation: open, update and close. First, to open a channel, both users have to lock up some money in a shared output (i.e., an output that is spendable if both users give their signature) in a transaction called the funding transaction or  $tx^f$ . From this output, they can create new transactions called state or  $tx^s$  which assign each of them a balance. Once the funding transaction is on the ledger, the users can exchange arbitrarily many new states (balance updates) in an off-chain manner, thereby realizing the update phase of the channel. Once they are done, they can close the channel by posting the final state to the ledger.

In this work, we use PCs in a black-box manner and refer the reader to [8], [28], [29] for more details. We abstract away from the implementation details and instead model the state of the channel as the outputs contained in a transaction  $tx^s$ , which is kept off-chain. For simplicity, we assume that this is the only state that the users can publish and abstract away from how the dishonest behavior is handled. In practice, it is possible that a dishonest user publishes a stale state of the channel and current constructions come with a way to handle this case (e.g., through a punishment mechanism that compensates the honest user [8]). We illustrate this abstraction in Figure 2.

## C. Payment channel networks

A payment-channel network (PCN) [28] is a graph where the nodes represent the users and the edges represent the PCs. The Lightning Network [32] is the state of the art in both PCs and PCNs for Bitcoin, and the largest PCN in terms of coins locked within its channel fundings, currently having around 81k channels, 19k active nodes and a total capacity of 3k BTC (around 130M USD).




Fig. 2: We abstract PCs using a squiggly line to hide details that are not needed in this work.  $P : x_P$  indicates that user  $P$  owns  $x_P$  coins in the state  $tx^s$ , written as  $(A, B)$ . The box containing  $x_A + x_B$  indicates the shared output of  $A$  and  $B$ .

In a PCN, any two users connected by a path of channels can perform what is called a multi-hop payment (MHP). Assume that there is a sender  $U_0$  who wants to pay  $\alpha$  coins to a receiver  $U_n$ , but they do not have a direct channel. Instead, they are connected by a path of channels going through intermediaries  $\{U_i\}_{i \in [1, n-1]}$ , such that any pair of neighbors  $U_j$  and  $U_{j+1}$  have a channel  $\gamma_j$ , for  $j \in [0, n-1]$ . A mechanism synchronizing all channels on the path is required for a payment, such that each channel is updated to represent the fact that  $\alpha$  coins moved from left to right. We give an example in Figure 10 in Appendix C.2.

## D. Blitz

There exist many different MHP protocols that synchronize the updates of channels. In particular, the Blitz [9] protocol is useful for this work. In Blitz, the PC updates are dependent on a transaction called  $tx^{er}$ , which acts as a global event. The PCs are synchronized in the following way: If  $tx^{er}$  is posted on-chain, the updates are reverted, otherwise, they are successful. In other words, the sender sets up a MHP conditioned on a “refund enabling” transaction  $tx^{er}$  in a way that the refund can be triggered, if anything goes wrong. If all channels participated honestly, the sender does not post  $tx^{er}$  and the MHP goes through (see Figure 3). In a bit more detail, Blitz consists of four operations:

- 1) Setup. The sender  $U_0$  creates a synchronization transaction  $tx^{er}$  as depicted in Figure 3b, which has an output  $\theta_{\epsilon_i}$  holding  $\epsilon$  coins for each user except the receiver  $U_n$ . The value  $\epsilon$  is set to the smallest possible value that the underlying blockchain allows (ideally zero); these outputs are merely to enable other transactions.
- 2) Open. Each channel sequentially, from sender to receiver, sets up a payment whose success or refund is conditioned on a time  $T$  or transaction  $tx^{er}$ , as conceptualized in Figure 3a and shown in detail in Figure 3c. In a nutshell, two users  $U_i, U_{i+1}$  update their channel  $\gamma_i$  to a state where the amount to be paid  $\alpha$  (more precisely  $\alpha_i$  which encodes a per-hop fee) coming from  $U_i$  can be spent as follows: Either by  $U_{i+1}$  using  $tx_i^p$  after time  $T$  or by  $U_i$  using  $tx_i^r$  if  $tx^{er}$  is posted on-chain. Since each  $tx_i^r$  uses the corresponding output  $\theta_{\epsilon_i}$  of  $tx^{er}$ , the UTXO model ensures that  $tx_i^r$  can only be posted if  $tx^{er}$  has been posted before.
- 3) Finalize. After the receiver has successfully set up the payment, she sends back a confirmation to the sender containing  $tx^{er}$ . If the sender receives a confirmation containing the  $tx^{er}$  she created in the setup phase within some time, she goes idle. Otherwise, she posts  $tx^{er}$ , initiating the refunds (see respond).
- 4) Respond. Every user  $U_i$  monitors the blockchain if  $tx^{er}$  appears. In case it appears before  $T$ , the user will publish the refund transaction  $tx_i^r$  for her channel  $\gamma_i$ . If the two users in  $\gamma_i$  collaborate, both updates and refunds can always be performed off-chain.

In this work, we utilize a slightly modified version of Blitz as a building block. We mark the modification in green in Figure 3b and describe it in Section V-C.

## E. State-of-the-art virtual channels

A virtual channel (VC) allows two users to establish a direct channel, without putting any transaction on-chain.




Fig. 3: (3a) Illustration of the Blitz synchronization protocol; (3b) Off-chain synchronization transaction spending from an output under  $U_0$ 's control and linking to the collateral in each channel. (i) Without the green part:  $tx^er$  in Blitz. (ii) With the green part:  $tx^vc$  used for funding the VC in this work; (3c) Two-party contract used within each channel

Indeed, the fundamental difference between a PC and a VC is that in a VC, the funding transaction  $tx^f$  does not go on-chain in the honest case. To still ensure that users do not lose their funds in case of dispute, this requires a new operation: In addition to the three operations open, update and close of PCs, we need the operation offload, which allows a user of the VC to put the funding transaction  $tx^f$  on-chain, transforming the VC into a PC in case of a dispute.

To understand how VCs work, let us look at an example following a state-of-the-art VC construction [25]. This example is depicted in Figure 4. Assume  $U_0$  and  $U_2$  want to construct a VC via  $U_1$ , i.e., there exist PCs  $(U_0, U_1)$  and  $(U_1, U_2)$ , and they wish to build a VC  $(U_0, U_2)$ . To open a VC, the main idea is to take the desired VC capacity  $\alpha$  and lock it in both channels, such that  $\alpha$  coins come from  $U_0$  and  $\alpha$  coins from the intermediary  $U_1$ . These  $2 \cdot \alpha$  coins are used both for funding the VC and as collateral; these coins can be spent in the following, mutually exclusive ways:

- (i) by putting the funding transaction  $tx^f$  on-chain, which simultaneously funds the VC and refunds the intermediary its collateral  $\alpha$ , or
- (ii) if both  $\alpha$  coins are not spent by a chosen punishment time  $t_{pun}$ ,  $U_0$  and  $U_2$  can each claim  $\alpha$  coins, which is the maximal amount they could hold in the VC

Clearly,  $U_1$ , who is part of both channels, is incentivized to




Fig. 4: Illustration of a VC construction over a single intermediary. The VC funding  $tx^f$  is rooted in the underlying channels is the only way for the intermediary to get its collateral back.  $tx^f$  and the punishment are mutually exclusive.




Fig. 5: Illustration of a rooted VC via multiple hops. The yellow lines indicate how the VC is rooted. All transactions connected to and to the left of  $tx^f$  need to be put on-chain in the case the rightmost VC is offloaded.

put  $tx^f$  on-chain, as this is the only way to get her collateral back. Simultaneously, the two end-users  $U_0$  and  $U_2$ , who are only part of one of the channels, are ensured that either  $tx^f$  goes on-chain, or else they receive the full  $\alpha$ .

Putting  $tx^f$  on-chain is called offloading and is a safety mechanism to ensure that users can claim their rightful balance in case of a dispute. Offloading can be initiated by either  $U_0$  or  $U_2$  (by closing their respective channel and threatening to take the collateral if  $U_1$  does not react), or by  $U_1$  by simply closing both channels. We emphasize that the money of  $tx^f$  comes from both underlying channels, i.e., it can only exist on-chain, if both underlying channels have been put on-chain (closed). We call this design a rooted VC.

If there is no dispute, the transactions depicted in Figure 4 remain off-chain and the underlying channels  $(U_0, U_1)$  and  $(U_1, U_2)$  remain open. The update of the VC requires no interaction of the intermediaries, the end-users simply update the channel  $(U_0, U_2)$  as they would a PC. Finally, to close the VC, the final balance of the VC has to be mapped into the base channels so that in the end both VC endpoints receive the latest balance of the VC and the intermediaries do not lose coins. Note that with the exception of offload, which requires at least one on-chain transaction (i.e., the funding), all other operations require no on-chain transaction. This single-intermediary idea can be used to construct a tree-like structure over a path of arbitrary intermediaries to get VCs of arbitrary length. We show this concept in Figure 5.

### III. THE DOMINO ATTACK

#### A. Reasons that lead to the attack

**Observation 1:** Balance security for VC endpoints. Independently of its inner workings, any VC construction must ensure that honest VC endpoints Alice and Bob can cash out the coins they hold in the VC (i.e., get their coins on-chain). As discussed in Section II-E, VCs are akin to payment channels (PCs), with the difference of having their funding transaction off-chain. This means that both endpoints can no longer directly claim their latest balance as in a PC. Instead, the VC funding transaction first needs to be put on-chain through the operation offload, which can be initiated by the VC endpoints and in some existing VC protocols [27] even by the intermediaries.

**Observation 2:** VC funding transaction is rooted in all underlying base channels. We recall that to enable the offload operation, the VC funding takes as inputs (either directly or indirectly, via intermediate transactions) outputs of each of the underlying base channels. We denote such a VC as being rooted in the base channels.<sup>2</sup> At a first glance, this seems the most natural approach since it allows both endpoints to offload the VC and the intermediaries to unlock their collateral. However, a rooted funding implies that it can be posted on-chain if and only if all underlying PCs are closed. This feature is the source of the Domino attack, as shown next.

#### B. Attack description

The Domino attack is essentially a DoS or griefing style attack. It follows directly from the two observations mentioned above and can proceed in the following phases: (i) an adversary controlling two nodes opens two PCs encasing a path of victim channels; (ii) the adversary opens a VC to herself via these victim paths; and (iii) she initiates the offloading of the VC.

The effect of this attack is to force the closure of every channel on this path, i.e., the two the attacker created and the channels on the victim path. Anyone not closing their channel risks losing their money. In stark contrast to payment protocols in PCNs such as Lightning or Blitz where closing one channel in the payment path still allows channels in the rest of the path to remain open, in current VC constructions there is no way that honest nodes can settle their channels honestly off-chain and keep them open. They are forced to close every channel, as the VC funding can only exist on-chain if all base channels are closed.

