

Shaduf: Non-Cycle Payment Channel Rebalancing

Zhonghui Ge, Yi Zhang, Yu Long* and Dawu Gu*
Shanghai Jiao Tong University
{zhonghui.ge, zy0105, longyu, dwgu}@sjtu.edu.cn

Abstract—A leading approach to enhancing the performance and scalability of permissionless blockchains is to use the payment channel, which allows two users to perform off-chain payments with almost unlimited frequency. By linking payment channels together to form a payment channel network, users connected by a path of channels can perform off-chain payments rapidly. However, payment channels risk encountering fund depletion, which threatens the availability of both the payment channel and network. The most recent method needs a cycle-based channel rebalancing procedure, which requires a fair leader and users with rebalancing demands forming directed cycles in the network. Therefore, its large-scale applications are restricted.

In this work, we introduce Shaduf, a novel non-cycle off-chain rebalancing protocol that offers a new solution for users to shift coins between channels directly without relying on the cycle setting. Shaduf can be applied to more general rebalancing scenarios. We provide the details of Shaduf and formally prove its security under the Universal Composability framework. Our prototype demonstrates its feasibility and the experimental evaluation shows that Shaduf enhances the Lightning Network performance in payment success ratio and volume. Moreover, our protocol prominently reduces users' deposits in channels while maintaining the same amount of payments.

I. INTRODUCTION

Blockchain has seen rapid development over the last decade, enabling secure transactions in a distributed and trust-less setting. Complex transaction logic is also supported with the help of so-called *smart contracts*, which can be regarded as program codes executed on the blockchain. However, the scalability issue still limits the deployment of blockchain, due to the inherent consensus mechanism to achieve a consistent view on transactions among all peers. For example, Bitcoin processes tens of transactions per second and requires around one hour to confirm a transaction.

A novel solution, called *payment channels*, has been introduced to address the scalability challenge. A payment channel allows two users to deposit funds to a co-maintained ledger and perform payments without broadcasting and recording them on the blockchain. It is guaranteed that users could refund coins in the two-party ledger to the blockchain at any time. Since payments require only approval from channel users rather than all peers of the blockchain, they are conducted quickly and

confirmed immediately, only limited by the communication bandwidth and latency.

Payment channels are extended to the network, in which users sharing no channels can perform multi-hop payments leveraging paths of existing channels. With other users as relays of payments, users do not have to open channels to each other, thereby reducing costs in maintaining channels. Lightning Network (LN) [21], for example, is a practical payment channel network deployed upon Bitcoin.

Users in payment channels face the problem of fund depletion [13]. As payment flows in both directions over a channel are not equal, funds gradually accumulate in the direction of the higher flow. Eventually, the channel becomes depleted and fails payments in that direction, undermining the off-chain network throughput. Current solutions for refunding the depleted channel fall into two categories. The first one is to refund the single channel using on-chain coins. A trivial approach is to close and reopen the channel, which requires two costly and time-consuming on-chain transactions. LOOP (see <https://lightning.engineering/loop/>) reduces the refund costs to one on-chain transaction. Both approaches need interactions with the blockchain each time of channel refunding, which can be regarded as “one-time on-chain execution, one-time refunding”. The second solution is to refund the channel by reallocating deposits in adjacent channels, also called “*rebalance*”. It recovers the channel from depletion without performing transactions on the blockchain, exemplified by Revive [13]. Relying on circular paths in the payment channel network, Revive instructs coins to flow along the path, therefore users' coins are rebalanced from well-funded channels to poorly-funded channels. The whole process is executed off-chain, thus being cost-free for unlimited times of rebalancing and relieving the transaction load on the underlying blockchain. These features make Revive act well in scenarios where users have stable and circular rebalancing relations.

Nevertheless, Revive suffers from low feasibility in large-scale applications, such as LN, because of the lack of *autonomy*. Specifically, to achieve a successful off-chain rebalancing, the following requirements need to be met. (1) For users with the demand of rebalancing, both the channels they want to rebalance and the direction of desired coin flows can form directed cycles in the payment channel network. (2) A fair leader is necessary to collect users' rebalancing demands, find the directed cycles, and generate the rebalancing transactions for cycles justly. (3) Once a rebalancing cycle is established, all cycle users would cooperate with the rebalancing procedure. If any of these requirements are not satisfied, Revive will fail. To make matters worse, even if the rebalancing succeeds, only the minimum rebalancing amount among the cycle users could be

* Corresponding authors.