**Example.** Assume an attacker controlling nodes  $U_0$  and  $U_4$  who wants to perform a Domino attack on the victim path  $U_1$ ,  $U_2$  and  $U_3$ , see Figure 5. If not already opened, the attacker opens the channels  $(U_0, U_1)$  and  $(U_3, U_4)$ . Then, she constructs a VC between her own nodes  $U_0$  and  $U_4$  recursively, as, e.g., established in the LVPC protocol [25]. After the attacker is done with this step, the transaction structure among different users is as in Figure 5. The attacker can now unilaterally force the closure of all underlying channels, i.e., the PCs  $(U_0, U_1)$ ,  $(U_1, U_2)$ ,  $(U_2, U_3)$  and  $(U_3, U_4)$  as well as the intermediate VCs  $(U_0, U_2)$ ,  $(U_0, U_3)$  and the offloading of  $(U_0, U_4)$ .

First,  $U_4$  closes the PC  $(U_3, U_4)$ , which she can do on her own. In the rooted VC example of Figure 5 (e.g., this could be LVPC), the output in the state of  $(U_3, U_4)$  which is used to fund the VC  $(U_0, U_4)$  goes to  $U_4$ , unless it is first consumed by the VC. This means that an honest  $U_3$  will lose money in the channel  $(U_3, U_4)$  to  $U_4$  by means of the punishment transaction on the bottom right in Figure 5 (dubbed Punish transaction in the LVPC protocol), unless she closes the channel  $(U_0, U_3)$  and claims its money by posting  $tx^f$ , i.e., the transaction funding the VC (i.e., offloading)  $(U_0, U_4)$ , dubbed Merge transaction in LVPC.

However, to post  $tx^f$  for  $(U_0, U_4)$ ,  $U_3$  first needs close  $(U_0, U_3)$ .  $U_3$  initiates the offloading by first closing  $(U_2, U_3)$ . This triggers a similar response from  $U_2$ , who is now at risk of losing the coins in  $(U_2, U_3)$ , unless she offloads  $(U_0, U_3)$  by putting the corresponding  $tx^f$ . But to do that,  $U_2$  first needs to close  $(U_1, U_2)$ . This is done, finally, by closing  $(U_0, U_1)$  which forces  $U_1$  to close also  $(U_0, U_1)$ .

In the end, all channels are closed. We illustrate this in the extended version of this work [10]. Let us clarify that by closing the underlying channels we mean that at least two transactions per channel have to be put on-chain, one for closing the channel and another one to spend the collateral locked for the VC. Due to the fact that LVPC first splits the channel into two subchannels before using one of them to fund the VC, closing the initial channel simultaneously spawns a new channel (i.e., the remaining subchannel) that has a capacity reduced by the amount put in the collateral funding the VC. The Domino attack works regardless of how the recursion was applied, as well as on Elmo [27]. In LVPC some  $(U_3$  in the example above) and in Elmo all intermediaries can carry out this attack. The Domino attack can also be launched if the attacker controls only one of the endpoints, assuming the other one agrees to open a VC with her over the victim path. We remark that LVPC and Elmo are modelled in the UC framework, however, their ideal functionalities explicitly allow for the Domino attack.

#### C. Quantitative analysis of the Domino attack

To quantify the severity of the Domino attack, we perform the following simulation. We take a current (March 2022) snapshot of the Lightning Network (LN) [5]. In this snapshot, there are 83k channels, 20k nodes and 3284 BTC (around 150M USD) locked in channels (of the largest connected component). The nodes' connectivity varies in the LN. There are leaf nodes having only one open channel, and there are nodes with almost 3000 channels. Additionally, entities can control multiple nodes. The entities can be linked by their alias, as pointed out in [33], something we follow in this simulation as well.

Clearly, differently connected nodes can launch the Domino attack with more or less devastating effect. The better connected a node is, the more channels can be closed down. Note that for this attack, it does not matter how many coins are locked in the channels under control of the attacker and not even the number of nodes the attacker control, but instead the number of open channels and the kind of paths which exist to another node under the attacker's control; the source and destination may be the same node.

<sup>2</sup>By base channel we mean either a PC or a VC that was used for opening a VC, to capture the fact that VCs can be constructed recursively.




Fig. 6: Simulated effect of the Domino attack.

Analyzing existing nodes. To measure the damage that can be caused by existing nodes in the LN with two or more open channels, we do the following. Assuming each node, or more precisely, each alias, is performing the Domino attack. This means, using the open channels the attacker tries to close as many channels as possible. Computing the optimal set of VCs the attacker would need to open to maximize the channels is computationally expensive and out of the scope of this simulation. Instead, we settle for a simpler heuristic. The attacker computes the cycle basis for a root node controlled by the attacker, yielding paths starting and ending at one node under the attacker’s control. The attacker chooses the longest one and proceeds to close the channels by performing the Domino attack. Now, on the new network with fewer channels, the attacker repeats these steps, until all of the attacker’s channels are closed and they can do no further damage.

We count the channels an attacker can close with this approach for each alias. Each node can close 1284 channels on average, which amounts to around 1.5% of all channels in the LN. However, note that around 8% of all nodes can close no channels at all, while the most well-connected entity can close around 53k channels, which more than 60% of the LN. We visualize our results in Figure 6a, where for a given interval of how many channels an entity can close, we show the percentage of nodes that falls into this category. The source code and raw results of this simulation can be found at [6].

To make matters worse, an attacker can target specific channels with this. This allows the attacker to perform attacks similar to Route Hijacking [36], a DoS attack where an attacker strategically places a channel in a topologically important location and announces low fees. Subsequently, users will route their payments through the attacker’s channel who can then drop the requests. In the worst case this can (temporarily) disconnect parts of the network from one another. In the Domino attack, an attacker can disconnect parts of the network directly, by closing all edges that connect the two subgraphs.

Analyzing newly placed nodes. In this second analysis, we let the attacker create new channels instead of assuming an existing node is corrupted. Clearly, without any restrictions, an attacker can do more damage than in the previous simulation, i.e., by opening the same (and more) channels as the the best performing node which had a bit less than 3000 channels. Taking a current average fee of 0.000031 BTC (1.27 USD) per transaction [2], this would cost an adversary around 0.186 BTC. In more detail, 0.093 BTC are needed for opening these channels and again 0.093 BTC for closing them after

establishing the according VC, triggering the Domino attack. Note that the latter amount is also needed if the channels are already there (in the previous simulation).

We therefore put some restrictions on the attacker. We assume that an adversary has a certain budget to spend on fees for establishing channels over the network. Further, the adversary constructs VCs of a length of up to  $n \in [2, 11]$  to herself, i.e., the adversary is the first and last node. We set the maximum VC length  $n$  to 11, the diameter of the LN snapshot, i.e., at this length every nodes can reach every other node.

The adversary needs to post 3 on-chain transactions per VC with the associated fees, two for establishing the two PCs encasing the victim path and one to close one of these channels. Further, for the VC itself, a certain minimum amount is needed to open it, similar to LN payments. However, since this amount is presumably not only very small, but also the adversary gets it back, we omit it in our simulation and say instead that the adversary performs this attack in sequence. Finally, we note that the effect of this attack is likely to be even more severe in reality, since in existing VC constructions, not only does the channel need to be closed, but subsequent transactions making up the rooted funding of the VC need to be posted as well.

We present our results in Figure 6b. Using only 1 BTC for fees, the adversary can close up to 97k honest channels, which is more than all channels in our LN snapshot (83k), and cause a cost of at least 6 BTC to the involved nodes. Budgets in the order of 0.2, 0.5, and 1 BTC are not unrealistic, as there are 1501, 799, and 453 nodes, respectively, holding this money within the LN, assuming equal balance distribution in the channels, i.e., 0.5% of nodes in the LN have enough balance to shut down the whole network. If we consider all Bitcoin addresses (even outside the LN), there exist 815k addresses owning 1 BTC or more [3].

We remark that since VCs are not used in practice, we cannot evaluate this in the real world. However, previous work has already shown the feasibility of similar DoS or griefing attacks and how they transfer to the real world [24]. For a discussion on why it is infeasible to deter this attack with fees, we refer to Appendix B. From our simulation it follows that this attack is too severe for the adaption of current VC solutions in PCNs such as the LN. In order to make VCs usable in practice, it is essential to prevent the Domino attack.

#### D. More drawbacks of current VC constructions

Unlimited lifetime. Existing VC constructions such as Elmo [27] offer VCs with an a priori unlimited lifetime. On a high level, unlimited lifetime of a VC means that if every party agrees (including endpoints and intermediaries), the VC can remain open potentially forever. While existing work highlights unlimited lifetime as a desirable feature for both PCs and VCs, we view it as a drawback in the context of VCs. Indeed, there is an important difference between VCs and PCs: in a VC funds are locked up not only by the endpoints, but also by the intermediaries of the underlying path. Without a lifetime, intermediaries could have their collateral locked up forever, unless they decide to go on-chain, which however forces them to close their PCs. Related to that, intermediaries should charge a fee proportional to the collateral and the time this collateral is

locked (analogously to the LN): without a lifetime, the second parameter cannot be estimated nor enforced without closing the base PCs.

We therefore propose a new approach: instead of having an a priori unlimited lifetime, we fix a certain lifetime at the point of creation. When this lifetime expires, users have the option to prolong it for another fixed lifetime if everyone agrees or to close it. Prolonging it means that the VC remains active and any applications hosted on top of can be kept on being used smoothly. In addition, every intermediary can charge a lifetime-based fee every time they prolong the VC. While all agree, they can repeat this process indefinitely. If one party wants to stop it, the party can unlock their funds without having to close any channel on-chain. We explain this concept in more detail in Section IV.

**Recursiveness.** The last issue we point out comes from how the VC funding is rooted in the underlying channels. In current VC constructions, the VC funding is built by recursively combining two channels at a time, forming a tree with the VC funding transaction being the root of the tree and the underlying channels being the leaves. This has two negative implications. First, in addition to closing all PCs (which requires at least one on-chain transaction per channel), i.e., the leaves of the tree, a linear number of transactions needs to go on-chain in order to offload a channel, i.e., the non-leaf nodes of the tree. Second, depending on how the recursiveness has been applied, the time it takes to offload a VC is also either linear (in case of an unbalanced tree) or logarithmic (in case of a balanced tree) in the number of underlying channels. In our construction, offloading involves only a constant number of on-chain transaction as elaborated in the next section.

**Lack of path privacy.** State-of-the-art VC constructions create the rooted funding by connecting outputs of pairs of channels in a recursive way. However, this requires interaction of some intermediaries with more than their direct neighbors on the path. In our construction, intermediaries on the path only learn about their direct neighbors in the honest case, exactly as in the Lightning Network.

#### IV. DONNER: KEY IDEAS

We describe the core ideas of Donner by assuming that a slight variant of the previously described Blitz construction is used as underlying MHP protocol. As detailed below, our construction is parameterized over it, so that other functionality-equivalent MHP protocols could be deployed instead.

**High level architecture.** Let us assume  $U_0$  and  $U_n$ , connected via  $U_i$  for  $i \in [1, n-1]$ , wish to open a bidirectional VC with capacity  $\alpha$  and time  $T$  fully funded by  $U_0$ . First,  $U_0$  starts with a slightly modified version of the Setup phase of a Blitz payment of  $\alpha$  coins, as explained in Section II-D, let us call it Setup\*. In this modified phase,  $U_0$  proceeds to create a transaction  $tx^{vc}$  as depicted in Figure 3b (this time, including the green part) instead of  $tx^{er}$ .  $tx^{vc}$  takes an input from  $U_0$  and creates an output holding  $\alpha$  coins and like in the Setup phase, an output holding  $\epsilon$  coins for each user except the receiver  $U_n$ . This transaction will serve two purposes: (i) it will be the funding of the VC and (ii) it will be used to synchronize a Blitz payment.

Next,  $U_0$  and  $U_n$  proceed to create the initial state (see Section II-B)  $tx^s$  of the VC using  $tx^{vc}$  as a funding. We emphasize that this process is exactly the same as for a PC, the only difference being that the funding transaction  $tx^{vc}$  has these additional outputs holding  $\epsilon$  and we do not intend to publish  $tx^{vc}$  on-chain. After this step is successful,  $U_0$  initiates the remaining phases of Blitz (Open, Finalize and Respond) using  $tx^{vc}$ . After completion, a Blitz payment of value  $\alpha$  is open between  $U_0$  and  $U_n$  conditioned on  $tx^{vc}$ , i.e., it is refunded if  $tx^{vc}$  is posted and otherwise successful after time  $T$ .

**Intuition security.** At this point, the VC is considered open and can be used exactly like a PC. The careful readers might be wondering why this VC is safe to use. After all, we detached the funding from the underlying PCs and removed the receiver  $U_n$ 's ability to offload the VC. However, the sender  $U_0$  did set up a Blitz payment to  $U_n$  of  $\alpha$  coins, which is the full capacity of the VC. By putting the VC funding inside the synchronization transaction of Blitz, we make the two actions offload the VC and refund the Blitz payment atomic. In other words, if  $U_0$  does not offload,  $U_n$  will automatically receive the full VC capacity via the payment after  $T$ .

**Getting rid of the Domino attack.** We recall the causes for the Domino attack: (i) the VC funding has to be enforceable on-chain by offloading and (ii) the VC funding is rooted in all underlying PCs. To prevent the attack, we got rid of (ii): The funding  $tx^{vc}$  comes solely from  $U_0$ , i.e., it is independent (or detached) from the PCs underlying the VC. The VC can be offloaded without closing the underlying PCs, simply by  $U_0$  posting  $tx^{vc}$ . Once posted, all PCs can be honestly settled, updating the PC to reflect the refund or the success of the Blitz payment, as in Blitz itself or other synchronization protocols.

**Closing the VC.** One of the most essential operations of the VC operation is closing the VC honestly, i.e., off-chain. This is challenging, because closing needs to proceed in a way, such that no one is at risk of losing funds. To solve this challenge, we first observe that if the receiver  $U_n$  already owns all  $\alpha$  coins in the VC, the VC endpoints need merely wait until the Blitz timeout  $T$  runs out. At this point, the Blitz payment will be successful automatically. But what about when  $U_n$  owns  $0 \leq \alpha' < \alpha$  coins in the VC? We need a protocol that atomically changes the value of the Blitz transaction from  $\alpha$  to  $\alpha'$ . To solve this issue, we introduce a new protocol, called synchronized modification, which given a payment of value  $\alpha$  tied to transaction  $tx^{vc}$  and a timeout  $T$ , allows for updating the payment to a value  $\alpha'$  such that  $0 \leq \alpha' < \alpha$ . This is illustrated in Figure 7.

**Synchronized modification works as follows.** We can update individual 2-party Blitz contracts to the new value  $\alpha'$  from right to left. An intermediary  $U_i$  is sure to not lose money, because the atomicity of Blitz ensures that in both the left ( $U_{i-1}, U_i$ ), having locked  $\alpha$ , and the right channel ( $U_i, U_{i+1}$ ), having locked  $\alpha'$ , the payment is either refunded or succeeds. In the former case,  $U_i$  does not lose money, as both payments are reverted. In the latter case,  $U_i$  gains  $\alpha$  while paying  $\alpha'$ , so  $U_i$  gets some money. We can incentivize the participation of intermediary users with fees. Alice is incentivized to publish  $tx^{vc}$  if the correct updates do not reach her (paying more money than she owes otherwise), thereby ensuring the atomicity of the synchronized modification. If all the channels are updated, they can simply go idle waiting for the payment to be successful

after  $T$ , or they can finalize this payment instantly by using the fast track functionality [9].

**Fair unlimited lifetime.** The timeout parameter  $T$  serves an additional purpose here: It is the lifetime of the VC. VC endpoints need to close the VC before  $T$  expires. Interestingly, we can use the aforementioned synchronized modification operation also for extending this lifetime. In particular, besides updating the contracts in each channel to a smaller amount, as shown in Figure 7, we can in fact update the timeout  $T$  in each channel. Before the initial timeout  $T$  expires, the VC endpoints can run a synchronized modification update from receiver to sender. If everyone agrees, they can update to the time  $T' > T$ , and intermediaries would charge a fee for this. Intuitively, users are incentivized to agree as they are fine to pay their money later (at  $T'$ ) to their right while receiving it earlier (at  $T$ ) on their left. This solves the problem of the a priori unlimited lifetime of prior VC constructions. The endpoints have the guarantee that the VC remains virtual until a pre-defined timeout, while the intermediaries have a guarantee that they can unlock their collateral after at most a pre-defined timeout without going on-chain and they can prolong it if everyone agrees for as long as they wish. Since the time for which the VC is prolonged is known, intermediaries can adopt a fee model that is based on time, which is not possible in existing solutions.

## V. DONNER: PROTOCOL DESCRIPTION

### A. Security and privacy goals

We informally define three security and three privacy goals for our VC construction. For formal definitions of these properties and proofs, we refer to the extended version of this work [10]. We mark security goals with an  $S$  and privacy goals with a  $P$ . Side channel attacks (e.g., probing and balance discovery) constitute a significant privacy threat for PCNs [26]. Here, we rule out side channels from the attacker model to reason about the leakage induced by the design of the VC construction itself.

**(S1) Balance security.** Honest intermediaries do not lose any coins when participating in the VC construction.




Fig. 7: Synchronized modification: Safely modify the contract tied to a transaction  $tx^{vc}$  in each channel atomically. Note that  $tx^{vc}$  is the same transaction in all three cases.

**(S2) Endpoint security.** No user can steal the sender's balance of the VC. Additionally, the receiver is always guaranteed to get at least its VC balance.

**(S3) Reliability.** No (possibly colluding) intermediaries can force two honest endpoints of a VC to close or offload the VC before the lifespan  $T$  of the VC expires.

**(P1) Endpoint anonymity.** In an optimistic VC execution, intermediaries cannot distinguish if their left (right) user is the sending (receiving) endpoint or merely an honest intermediary connected to the sending (receiving) endpoint via other, non-compromised users.

**(P2) Path privacy.** In an optimistic VC execution, intermediaries do not learn any identifiable information about the other intermediaries, except for their direct neighbors.

**(P3) Value privacy.** The users on the path learn only about the initial and the final balance of the VC, not the value of the individual payments.

The careful readers may have noticed that P1 and P2 hold only for the optimistic case. Indeed, like in any other off-chain protocol (e.g., the Lightning Network), the channels have to go on-chain in order to resolve disputes in the worst case. This means that anyone observing the blockchain can reconstruct the path. Note, however, that this happens rarely, as the optimistic case is less costly for the participants. Designing off-chain protocols that achieve privacy even in case of disputes is an interesting open question.

### B. Assumptions and prerequisites

**Digital signatures.** A digital signature scheme is a tuple of algorithms  $\Sigma := (\text{KeyGen}, \text{Sign}, \text{Vrfy})$ . On a high level,  $(pk, sk) \leftarrow \text{KeyGen}(\lambda)$  is a PPT algorithm that on input a security parameter  $\lambda$  generates a keypair  $(pk, sk)$ . The public key  $pk$  is publicly known, while the secret key  $sk$  is only known to the user who generated that keypair.  $\sigma \leftarrow \text{Sign}(sk, m)$  is a PPT algorithm that on input a secret key  $sk$  and a message  $m \in \{0, 1\}^*$  generates a signature  $\sigma$  of  $m$ . Finally,  $\{0, 1\} \leftarrow \text{Vrfy}(pk, \sigma, m)$  is a DPT algorithm that on input a public key  $pk$ , a message  $m$  and a signature  $\sigma$  outputs 1 iff the signature is a valid authentication tag for  $m$  w.r.t.  $pk$ . We use a EUF-CMA secure [23] signature scheme  $\Sigma$  as a black-box throughout this work.

**Payment channel notation.** We model each payment channels as a tuple:  $\gamma := (\text{id}, \text{users}, \text{cash}, \text{st})$ . The attribute  $\gamma.\text{id} \in \{0, 1\}^*$  uniquely identifies a channel;  $\gamma.\text{users} \in P^2$  identifies the two parties involved in the channel out of the set of all parties  $P$ . Moreover,  $\gamma.\text{cash} \in \mathbb{R}_{\leq 0}$  denotes the total monetary capacity (i.e., the coins) of the channel and the current state is stored as a vector of outputs of  $tx^{\text{state}}$ :  $\gamma.\text{st} := (\theta_1, \dots, \theta_n)$ . In this work, we use channels in paths from a sender to a receiver. For simplicity, we say that  $\gamma.\text{left} \in \gamma.\text{users}$  refers to the user closer to the sender, while  $\gamma.\text{right} \in \gamma.\text{users}$  refers to the user closer to the receiver. The balance of both users can always be inferred from the current state  $\gamma.\text{st}$ . For convenience, we say that  $\gamma.\text{balance}(U)$  gives the coins owned by  $U \in \gamma.\text{users}$  in this channel's latest state  $\gamma.\text{st}$ . Finally, we define a channel skeleton  $\gamma$  for a channel  $\gamma$ , as  $\gamma := (\gamma.\text{id}, \gamma.\text{users})$ .

**Ledger and channels.** We use the ledger (or blockchain) and a PCN (both introduced in Section II) as black-boxes in our

construction. The ledger keeps a record of all transactions and balances and is append-only. The PCN supports opening, updating and closing of PCs. We assume the PCs involved in VCs to be already open. We interact with ledger and PCN through the following procedures.

**publishTx( $\bar{tx}$ ):** The transaction  $\bar{tx}$  is posted on-chain after at most  $\Delta$  time (the blockchain delay), if it is valid.

**updateChannel( $\bar{v}_i, tx_i^{state}$ ):** This procedure initiates an update in the channel  $\bar{v}_i$  to the state  $tx_i^{state}$ , when called by a user  $i \in \bar{v}_i.users$ . The procedure terminates after at most  $t_u$  time and returns (update-ok) in case of success and (update-fail) in case of failure to both users. We call this function also to update our VC hosted on  $tx^{vc}$ .

**closeChannel( $\bar{v}_i$ ):** This procedure closes the channel  $\bar{v}_i$ , when called by a user  $i \in \bar{v}_i.users$ . The latest state transaction  $tx_i^{state}$  appears on the ledger after at most  $t_c$  time.