TABLE I: Comparison of channel refunding solutions.

	# On-Chain TXs for n Times Refunding	Wait Time for 1 Time Refunding	Leader	Bitcoin-compatible	Usability		
					Non-cycle Users		Cycle Users
					Single-channel Users	Multiple-channel Users	
Close-and-Reopen	$2n$	2Δ	No need	Yes	●	●	●
LOOP	n	Δ	No need	Yes	●	●	●
Revive	0	instantly ¹	Fair ²	Yes ³	○	○	◐
Shaduf	1 ⁴	instantly ¹	No need	No	○ ⁵	●	●

Δ denotes the blockchain delay, which is the maximum time required for one transaction to be recorded on the blockchain.

○: unusable; ◐: usable when requirements are met; ●: usable.

¹ Instantly when involved users cooperate. In Revive, it includes users along the cycle. In Shaduf, it includes two adjacent users.

² The leader treats users' rebalancing requests equally.

³ With the atomic multi-channel update protocol introduced in [9], Revive could be deployed to Bitcoin-compatible platforms.

⁴ For one single pair of channels, Shaduf achieves 1-time on-chain binding, n -times off-chain refunding.

⁵ Although the single-channel users cannot be off-chain rebalanced, their channels could help other users to rebalance. In the LN snapshot used in the evaluation (Section VI-B), the single-channel users account for around 45%.

achieved. In other words, one user's rebalanced amounts are limited by the amounts of other users in the cycle.

Contributions and roadmap. This work addresses the aforementioned shortcomings with a protocol we call *Shaduf*, which overcomes the limitations of Revive and can be applied to more general rebalancing scenarios. Our contributions include,

- *Novel solution for payment channel rebalancing.* We introduce Shaduf, a non-cycle off-chain channel rebalancing scheme, which allows users to perform unlimited times of off-chain coin shifts between channels after *one-time on-chain binding*. Compared with Revive, Shaduf does not rely on cycles in the network. Even better, since only the user's own channels and adjacent users are involved, Shaduf is resistant to uncooperative participants and not restricted by others' demands. The comparison between current approaches and Shaduf is shown in Table I. An additional benefit brought by the autonomy of rebalancing is that Shaduf reduces the deposits of channel users and maintains the payment amount (Section III and V).
- *Formal security proof.* We formalize the security properties of Shaduf using the ideal/real world paradigm in the Universal Composability (UC) framework and provide the formal security proof according to our definition (Section IV and V, and Appendix B).
- *Implementation and valid evaluation.* We provide a proof-of-concept implementation of Shaduf on Ethereum to show the required one-time binding is cheap compared with the unlimited times of rebalancing brought by it. Besides, an experimental evaluation demonstrates that Shaduf outperforms LN and Revive. Particularly, under different application strategies, Shaduf increases by 6% - 28% in terms of payment success ratio and volume, and reduces users' deposits by 37% - 75% while maintaining the same network performance (Section VI).

II. PRELIMINARIES

A. Blockchain and Smart Contracts

The blockchain is an append-only transaction ledger and maintains users' balances. Let Δ be the blockchain delay, i.e., the time for a transaction to be recorded on the blockchain.

A smart contract is a piece of code executed on the blockchain, receiving coins and handling them according to predefined rules when triggered by users or other contracts.

Our scheme is constructed upon Turing-complete blockchains such as Ethereum.

B. Payment Channel

There are three phases in the lifetime of a payment channel: open, update, and close, as shown in Figure 1.

1) *Open:* Two users, Alice and Bob, open a channel by depositing funds into it. Firstly, Alice sends a coins to the contract as her deposits in the channel. Upon receiving the coins, the contract informs Bob of the opening event and waits Δ time for his confirmation. If Bob confirms, the channel is opened and the sent along b coins serve as his deposits in the channel. Otherwise, Alice retrieves the coins she deposited.

2) *Update:* Two users update the channel state in pace with off-chain payments. The channel state includes the allocation of the channel capacity, i.e., a total of $a + b$ coins, between Alice and Bob, and a version that indicates the number of updates and increases by 1 each update. A state is valid when it is confirmed by both users, that is, users have exchanged their signatures on the state.

Any user can initiate the update. For instance, when Alice performs a payment, she increases the channel version, updates the balance allocation, and sends the updated state to Bob along with her signature on it. If Bob approves the state and replies with his signature on it, the state becomes valid. The process is completed without interacting with the blockchain.

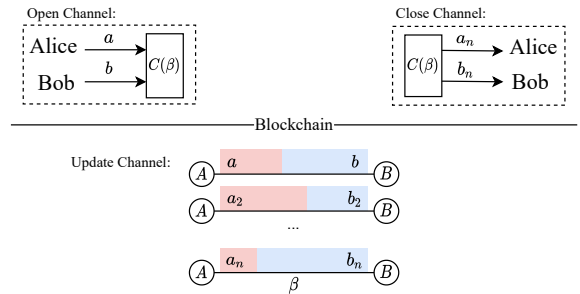


Fig. 1: Three phases of a payment channel, exemplified by β whose users are Alice and Bob. Dashed boxes indicate on-chain transactions and colored boxes represent the balances of users in the channel. Two users first escrow a and b coins to the contract $C(\beta)$ respectively, thereby opening the channel. After an arbitrary number of payments, they close the channel, and the latest balances, a_n and b_n , are refunded.




Fig. 2: Two forms of multi-hop payments in the payment channel network: payment via a path and via a hub. Lines between users denote channels.

3) *Close*: Users could close the channel and withdraw their balances to the blockchain. In particular, it happens when users enter into a dispute, such as when Alice receives no response after sending the updated state to Bob. The user who wants to close the channel, suppose Alice, submits the request to contract along with the latest valid channel state in her view. The contract informs Bob of the closing and waits Δ for his latest channel state. Comparing received valid states, the contract decides on the more recent one, i.e., the state with a higher version, and dispenses funds accordingly.

C. Payment Channel Network

The payment channel network (PCN) is constituted by payment channels, where users could perform payments taking others as intermediaries. Figure 2a presents a typical payment channel network structure, where channels are established between users who perform payments frequently. For example, C opens channels with B , D , and E . The users who share no channels, such as A and E , can perform payments via a multi-hop path connecting them. In this figure, A could pay E via paths $A-B-C-E$ and $A-D-C-E$. Here we ignore mechanisms guaranteeing the security of the multi-hop payments and only consider the path. When all users along the path have sufficient balances, the payment can be completed.

However, it is non-trivial to find a fund-sufficient path between users, because the balances of users in a channel are invisible to others. Payment hub relieves the burden in finding paths, as shown in Figure 2b. Users setup channels with the hub, which in turn links any pair of users. Although the path is easy to be settled, H needs to escrow a huge deposit to channels to afford payments between these users, impeding its deployment in large-scale networks.

D. Rebalancing in PCN

Payment channels face the challenge of fund depletion, where funds in the channel (i.e., channel capacity) flow to one user. It reduces a channel from bi-directional to unidirectional. Refunding the depleted channel using on-chain coins needs interactions with the blockchain, which is costly and requires at least Δ time wait before the channel is refunded. Considering the case where one user serves in multiple channels, Revive [13] proposes to refund a channel leveraging coins in adjacent channels without interactions with the blockchain.

We illustrate how Revive works by taking the directed cycle $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$ in Figure 2a as an example. Assume each user has adequate coins in the channel with the next user and lacks funds in the channel with the previous user. Therefore,

each user has the demand to rebalance coins between its two channels. Revive performs the coin flow along the directed cycle, where each user sends coins to the next user and receives coins from the previous one. To enable the cooperation of cycle participants, only the minimum rebalancing demand of them can be satisfied. After that, for each user, there are balanced coins in its adjacent channels and the sum of coins in the two channels remains unchanged.

Compared with Close-and-Reopen and LOOP, while Revive is attractive for the cost-free feature and unlimited times of rebalancing, it suffers from the lack of autonomy. Firstly, it is only possible when users with the rebalancing demands could form directed cycles in the network. Users, E , F , G in Figure 2a and all users in Figure 2b, can neither take the solution nor facilitate others' rebalancing. Users, A , B , C , D in Figure 2a could be rebalanced only when all of them have the rebalancing demands in the consistent direction. Secondly, the achievement and amount of rebalancing are both restricted by the cycle participants, including the non-adjacent ones. Therefore, Revive's application in PCN is limited.


III. SHADUF OVERVIEW

While current solutions in rebalancing channels rely on a set of participants who can form a cycle in the network, Shaduf allows one user to shift coins between adjacent channels in a non-cycle setting. In a nutshell, after one-time on-chain binding, a user can shift coins between two *bound channels* directly, for unlimited times and in both directions.