**preCreate( $tx^{vc}$ , index,  $U_0, U_n$ ):** Pre-creates the VC  $\bar{v}_{vc}$ , exchanging the initial state transactions with the other user in  $\bar{v}_{vc}.users := (U_0, U_n)$  based on the output identified by index of the funding transaction  $tx^{vc}$  that remains off-chain for now. It finally returns  $\bar{v}_{vc}$ .

**Assumptions and remarks.** In our construction, we assume that every user  $U$  has a public key  $pk_U$  to receive transactions. Additionally, we assume that honest parties stay online for the duration of the protocol, like in the Lightning Network. A path finding algorithm to identify a payment path can be called by  $pathList \leftarrow GenPath(U_0, U_n)$ . This will return a path in the PCN from  $U_0$  to  $U_n$ . Path finding algorithms are orthogonal to the problem tackled in this paper and we refer the reader to [34], [35] for more details. Finally, we assume fee to be a publicly known value charged by every user. Note that in practice, every user can charge an individual fee. We reuse the pseudo-code definitions of Setup, Open, Finalize and Respond from [9] in Figure 8.

### C. Detailed construction and pseudocode

Recall the setting, where  $U_0$  and  $U_n$ , connected via  $U_i$  for  $i \in [1, n-1]$ , wish to open a bidirectional VC with capacity  $\alpha$  fully funded by  $U_0$ . We consider the different phases of Donner: OpenVC, UpdateVC, CloseVC, ProlongVC and Respond. We show the used macros in Figure 8(a), the procedure for updating individual PCs for the close or prolong VC phase in Figure 8(b), and the whole protocol in Figure 8(c). For completeness, we explain the protocol including the operations of Blitz [9] below in prose, while in Figure 8(c) we show a modularized protocol based on the operations setup, open, finalize and respond. We remark that in this work, we could use any other construction providing the same functionality, e.g., this can be achieved by smart contract enabling UTXO-based chains such as the EUTXO model used in Cardano [15]. For better readability we simplify the protocol, e.g., we omit ids required for routing VCs concurrently. For the formal protocol description in the UC framework, we defer to the extended version of this work [10].

**OpenVC.** This phase makes use of a modified Blitz Setup phase (Setup\*) and Open/Finalize of Blitz. Setup\*: The sender  $U_0$  starts by creating a transaction  $tx^{vc}$  that contains an

output  $\theta_{vc}$  holding  $\alpha$  coins spendable under the condition  $MultiSig(U_0, U_n)$  and  $n$  outputs  $\theta_{e_i}$  holding  $\epsilon$  coins each spendable under the condition  $OneSig(U_i) + RelTime(t_c + \Delta)$ , one for every user  $U_i$  for  $i \in [0, n-1]$ . Spending from  $\theta_{vc}$ ,  $U_0$  and  $U_n$  create commitment transactions for the VC with  $\bar{v}_{vc} := preCreate(tx^{vc}, 0, U_0, U_n)$ . This function pre-creates the VC  $\bar{v}_{vc}$ , exchanging the initial state transactions with the other user in  $\bar{v}_{vc}.users := (U_0, U_n)$  based on the output with index 0 of the funding transaction  $tx^{vc}$  that remains off-chain for now. It finally returns  $\bar{v}_{vc}$ .

**Open (Blitz):** Then, each pair of users from  $U_0$  to  $U_n$  performs 2pSetup of [9], which we briefly summarize as follows. Sender  $U_0$  presents its neighbor  $U_1$  with  $tx^{vc}$  and an update of their channel to a state, where  $\alpha$  coins of  $U_0$  are spendable under the condition  $\phi = (OneSig(U_0) \boxtimes AbsTime(T)) \boxtimes (MultiSig(U_0, U_1) \boxtimes RelTime(\Delta))$ . Passing along  $tx^{vc}$  does not violate privacy, due to the usage of stealth addresses, see Appendix D.1.

Before actually updating the channel,  $U_1$  gives  $U_0$  its signature for  $tx_0^r$ .  $tx_0^r$  takes as inputs the output holding  $\alpha$  of the aforementioned proposed state update and the output  $\theta_{e_0}$  of  $tx^{vc}$  holding  $\epsilon$  under  $U_0$ 's control. After receiving the signature, they perform this update and revoke their previous state. In the same fashion,  $U_1$  continues this procedure with its neighbor  $U_2$  and this continues with its neighbor until the receiver  $U_n$  has successfully updated its channel with its left neighbor  $U_{n-1}$ . Then,  $U_n$  sends a confirmation to  $U_0$  (Finalize).

**UpdateVC.** At this point the VC  $\bar{v}_{vc}$  is considered to be open and ready to be used. An update can be performed by creating a new state  $tx_i^{state}$  and calling  $updateChannel(\bar{v}_{vc}, tx_i^{state})$ . This function updates the VC  $\bar{v}_{vc}$ , changing the latest state transaction to  $tx_i^{state}$  and revoking the previous one. In case of a dispute, the users wait until the VC is offloaded. At this time, the VC is closed.

In the beginning, the whole balance lies with  $U_0$ , but once the balance is redistributed, the channel is usable in both directions. Should they wish to construct a channel where they both hold some balance initially, they can start the construction in the other direction for a second time, as we discuss in Appendix B. When they have rebalanced the money inside the VC and definitely before time  $T$ , they proceed to the next phase, the closing phase.

**CloseVC/ProlongVC (Synchronized modification).** For closing the VC, assume its final balance is  $\alpha - \alpha'$  belonging to  $U_0$  and  $\alpha'$  to  $U_n$  (and  $T' = T$ ). For prolonging the lifetime, assume the new time is  $T' > T$  (and  $\alpha' := \alpha$ ). In either case, pairs of users from perform the new functionality 2pModify from right to left, which we outline as follows.  $U_n$  starts the following update process with its left neighbor  $U_{n-1}$ .  $U_n$  presents a state, where (instead of  $\alpha$ ) only  $\alpha'$  coins from  $U_{n-1}$  are spendable under the condition  $\phi = (OneSig(U_n) \boxtimes AbsTime(T)) \boxtimes (MultiSig(U_{n-1}, U_n) \boxtimes RelTime(\Delta))$  (closing) or the time in this condition is changed to  $T'$  (prolong). For this new state,  $U_n$  creates a transaction  $tx_{n-1}^r$  spending this output and the output of  $tx^{vc}$  belonging to  $U_{n-1}$  and gives its signature for this new  $tx_{n-1}^r$  to  $U_{n-1}$ . After  $U_{n-1}$  checks that the new state and new  $tx_{n-1}^r$  are correct, they update their channel to this new state and revoke the previous one (cf. Figure 8(b)).

User  $U_{n-1}$  continues this process with its left neighbor  $U_{n-2}$  and so on, until the sender  $U_0$  is reached.  $U_0$  checks that the balance in the state update is actually the balance that  $U_0$  owes  $U_n$  in the VC,  $\alpha'$ . If it is not the same, or no such request reaches the sender,  $U_0$  simply publishes  $tx^{vc}$  on-chain and claims  $tx^r$  before the currently active timeout  $T$  expires. In the case where the correct request reaches the sender, they can either continue using the VC until  $T'$  (prolong) or in the case of closing, they wait until  $T$  expires, at which the money  $\alpha'$  automatically moves from left to right to the receiver, or they perform the fast-track mechanism of [9] to immediately unlock their funds (cf. Appendix B). VC endpoints do not need to wait until  $T$ , but can close the VC well before if they wish to do so.

**Respond.** This phase corresponds to the phase with the same name of Blitz, which proceeds thus. Participants have to monitor the ledger if  $tx^{vc}$  is published. In case it is published and its outputs are spendable before  $T$ , each user  $U_i$  for  $i \in [0, n-1]$  needs to refund the money they staked in their right channel. They can either do this off-chain if their right neighbor is cooperating or in the worst case, forcefully on-chain via  $tx^r$ . Similarly, after time  $T$  has expired without  $tx^{vc}$  being published on-chain, each user  $U_i$  for  $i \in [1, n]$  can claim the money from their left channel. Again, this can happen honestly off-chain or forcefully via  $tx^p$ .

**Remarks.** Because we detached the funding transaction from the underlying channels, we additionally get rid of the other issues presented in Section III-D. Since the funding can be published independently from the channels and the collateral outcome depends on the funding, we give back the possibility to intermediaries to resolve their channels honestly. Additionally, as the funding is not constructed by combining the outputs of the underlying channels in sequence, we eliminate the additional linear on-chain transactions (needing only one) and reduce the linear (or logarithmic) time delay for publishing the funding transaction to a constant. Further, as we discuss in Section VI, Donner achieves a better level of privacy. For a formal privacy treatment as well as an illustration of the full Donner construction along with the offload operation, we defer the reader to the extended version of this work [10].

## VI. SECURITY ANALYSIS

### A. Informal security analysis

**Balance security.** When the VC is opened, a Blitz [9] collateral payment is simultaneously opened from sender to receiver. A Blitz payment provides balance security to the intermediaries. An intermediary is merely involved in a payment, the outcome of which is atomically determined by whether or not  $tx^{vc}$  is posted. For both of these outcomes, the intermediary does not lose money. As already argued in Section IV the synchronized modification operation does not put an intermediary at risk.

**Endpoint security.** An honest sender can always enforce the VC that holds its correct balance by posting  $tx^{vc}$  and thereby offloading the VC. By doing so, the refunding of the collateral along the path is triggered, including the one of the sender itself. This means that in case of a dispute or someone not cooperating, the sender can always use the offloading before  $T$  to ensure its balance. An honest receiver will get its rightful

(a) Macros: **genState** $(\alpha_i, T, \gamma_i)$ : Generates and returns a new channel state carrying transaction  $tx_i^{state}$  from the given parameters. **genPay** $(tx_i^{state})$ . Returns  $tx^p$ , which takes  $tx_i^{state}.output[0]$  as input and creates a single output  $:= (\alpha_i, OneSig(U_{i+1}))$ . **genRef** $(tx_i^{state}, tx^{vc}, \theta_{\epsilon_i})$ . Return  $tx^r$ , which takes as input  $tx_i^{state}.output[0]$  and  $\theta_{\epsilon_i} \cdot tx^{vc}.output$ . The calling user  $U_i$  makes sure that this output belongs to an address under  $U_i$ 's control. It creates a single output  $tx^r.output := (\alpha_i + \epsilon, OneSig(U_i))$ , where  $\alpha_i, U_i, U_{i+1}$  are taken from  $tx_i^{state}$ .

(b) 2-party operation: **2pModify** $(\gamma_i, tx^{vc}, \alpha_i, T')$

Let  $T$  be the timeout,  $\alpha_i$  the amount and  $\theta_{\epsilon_i}$  be the output used for the two party contract set up between  $U_{i-1}$  and  $U_i$ , known from **2pSetup** executed in **Open** [9] phase.

$U_i$ :  $tx_{i-1}^{state'} := \text{genState}(\alpha_i', T', \gamma_{i-1})$ ,  $tx_{i-1}^r := \text{genRef}(tx_{i-1}^{state'}, \theta_{\epsilon_{i-1}})$ , then send  $(tx_{i-1}^{state'}, tx_{i-1}^r, \sigma_{U_i}(tx_{i-1}^r))$  to  $U_{i-1}$   $// \theta_{\epsilon_{i-1}}$  known as  $\theta_{\epsilon_x}^{i-1}$  from **2pSetup**  $U_{i-1}$  upon  $(tx_{i-1}^{state'}, tx_{i-1}^r, \sigma_{U_i}(tx_{i-1}^r))$ :

1) Extract  $\alpha_i'$  and  $T'$  from  $tx_{i-1}^{state'}$ . Check that  $\alpha_i' \leq \alpha_i$ ,  $T' \geq T$  and  $tx_{i-1}^{state'} = \text{genState}(\alpha_i', T', \gamma_{i-1})$ . If  $U_{i-1} = U_0$ , ensure that  $\alpha_i \leq x_{i-1} \cdot n \cdot \text{fee}$  where  $x$  is the final balance of  $U_n$  in the virtual channel. Check that  $\sigma_{U_i}(tx_{i-1}^r)$  is a correct signature of  $U_i$  for  $tx_{i-1}^r$ . Check that  $tx_{i-1}^r = \text{genRef}(tx_{i-1}^{state'}, \theta_{\epsilon_{i-1}}) // \alpha_i, T'$  and  $\theta_{\epsilon_{i-1}}$  from **2pSetup**<sup>1</sup>

2) **updateChannel** $(\gamma_i, tx_{i-1}^{state'})$

3) If, after  $t$  time has expired, the message (update-ok) is returned, replace variables  $tx^{state}$  and  $tx_{i-1}^r$  with  $tx_{i-1}^{state'}$  and  $tx_{i-1}^r$ , respectively. Return  $(\gamma_i, \alpha_i', T')$ . Else, return  $\gamma_{i-1}$

$U_i$ : Upon (update-ok), replace variables  $tx^{state}$ ,  $tx_{i-1}$  and  $tx_{i-1}^p$  with  $tx_{i-1}^{state'}$ ,  $tx_{i-1}^p$  and  $tx_{i-1} := \text{genPay}(tx_{i-1}^{state'})$ , respectively.

(c) Protocol:

**OpenVC**

(i) **Setup**<sup>2</sup> (see also Appendix D, Figure 11), as in [9], except:

- Create  $tx^{vc}$  instead of  $tx^{er}$  as shown in Figure 3b
- $\gamma_{vc} := \text{preCreate}(tx^{vc}, 0, U_0, U_n)$  together with  $U_n$  after creating  $tx^{vc}$ , to create the VC commitment transactions.

(ii) **Open** and **Finalize** (see also Appendix D, Figure 11) as in [9]

**UpdateVC**

Either user  $U_i$  or  $\gamma_{vc}$ -users can update the VC  $\gamma_{vc}$  by creating a new state  $tx_i^{state}$  and calling **updateChannel** $(\gamma_{vc}, tx_i^{state})$ .

**CloseVC/ProlongVC** (synchronized modification)

(i) **InitClose/InitProlong**

$U_n$ : Let  $\alpha_i'$  be the final balance of  $U_n$  in the virtual channel and  $T' = T$  (Close) or let  $T' > T$  be the new lifetime of the VC and leave  $\alpha_i' = \alpha_i$  (Prolong). Execute **2pModify** $(\gamma_i, tx^{vc}, \alpha_i', T')$   $U_{i-1}$  upon  $(\gamma_i, \alpha_i, T)$ : If  $U_{i-1} = U_0$ , let  $\alpha_{i-1} := \alpha_i + \text{fee}$ . Execute **2pModify** $(\gamma_{i-2}, tx^{vc}, \alpha_{i-1}, T')$

(ii) **Emergency-Offload**

$U_0$ : If  $U_0$  has not successfully performed **2pModify** with the correct value  $\alpha'$  (plus fee for each hop) until  $T - t_c - 3\Delta$ , publish  $tx(tx^{vc}, \sigma_{U_0}(tx^{vc}))$ . Else, update  $T := T'$

**Respond** (see also Appendix D, Figure 11) as in [9]

Fig. 8: (a) macros, (b) 2-party operation, (c) protocol

balance either when the channel is offloaded or, if it is not, after time  $T$  through the collateral, which is moved from left to right along the path.

**Reliability.** Only the sender is able to offload the VC. This

means that if sender and receiver are honest, no one can force them to offload the VC before T.

**Endpoint anonymity and path privacy.**  $\text{tx}^{\text{VC}}$  is constructed, as in Blitz, based on fresh and stealth addresses and the endpoints of the VC rely on fresh addresses too. Hence, an intermediary observing  $\text{tx}^{\text{VC}}$  learns no meaningful information about the sender, the receiver, and the path. This holds only in the optimistic case. In the pessimistic case, it might be possible to link (parts of) the path to  $\text{tx}^{\text{VC}}$  and also link the VC to sender/receiver, like in any other off-chain protocol, including the Lightning Network.

**Value privacy.** Similarly to how payments between users of a payment channel (PC) are known only to those users, also VC updates are only known to the endpoints. There occur no on-chain transactions in the optimistic case throughout the protocol. Any two users connected in the PC network can open a VC, and apart from their open and close balance, the amount and nature of the individual updates remains known only to them, even in the pessimistic case.

## B. Security model

We rely on the synchronous, global universal composability (GUC) framework [14] to model the Donner protocol. We make use of some preliminary functionalities commonly used in the literature [7]–[9], [19], [20]. The global ledger  $L$  is maintained by the functionality  $G_{\text{Ledger}}$ , which is parameterized by a signature scheme  $\Sigma$  and a blockchain delay  $\Delta$ , i.e., an upper bound on number of rounds it takes for a valid transaction to appear on  $L$ , after it is posted. The notion of time (or computational rounds) is modelled by  $G_{\text{clock}}$  and the communication by  $F_{\text{GDC}}$ . Finally, a functionality  $F_{\text{Channel}}$  handles the creation, update and closure of PCs as well as the preparation and update of the VCs.

We define an ideal functionality  $F_{\text{VC}}$  that models the idealized behavior of our VC protocol, stipulating input/output behavior, impact on the ledger as well as possible attacks by adversaries. In the ideal world,  $F_{\text{VC}}$  is a trusted third party. Additionally, we formally define the real world hybrid protocol  $\Pi$  and show that  $\Pi$  emulates (or realizes)  $F_{\text{VC}}$ . For this, we describe a simulator  $S$  that translates any attack of any adversary on  $\Pi$  into an attack on  $F_{\text{VC}}$ .

To show that the protocol  $\Pi$  realizes  $F_{\text{VC}}$ , we need to show that no PPT environment  $E$  can distinguish between interacting with the real world and interacting with the ideal world with non-negligible probability. This implies, that any attack that is possible on the protocol is also possible on the ideal functionality. Intuitively, it suffices to output the same messages and add the same transaction to the ledger in both the real and the ideal world in the same rounds. We refer to the extended version of this work [10] for the preliminaries, the ideal functionality, the formal protocol, the simulator, the formal proof of Theorem 1 and the formalization of the security and privacy goals of Section V-A as well as the proof that  $F_{\text{VC}}$  has these properties.

**Theorem 1.** For functionalities  $G_{\text{Ledger}}$ ,  $G_{\text{clock}}$ ,  $F_{\text{GDC}}$ ,  $F_{\text{Channel}}$  and for any ledger delay  $\Delta \in \mathbb{N}$ , the protocol  $\Pi$  UC-realizes the ideal functionality  $F_{\text{VC}}$ .

TABLE II: Communication overhead of Donner for the whole path (not per party) for the different operations, assuming a VC across  $n$  channels. In the pessimistic offload,  $k \in [0, n]$  is the number of channels where there is a dispute. Only in the Offload case transactions are posted on-chain.

|                       | # txs           | size (bytes)                        | on-chain cost (USD)                  |
|-----------------------|-----------------|-------------------------------------|--------------------------------------|
| Open                  | $4 \cdot n + 2$ | $34 \cdot n^2 + 1240 \cdot n + 695$ | 0                                    |
| Update                | 2               | 695                                 | 0                                    |
| Close                 | $3 \cdot n$     | $1048 \cdot n$                      | 0                                    |
| Offload (Optimistic)  | 1               | $192 + 34 \cdot n$                  | $0.25 + 0.04 \cdot n$                |
| Offload (Pessimistic) | $3k + 1$        | $1048 \cdot k + 192 + 34 \cdot n$   | $1.36 \cdot k + 0.25 + 0.04 \cdot n$ |

## VII. EVALUATION AND COMPARISON

**Communication overhead.** We implemented a small proof-of-concept that creates the raw Bitcoin transactions necessary for Donner [1]. We use the library `python-bitcoin-utils` and Bitcoin Script to build the transactions and tested their compatibility with Bitcoin by deploying them on the testnet. We show the results for the operations Open, Update, Close, Offload in Table II. For transactions that go on-chain, we provide additionally the expected cost in USD at the time of writing. For this evaluation we assume generalized channels [8] as the underlying payment channel (PC) protocol, but note that Donner is also compatible with Lightning channels (as we discuss at the end of this section).

For opening a virtual channel (VC), each of the  $n$  underlying PCs needs to exchange 4 transactions:  $\text{tx}^{\text{VC}}$ ,  $\text{tx}_i^{\text{r}}$  and two transactions for updating the state. Since  $\text{tx}^{\text{VC}}$  has an output for every intermediary and the sender, its size increases with the number of channels on the path  $n$  and is  $192 + 34 \cdot n$  bytes.  $\text{tx}_i^{\text{r}}$  has a size of 306 bytes, and a channel update to a state holding this contract is 742 bytes.  $\text{tx}_i^{\text{p}}$  does not need to be exchanged, since the left user of a channel can generate it independently. This totals to  $1240 + 34 \cdot n$  bytes of off-chain communication per channel for the open phase. Then, we require to exchange the initial state of the VC, which is 2 transactions or 695 bytes. This totals  $4 \cdot n + 2$  transactions or  $34 \cdot n^2 + 1240 \cdot n + 695$  bytes for the path.

For honestly closing a VC, the payment needs to be updated from right to left. However,  $\text{tx}^{\text{VC}}$  does not need to be exchanged anymore, so we only need to exchange 3 transactions or 1048 bytes for each of the  $n$  underlying channels. To update a VC, the two endpoints need to exchange 2 transactions with 695 bytes, the same as a PC update.

Finally, for offloading, only the transaction  $\text{tx}^{\text{VC}}$  needs to be posted on-chain and nothing per channel. This means  $192 + 34 \cdot n$  bytes and costs  $0.25 + 0.04 \cdot n$  USD. Note that if individual users on the path do not collaborate, regardless if the VC is offloaded or successfully closed, these channels may need to be closed as well. We argue that this is also the case during the normal PC execution, e.g., when routing multi-hop payments. However, for every channel that does need to be closed, the three transactions exchanged in the close phase need to be posted additionally. If there are  $k$  channels with such a dispute, this results in a total of  $3k + 1$  transactions or  $1048 \cdot k + 192 + 34 \cdot n$  bytes, which costs  $1.36 \cdot k + 0.25 + 0.04 \cdot n$  USD for the whole path. We mark this as the pessimistic case in Table II.

**Efficiency comparison.** We compare our construction to LVPC [25] and Elmo [27] (cf. Table III), the only current

TABLE III: Comparison of LVPC, Elmo and Donner for a VC over from  $U_0$  to  $U_n$ .<sup>1</sup>In the pessimistic offload in Donner,  $k \in [0, n]$  is the number of channels where there is a dispute.

|                          |                     | # txs             | off-chain |
|--------------------------|---------------------|-------------------|-----------|
| Open                     | LVPC [25]           | $7 \cdot (n - 1)$ | ✓         |
|                          | Elmo [27]           | $\Theta(n^3)$     | ✓         |
|                          | Donner              | $4 \cdot n + 2$   | ✓         |
| Update                   | LVPC [25]           | 2                 | ✓         |
|                          | Elmo [27]           | 2                 | ✓         |
|                          | Donner              | 2                 | ✓         |
| Close                    | LVPC [25]           | $4 \cdot (n - 1)$ | ✓         |
|                          | Elmo [27]           | $3 \cdot n + 3$   | ✗         |
|                          | Donner              | $3 \cdot n$       | ✓         |
| Offload<br>(Optimistic)  | LVPC [25]           | $5 \cdot (n - 1)$ | ✗         |
|                          | Elmo [27]           | $3 \cdot n + 1$   | ✗         |
|                          | Donner              | 1                 | ✗         |
| Offload<br>(Pessimistic) | LVPC [25]           | $5 \cdot (n - 1)$ | ✗         |
|                          | Elmo [27]           | $3 \cdot n + 1$   | ✗         |
|                          | Donner <sup>1</sup> | $3 \cdot k + 1$   | ✗         |

Bitcoin-compatible VC solutions over multiple hops. As already mentioned, LVPC and Elmo have rooted VC funding transactions. We evaluate, in particular, the off-chain and on-chain costs of the core VC operations (open, update, close, and offload), concluding that Donner is better in each case.

LVPC is constructed recursively; there are different ways of doing the recursion. Each combination leads to the same minimum number of VCs required for a path of  $n$  base channels: One for each of the  $n - 1$  intermediaries. The storage overhead per intermediary is linear in the number of layers on top of a user, which in turn is in the worst case linear and in the best case logarithmic in the path length (more on this in the extended version of this work [10]).

In the open phase across the whole path, Donner requires  $4 \cdot n + 2$  off-chain transactions for the whole path. In LVPC, 7 off-chain transactions per VC are needed, so  $7 \cdot (n - 1)$ . Similar, for closing, we need to store 4 transactions per VC in LVPC, so  $4 \cdot (n - 1)$ . Elmo requires to store  $n - 2 + \chi_{i=2} + \chi_{i=n-1} + (i - 2 + \chi_{i=2})(n - i - 1 + \chi_{i=n-1}) \in \Theta(n^2)$  (where  $\chi_p$  is 1 if  $P$  is true and 0 otherwise) for the  $i^{\text{th}}$  intermediary (and 1 for the endpoints), resulting in a storage overhead of  $\Theta(n^3)$  for the whole path. Closing honestly (i.e., off-chain) is not defined for Elmo, so it needs to be closed on-chain, resulting in 2 transactions per channel ( $n$ ) for closing plus 1 transaction per user ( $n + 1$ ) plus 2 transactions to close the VC or  $3 \cdot n + 3$  on-chain transactions. Donner requires the close operation per underlying channel, so  $3 \cdot n$  transactions. The update phase is the same in all constructions.

The interesting case again is the offload case. As we already pointed out, a fully rooted, recursive VC construction requires to close all underlying channels. This means in LVPC, we require 2 transactions per underlying channel, of which we have  $n$  PCs and  $n - 2$  VCs (all but the topmost one). Additionally, we need to publish  $n - 1$  funding transactions of the VCs including the topmost one. This results in  $2 \cdot (2n - 2) + n - 1 = 5 \cdot n - 5$  transactions that have to be posted on-chain along with the fact that all involved channels have to be closed in the case of a dispute. In Elmo, we need  $3 \cdot n + 1$ , i.e., the number of transactions to close minus the 2 transactions required to put the VC state on-chain.

In Donner, only 1 transaction has to be posted on-chain. For the pessimistic offload, there need to be  $3 \cdot k + 1$  transactions posted in Donner, where  $k$  is the number of channels where there is a dispute. We show an application example in the extended version of this work [10], where we analyze how Donner can be used to connect a node better to a network via VCs, compared to no VCs and LVPC.

Compatibility with LN channels. To simplify the formalization of this work, we built our VC construction on top of generalized payment channels (GC) [8], which have one symmetric channel state. However, it is also possible to construct Donner on top of LN channels, which have two asymmetric channel states. The (one-hop) BCVC [7] constructions rely on GCs as well, while the recursive LVPC [25] relies on simple channels that have only one state, but each update reduces the limited lifetime of the channel. (Elmo [27] needs the opcode ANYPREVOUT that is not supported in Bitcoin or in the LN.)

As LN channels are the only ones deployed in practice so far, it is interesting to investigate the effect of building VCs on top of LN channels. We point out that building Donner on top of LN channel is not difficult, as the collateralization in the underlying base channels is similar to a MHP. In fact, the only two differences for implementing Donner on top of LN channels instead of GCs is that (i) for each of the two asymmetric states per channel we now need to create a  $\text{tx}_i^j$  transaction, so two instead of one, and (ii) a punishment mechanism has to be introduced per output instead of per state (e.g., similar to how HTLCs are handled in LN).

The LVPC construction is not as straightforward to implement on top of LN channels. Similarly to Donner, we need to introduce a punishment mechanism (ii). However, the more difficult part is handling the two asymmetric states (i). Since the VC needs to be able to be posted regardless of which of the two states are posted, there needs to be a unique funding transaction (called Merge in [25]) for each possible combination of states in the underlying channels. This implies that in a LVPC like construction which is built on top of LN channels, the storage overhead per party is exponential in the layers of VCs that are constructed over this party. In fact, using channels with duplicated states this exponential growth is present in every rooted, recursive VC construction. This follows from the evaluation in [8]. For each of these exponentially many copies of the VC, commitment transactions need to be exchanged for an update, so there is an exponential communication overhead too. Note that the storage overhead for Donner on top of LN channels is constant as is the communication overhead for updates.

## VIII. CONCLUSION

Payment channel networks (PCNs) have emerged as successful scaling solutions for cryptocurrencies. However, path-based protocols are tailored to payments, excluding novel and interesting non-payment applications such as Discreet Log Contracts, while creating direct PCs on-demand is expensive, slow and infeasible on a large scale. VCs are among the most promising solutions. We show that all existing UTXO-based constructions are vulnerable to the Domino attack, which fundamentally undermines the underlying PCN itself.

Hence we introduce a new VC design, the first one to be secure against the Domino attack, besides the only one achieving path privacy and a time-based fee model. Our performance analysis demonstrates that Donner is more efficient: It only requires a single on-chain transaction to solve disputes, as opposed to a number that is linear in the path length, and the storage overhead is constant too, as opposed to linear.

Overall, Donner offers an easy-to-adopt, LN-compatible VC construction enabling new applications such as Discreet Log Contracts or fast and direct micropayments, without the need to create a direct PC. Unlike the underlying PCNs, the VCs are not susceptible to liveness and privacy attacks by the intermediaries and do not require fees per payment.

**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) 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 and COMET K1 ABC; by the Vienna Business Agency through the project Vienna Cybersecurity and Privacy Research Center (VISPR); by CoBloX Labs; by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association through the Christian Doppler Laboratory Blockchain Technologies for the Internet of Things (CDL-BOT); by the National Science Foundation (NSF) under grant CNS-1846316; by Madrid regional government as part of the program S2018/TCS-4339 (BLOQUES-CM) co-funded by EIE Funds of the European Union; by the project HACRYPT (N00014-19-1-2292); by grant IJC2020-043391-I/MCIN/AEI/10.13039/501100011033 and European Union NextGenerationEU/PRTR; by PRODIGY Project (TED2021-132464B-I00) funded by MCIN/AEI/10.13039/501100011033/ and the European Union NextGenerationEU/PRTR; by SCUM Project (RTI2018-102043-B-I00) MCIN/AEI/10.13039/501100011033/ERDF A way of making Europe.

## REFERENCES

- [1] “Donner vc evaluation of the communication overhead,” 2021, <https://github.com/donner-vc/overhead>.
- [2] “Bitcoin avg. transaction fee historical chart,” Jan. 2022, <https://bitinfocharts.com/comparison/bitcoin-transactionfees.html>.
- [3] “Bitcoin rich list,” Jan. 2022, <https://bitinfocharts.com/top-100-richest-bitcoin-addresses.html>.
- [4] “Bolt #2: Peer protocol for channel management,” Mar. 2022, <https://github.com/lightning/bolts/blob/master/02-peer-protocol.md#rationale-7>.
- [5] “Ln snapshot,” Jan. 2022, <https://ln.fiatjaf.com/>.
- [6] “Simulation of domino attack,” 2022, <https://github.com/donner-vc/simulation>.
- [7] L. Aumayr, O. Ersoy, A. Erwig, S. Faust, K. Hostakova, M. Maffei, P. Moreno-Sanchez, and S. Riahi, “Bitcoin-Compatible Virtual Channels,” in *IEEE Security and Privacy*, 2021.
- [8] —, “Generalized channels from limited blockchain scripts and adaptor signatures,” in *Asiacrypt*, 2021.
- [9] L. Aumayr, P. Moreno-Sanchez, A. Kate, and M. Maffei, “Blitz: Secure Multi-Hop Payments Without Two-Phase Commits,” in *USENIX Security*, 2021.
- [10] —, “Breaking and Fixing Virtual Channels: Domino Attack and Donner,” *Cryptology ePrint Archive*, Report 2021/855, 2021, <https://eprint.iacr.org/2021/855>.
- [11] I. Bentov and R. Kumaresan, “How to use bitcoin to design fair protocols,” in *CRYPTO*, 2014.
- [12] C. Burchert, C. Decker, and R. Wattenhofer, “Scalable funding of bitcoin micropayment channel networks,” in *Stabilization, Safety, and Security of Distributed Systems*, 2017, pp. 361–377.
- [13] J. Camenisch and A. Lysyanskaya, “A formal treatment of onion routing,” in *Advances in Cryptology CRYPTO*, 2005, pp. 169–187.
- [14] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, “Universally composable security with global setup,” in *TCC*, vol. 4392, 2007, pp. 61–85.
- [15] M. M. Chakravarty, J. Chapman, K. MacKenzie, O. Melkonian, M. Peyton Jones, and P. Wadler, “The extended utxo model,” in *FC*. Springer, 2020, pp. 525–539.
- [16] G. Danezis and I. Goldberg, “Sphinx: A compact and provably secure mix format,” in *IEEE Security and Privacy*, 2009.
- [17] “DLC over Lightning,” [dlc-dev] Mailing List, Nov. 2021, available at <https://mailmanlists.org/pipermail/dlc-dev/2021-November/000091.html>.
- [18] T. Dryja, “Discreet Log Contracts,” 2017, available at <https://adiabat.github.io/dlc.pdf>.
- [19] S. Dziembowski, L. ECKEY, S. Faust, J. Hesse, and K. Hostakova, “Multi-party Virtual State Channels,” in *Eurocrypt*, 2019, pp. 625–656.
- [20] S. Dziembowski, L. ECKEY, S. Faust, and D. Malinowski, “Perun: Virtual payment hubs over cryptocurrencies,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 106–123.
- [21] S. Dziembowski, S. Faust, and K. Hostakova, “General State Channel Networks,” in *Computer and Communications Security, CCS*, 2018.
- [22] C. Egger, P. Moreno-Sanchez, and M. Maffei, “Atomic multi-channel updates with constant collateral in bitcoin-compatible payment-channel networks,” in *ACM CCS*, 2019, p. 801–815.
- [23] S. Goldwasser, S. Micali, and R. L. Rivest, “A digital signature scheme secure against adaptive chosen-message attacks,” *SIAM Journal on computing*, vol. 17, no. 2, pp. 281–308, 1988.
- [24] J. Harris and A. Zohar, “Flood & loot: A systemic attack on the lightning network,” in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, ser. AFT ’20, 2020.
- [25] M. Jounen, M. Larangeira, and K. Tanaka, “Lightweight Virtual Payment Channels,” in *19th International Conference on Cryptology and Network Security (CANS)*, 2020.
- [26] G. Kappos, H. Yousaf, A. Piotrowska, S. Kanjalkar, S. Delgado-Segura, A. Miller, and S. Meiklejohn, “An empirical analysis of privacy in the lightning network,” in *FC*, 2021, pp. 167–186.
- [27] A. Kiayias and O. S. T. Litos, “Elmo: Recursive virtual payment channels for bitcoin,” <https://eprint.iacr.org/2021/747>.
- [28] G. Malavolta, P. Moreno-Sanchez, A. Kate, M. Maffei, and S. Ravi, “Concurrency and privacy with payment-channel networks,” in *ACM CCS*, 2017.
- [29] G. Malavolta, P. Moreno-Sanchez, C. Schneidewind, A. Kate, and M. Maffei, “Anonymous multi-hop locks for blockchain scalability and interoperability,” in *NDSS Symposium*, 2019.
- [30] A. Miller, I. Bentov, R. Kumaresan, and P. McCorry, “Sprites: Payment channels that go faster than lightning,” *CoRR*, vol. abs/1702.05812, 2017. [Online]. Available: <http://arxiv.org/abs/1702.05812>
- [31] “Perun network,” 2020, <https://perun.network/>.
- [32] J. Poon and T. Dryja, “The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments,” Jan. 2016, draft version 0.5.9.2, available at <https://lightning.network/lightning-network-paper.pdf>.
- [33] M. Romiti, F. Victor, P. Moreno-Sanchez, P. S. Nordholt, B. Haslhofer, and M. Maffei, “Cross-layer deanonymization methods in the lightning protocol,” in *FC*. Springer, 2021, pp. 187–204.
- [34] S. Roos, P. Moreno-Sanchez, A. Kate, and I. Goldberg, “Settling payments fast and private: Efficient decentralized routing for path-based transactions,” in *NDSS Symposium*, 2018.
- [35] V. Sivaraman, S. B. Venkatakrishnan, K. Ruan, P. Negi, L. Yang, R. Mit-tal, G. C. Fanti, and M. Alizadeh, “High throughput cryptocurrency routing in payment channel networks,” in *NSDI*, 2020, pp. 777–796.

- [36] S. Tochner, A. Zohar, and S. Schmid, "Route hijacking and dos in off-chain networks," ser. AFT, 2020.
- [37] N. Van Saberhagen, "Cryptonote v 2.0 (2013)," URL: <https://web.archive.org/web/20201028121818/https://cryptonote.org/whitepaper.pdf>. White Paper.

## APPENDIX A

### WHEN TO USE VIRTUAL CHANNELS

In the state of the art on off-chain protocols, we can distinguish between generic 2-party applications and simple payments. The former require a direct channel between the parties and therefore it is interesting to compare VCs and direct PCs in this setting. In the latter, PCNs have already been shown to offer improvements over constructing a direct channel and therefore it is worth to compare VC against PCN payments. Next, we highlight use cases of VCs in these two settings.

**VCs vs PCs for 2-party applications.** Imagine that two arbitrary users that do not share a PC or a VC decide to execute a 2-party application between them. The first disadvantage of using a PC over a VC is that over their lifespan they would pay twice as many fees per on-chain transaction (i.e., to open and close the channel). At the current average Bitcoin transaction cost of 4100 satoshi (or 0.000041 BTC or 1.73 USD), the overall cost would be 8200 satoshi (3.46 USD).

Since VCs are currently not being used in practice, there is no fee model for them. To put the cost of opening a VC into perspective, we can compare it to payments over the PCN. Say Alice and Bob are connected by a path of payment channels that has 3 hops (we take the average shortest distance of a current LN snapshot). Taking the current average fees of the LN, and, say, an average transaction amount of 50,000 satoshi (21.10 USD), Alice and Bob could perform 1115 payments in the LN for the same fee of 8200 satoshi (3.46 USD). This means that in this example, the fees paid to intermediaries for operating a VC, i.e., opening and closing, is cheaper in terms of fees, if these VC operating fees are less than the fees of 1115 LN payments.

More generally, we can compare the cost of VC versus PC as follows. We introduce  $x$  as a factor by which VCs are more expensive than PCN payments. A VC channel is cheaper if  $l \cdot (BF + RF \cdot a) \cdot x < 2 \cdot TF$  holds, where  $l$  is the number of hops in the path between the two VC endpoints. Further  $BF$  and  $RF$  are the two types of fees charged in PCN implementation such as the LN, where  $BF$  is a base fee charged by intermediaries for forwarding payments and  $RF$  a relative fee based on the payment amount. We compare this to the transactions fee on-chain  $TF$ , paid twice in the lifespan of a PC. For instance, taking the concrete values from the example above we can write the following:  $3 \cdot (1 + 0.000029 \cdot a) \cdot x < 8200$ .

Secondly, creating direct PCs on-demand for applications such as Discreet Log Contracts instead of VCs is again not scalable. Doing so would incur a continuous on-chain transaction load for opening and closing channels. This is against the purpose of PCs and PCNs, which aim at reducing the on-chain load.

Finally, and perhaps still more importantly, it is not possible to open a short-lived PC, since it requires to wait for the confirmation of the funding transaction on the blockchain, which is around 1 hour in Bitcoin. So for applications that

are time-critical, direct PCs are not an option. Applications such as betting on a sports event, say, half an hour before they end are simply impossible with direct PCs.

**VCs vs PCN payments.** Due to the limited transaction size in Bitcoin, current Lightning channels are limited to hold 483 concurrent payments, which becomes especially critical in a micropayment setting. VCs can be used to overcome this issue. Simply, instead of a payment, an output can be used to collateralize a VC, which in turn can be used to again hold 483 payments or further VCs, effectively helping to mitigate this limitation.

In terms of fees, VCs are more desirable than payments over a PCN in the context of micropayments. This is due to the fact that in a PCN, the intermediaries charge a fee for every payment, while for a VC, the fee is charged only once. We can therefore say that a VC is cheaper, if the (simplified) inequality  $l \cdot (BF + RF \cdot a) \cdot x < l \cdot (n \cdot BF + RF \cdot a)$  holds, where similar to above we use the base fee  $BF$  and relative fee  $RF$  of the LN.  $a$  is the sum of the amounts of all micropayments,  $n$  the number of micropayments and  $x$  again the factor by which a VC is more expensive than a payment. We stress that for any given  $x$  there is a number of payments  $n$ , such that the use of VC becomes cheaper than payments over the PCN, because the base fee  $BF$  is paid for each of the  $n$  micropayments in the PCN setting and only once in the VC setting.

**Offline users.** Routing multi-hop payments (MHPs) through the network requires active participation from the intermediaries. However, users may want to go offline and then cannot route MHPs. To still lend their capacity in a productive way and generate some fees, they can allow other nodes to build a VC over them, using watchtowers to ensure their balance.

## APPENDIX B

### EXTENDED COMPARISON AND DISCUSSION

#### 1. Extended comparison to the state of the art in VCs

Dziembowski et al. [20] proposed the first construction of VCs over a single intermediary. Recursive constructions [21] followed up allowing for creating VCs on top of other VCs (or a pair composed of a VC and a PC), thereby supporting arbitrarily many intermediaries. Dziembowski et al. [19], [31] further extended the expressiveness of VCs proposing the notion of multi-party VCs, where a set of  $n$  participants build an  $n$ -party channel recursively from their pair-wise payment/virtual channels. Unfortunately, all the aforementioned constructions rely on the expressiveness of Turing-complete scripting languages such as that of Ethereum and are based on the account model instead of Unspent Transaction Output (UTXO) model; thus, they are incompatible with many of the cryptocurrencies available today, including Bitcoin itself. Aumayr et al. [7] have later shown how to design a Bitcoin-compatible VC through a carefully crafted cryptographic protocol in the UTXO model, supporting however only one intermediary.

Jourenko et al. [25] have recently introduced the first Bitcoin-compatible construction over multiple intermediaries, called Lightweight Virtual Payment Channels (LVPC), where a VC over one hop is applied recursively to achieve a VC between two users separated by a path of any length. More recently, Kiayias and Litos have introduced Elmo [27], a VC

TABLE IV: Comparison to other virtual channel protocols. We denote dispute as the case where a party needs to enforce their VC funds or be compensated. In the UTXO case, this means offloading.  $\boxtimes$  by synchronizing all channels, this time can be reduced to  $\Theta(\log(n))$ .  $^\dagger$  for single-hop constructions  $n$  is constant, however, since the action/storage overhead/time delay is per user, we write  $\Theta(n)$ .  $^\ddagger$  This depends on using indirect/direct dispute.

|                            | Perun [20]          | GSCN [21]             | MPVC [19]                                | BCVC-V/BCVC-NV [7]  | LVPC [25]             | Elmo [27]             | Donner      |
|----------------------------|---------------------|-----------------------|------------------------------------------|---------------------|-----------------------|-----------------------|-------------|
| Scripting req.             | Ethereum            | Ethereum              | Ethereum                                 | Bitcoin             | Bitcoin               | Bitcoin + ANYPREVOUT  | Bitcoin     |
| Multi-hop                  | no                  | yes                   | yes                                      | no                  | yes                   | yes                   | yes         |
| Domino attack              | no                  | no                    | no                                       | yes                 | yes                   | yes                   | no          |
| Path privacy               | no                  | no                    | no                                       | no                  | no                    | no                    | yes         |
| Storage Overhead per party | $\Theta(n)^\dagger$ | $\Theta(n)^\boxtimes$ | $\Theta(n)^\boxtimes/\Theta(1)^\ddagger$ | $\Theta(n)^\dagger$ | $\Theta(n)^\boxtimes$ | $\Theta(n^3)$         | $\Theta(1)$ |
| Time-based fee model       | yes                 | yes                   | yes                                      | no/yes              | no                    | no                    | yes         |
| Unlimited lifetime         | no                  | no                    | no                                       | yes/no              | no                    | yes                   | yes         |
| Off-chain closing          | yes                 | yes                   | yes                                      | yes                 | yes                   | no                    | yes         |
| Dispute: txs on-chain      | 1                   | 1                     | 1                                        | $\Theta(n)^\dagger$ | $\Theta(n)$           | $\Theta(n)$           | 1           |
| Dispute: time delay        | $\Theta(n)^\dagger$ | $\Theta(n)^\boxtimes$ | $\Theta(n)/1^\ddagger$                   | $\Theta(n)^\dagger$ | $\Theta(n)^\boxtimes$ | $\Theta(n)^\boxtimes$ | 1           |

construction that does not rely on creating intermediate VCs, by instead relying on scripting functionality not present in Bitcoin, i.e., the opcode ANYPREVOUT. In Table IV we compare Donner to existing VC protocols, including those that rely on a Turing-complete scripting language or are limited to a single intermediary.

## 2. Extended discussion

**Detering the Domino attack with fees.** One might think that the Domino attack could be deterred by fees. I.e., intermediaries charge fees high enough to be compensated for having to close and reopen their channel, as well as having to claim the collateral put into the VC, in total at least three transaction per intermediary, in addition to the fees charged for the VC usage. It becomes clear, that this is an infeasible deterrence strategy and is in opposition to the aim of VCs to provide scalable and cheap payments: No user would pay three times an on-chain fee per intermediary for a VC. They would simply post an on-chain transaction or open a new direct PC.

**Unidirectionally funded.** Similar to current PCs in the Lightning Network, our VCs are only funded by  $U_0$ , whom we call the sending endpoint or sender. User  $U_n$  is the receiving endpoint or receiver and the intermediaries are  $\{U_i\}_{i \in [1, n-1]}$ . Even though the VC is only funded by  $U_0$ , once some money has been moved, they can use the channel also in the other direction. Moreover, if they want to have a channel funded from both endpoints, they can simply construct another channel from  $U_n$  to  $U_0$ .

**Choosing the lifetime.** The lifetime  $T$  is chosen by the two endpoints of the VC, depending on how long they plan to use the VC. They propose this to the intermediaries who can, based on this time and the amount they need to lock as a collateral, charge a fee. Note that this  $T$  has to be larger than the time it takes to settle the Blitz contracts,  $T \geq 3\Delta + t_c$ , where  $\Delta$  is an upper bound on the time it takes for a valid transaction to appear on the ledger (i.e., modelling the block delay as mentioned in Section II) and  $t_c$  is the time it takes to close a channel. Intermediaries can prolong the lifetime if they agree and they can charge a fee based on time and amount.

**Properties inherited from Blitz.** The fee mechanism of Blitz can be reused here as well, i.e., the intermediary nodes forward fewer coins than they receive. Additionally, the outputs  $\epsilon$  of  $tx^{vc}$  represent a small number. Since they cannot be 0, they

are the smallest possible value, one dust (546 satoshi), i.e., something that is insignificant in value to the sender. If a VC is closed (honestly) before the lifetime expires, parties do not need to wait until the lifetime expires to unlock their money. They can unlock it right away by using the fast track mechanism of Blitz. We refer the reader for these details to [9]. Finally, reusing the stealth address and onion routing mechanism as in [9] we achieve our desired privacy properties.

## APPENDIX C EXTENDED BACKGROUND

### 1. Transaction graphs

In this section we give a more in-depth explanation and example (Figure 9) of our transaction graph notation. Rounded rectangles represent transactions, if they have a single border it means they are off-chain, with a double border on-chain. Incoming arrows to a transaction represent its inputs. The boxes within transactions denote outputs, the outgoing arrows define how an output can be spent.

More specifically, below an arrow we write who can spend the coins. This is usually a signature that verifies w.r.t. one or more public keys, which we denote as  $\text{OneSig}(pk)$  or  $\text{MultiSig}(pk_1, pk_2, \dots)$ . Above the arrow, we write additional conditions for how an output can be spent. This could be any script supported by the scripting language of the underlying blockchain, but in this paper we only use relative and absolute time-locks. For the former, we write  $\text{RelTime}(t)$  or simply  $+t$ , which signifies that the output can be spent only if at least  $t$  rounds have passed since the transaction holding this output was accepted on the blockchain. Similarly, we write  $\text{AbsTime}(t)$  or simply  $\geq t$  for absolute time-locks, which means that the transaction can be spent only if the blockchain is at least  $t$  blocks long. A condition can be a disjunction of conjunction of subconditions.

### 2. Synchronization example

A multi-hop payment (MHP) allows to transfer coins from  $U_0$  to  $U_n$  through  $\{U_i\}_{i \in [1, n-1]}$  in a secure way, that is, ensuring that no intermediary is at risk of losing money. A mechanism synchronizing all channels on the path is required for a payment, such that each channel is updated to represent the fact that  $\alpha$  coins moved from left to right. We give an example of what we mean in Figure 10.




Fig. 9: (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  $x$  and can be spent by a transaction whose witness satisfies the output condition  $\phi_1 \wedge \phi_2 \wedge (\phi_3 \wedge \phi_4)$ . The input of  $tx$  is not shown.

$$U_0 \xrightarrow[3,16]{7,12} U_1 \xrightarrow[4,6]{8,2} U_2 \xrightarrow[7,11]{11,7} U_3 \xrightarrow[5,4]{9,0} U_4$$

Fig. 10: Example of a MHP in a PCN. Here,  $U_0$  pays 4 coins (disregarding any fees) to  $U_4$ , via  $U_1$ ,  $U_2$  and  $U_3$ . The lines represent payment channels. We write balances as  $(x, y)$ , where  $x$  is the balance of the user on the right, and  $y$  the balance of the user on the left. Above we write the channel balances before and below after the payment. In an MHP, this change of balance should happen atomically in every channel (or not at all).

## APPENDIX D

### EXTENDED MACROS, PREREQUISITES AND PROTOCOL

In this section, we discuss the prerequisites for stealth addresses and onion routing. We give extended pseudo-code for the used subprocedures used in our protocol. Further, we spell out the full protocol pseudocode, including the parts taken from. For the protocol see Figure 11, for the two party protocols used therein see Figure 12. To be transparent about the similarities to [9] and highlight the novelties of this work, we mark the latter in green color.

#### Subprocedures

checkTxIn( $tx^{in}$ ,  $n$ ,  $U_0$ ,  $\alpha$ ):

- 1) Check that  $tx^{in}$  is a transaction on the ledger  $L$ .
- 2) If  $tx^{in}.output[0].cash \geq n \cdot \epsilon + \alpha$  and  $tx^{in}.output[0].\phi = OneSig(U_0)$ , that is spendable by an unused address of  $U_0$ , return  $\top$ . Otherwise, return  $\perp$ . When using this transaction (to fund  $tx^{vc}$ ), the sender will pay any superfluous coins back to a fresh address of itself.

checkChannels(channelList,  $U_0$ ):

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  $U_n$ .

checkT( $n$ ,  $T$ ):

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

genTxVc( $U_0$ , channelList,  $tx^{in}$ ):

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

genState( $\alpha_i$ ,  $T$ ,  $\gamma_i$ ):

- 1) For the users  $U_i := \gamma_i.left$  and  $U_{i+1} := \gamma_i.right$ , create the output vector  $\theta_i := (\theta_0, \theta_1, \theta_2)$ , where
  - $\theta_0 := (\alpha_i, (MultiSig(U_i, U_{i+1})) \wedge RelTime(T)) \wedge (OneSig(U_i) \wedge AbsTime(T))^{i+1}$
  - $\theta_1 := (x_{i+1} - \alpha_i, OneSig(U_i))$
  - $\theta_2 := (x_{U_{i+1}}, OneSig(U_{i+1}))$
 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  $tx^{state}$  be a channel transaction carrying the state with  $tx^{state}.output = \theta_i$ . Return  $tx^{state}_i$ .

checkTxVc( $U_i$ ,  $a$ ,  $b$ ,  $tx^{vc}$ ,  $rList$ ,  $onion_i$ ):

- 1)  $x := GetRoutingInfo(onion_i, U_i)$ . If  $x = \perp$ , return  $\perp$ . If  $U_i$  is the receiver and  $x = H(tx^{vc})$ , return  $(\top, \top, \top, \top)$ . Else, if  $x = (U_i, H(tx^{vc}), onion_i)$ , return  $\top$ .
- 2) For all outputs  $(cash, \phi) \in tx^{vc}.output$  except output with index 0 it must hold that:
  - $cash \neq OneSig(pk_x) \wedge RelTime(t_c + \Delta)$  for some identity  $pk_x$
- 3) For exactly one output  $\theta := (\epsilon, OneSig(U_i)) \wedge RelTime(t + \Delta) \in tx^{vc}.output$  and one element  $R_i \in rList$  it must hold that:
  - Let  $pk_f$  be the corresponding public key of  $OneSig(U_i)$
  - $sk_f := GenSk(a, b, pk_f, R_i)$  must be the corresponding secret key of  $pk_f$
- 4) If the checks in 2 or 3 do not hold, return  $\perp$
- 5) Return  $(sk_f, \theta_{\epsilon}, R_i, U_{i+1}, onion_{i+1})$

### 1. Prerequisites

**Stealth addresses.** In order to hide the underlying path, we use stealth addresses [37] for the outputs in the transaction  $tx^{vc}$ . On a high level, every user  $U$  controls two private keys  $a$  and  $b$ . The respective public keys  $A$  and  $B$  are publicly known. A sender can use these public keys controlled by  $U$  to create a new public key  $P$  and a value  $R$ . The user  $U$  and only the user  $U$  knowing  $a$  and  $b$  can use  $R$ ,  $P$  together with  $a$  and  $b$  to construct the private key  $p$ . In particular, also the sender is unaware of  $p$ . This new one-time public key is unlinkable to  $U$  by anyone observing only  $R$  and  $P$  [37].

**Onion routing.** Like in the Lightning Network, we rely on onion routing [13] techniques like Sphinx [16] to allow users communicate anonymously with each other. This allows users to route the VC correctly through the desired path, while ensuring that intermediaries remain oblivious to the path except for their direct neighbors. On a high level, an onion is a layered encryption of routing information and a payload. Each user




Fig. 11: Pseudocode of the protocol.

in turn can peel off one layer, revealing the next user on the path, the payload meant for the current user and another onion, which is designated for the next user. For simplicity, we use onion routing by calling the following two functions:  $\text{onion} \leftarrow \text{CreateRoutingInfo}(\{U_i\}_{i \in [1, n]}, \{\text{msg}_i\}_{i \in [1, n]})$  generates an onion using the public keys of users  $\{U_i\}_{i \in [1, n]}$  on the path, while the procedure  $\text{GetRoutingInfo}(\text{onion}_i, U_i)$  returns the tuple  $(U_{i+1}, \text{msg}_i, \text{onion}_{i+1})$  when called by the correct user  $U_i$ , or  $\perp$  otherwise.




Fig. 12: Protocol for 2-party channel update.