In this section, we first demonstrate the basic construction and introduce the one-time channel binding. Then we describe the rebalancing procedures and extend the structure in the payment channel network. The formal and detailed description of Shaduf is presented in Section IV and V.

A. Basic Construction Idea

Consider the basic structure where one user opens two channels with different users, as shown below.



User I setup channel β with L , and channel γ with R . Currently, I owns sufficient coins, say 18 coins, in β and 0 coins in γ . Since the money has been run out, I cannot initiate payments in γ . In Shaduf, I refunds it leveraging coins in β , depicted in the below example.



I shifts 5 coins from β to γ , with which I could perform further payments in γ . The coin shift is completed off-chain between the two channels directly, without the involvement of the blockchain or other participants. Consequently, the rebalancing is executed at no cost and of high efficiency. Moreover, it can be extended to multiple times of coin shift in both directions. For example, I can shift coins from β to γ and shift them back later then.

During the process, measures should be taken to guarantee the balance security, i.e., honest users would not lose money [7], [16]. For instance, to guarantee that there are sufficient coins to be shifted out of one channel and these coins are accepted in another channel for further payments, each coin shift needs approval from the other two users. In addition, to guarantee the off-chain rebalancing takes effect, the off-chain shifted coins should be on-chain delivered at the time of claim.

B. Necessity of One-Time Channel Binding

Simple as it seems, the basic scheme suffers from the “double-shifting attack”, where I colludes the other channel user and shifts coins to multiple channels. For example, I shifts 5 coins from β to γ and 21 coins to another channel η , which is composed of I and P . A total of 26 coins are shifted out of β , while there are only 25 coins in it, leading to 1 coin loss at the time of claim in either γ or η .

Different from other off-chain operations that only concern the balance allocation within one channel, the coin shifts in our scheme change the funds reserved in the two channels. More concretely, fund distribution between the two channels. To solve this problem, it needs to be declared on the blockchain before performing off-chain coin shifts that funds in β and γ can be redistributed along with I 's coin shifts between them. The process is called “bind”, whose introduction maintains the low cost of basic construction and guarantees the balance security. Two channels are bound when the bind is recorded on the blockchain.

After the one-time bind, multiple times of off-chain coin shifts in both directions are allowed. If any user wants to claim the shifted coins, he could request the blockchain to “unbind” the two channels. We refer to Section V for the detailed description of the bind and unbind processes.

C. Rebalancing Procedures after Binding

After the channel binding, unlimited times of rebalancing between two bound channels are enabled. In this way, a record of coin shifts, called bind state, needs to be maintained by the three users. Each time performing the coin shift, users need to update the bind state and channel state, as shown in Figure 3.

Before describing the two phases, we first clarify the contents of the bind state and channel state. Besides the resulting coin shift between channels, two parts are included in a bind state: the bind version which is increased by 1 each shift to indicate the number of shifts, and the latest coin shift for the security consideration which is explained later. In the channel state, besides the channel version and balances of users, the bind version is also included to indicate the latest bind state, since the shifted coins are also part of the channel funds.

In the first phase, three users achieve agreement on the updated bind state, which is initiated by I . Assume that I wants to shift c coins from β to γ . I updates the bind state, including recalculating the total coin shift, recording this shift as the latest one, and increasing the bind version by 1. Then I sends the new state to L and R for their approval (Step (1)). From the view of L , it needs to be verified that I has more than c coins, while R needs to ensure that there are no less than c coins in β currently. Since the channel capacity is public and

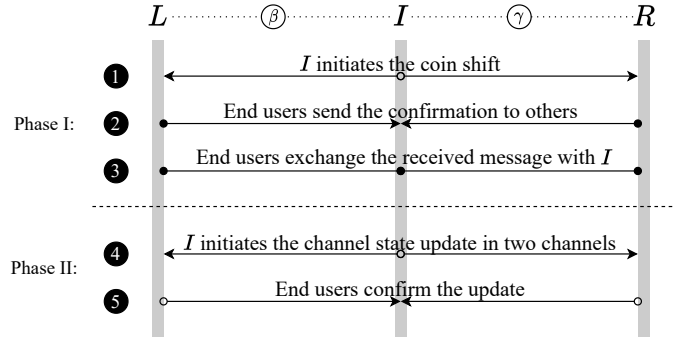


Fig. 3: Flowchart of rebalancing between bound channels β and γ , which are composed by users $\{L, I\}$ and $\{I, R\}$, respectively. The first phase is bind state update, initiated by I . The second phase is channel state update, initiated by I in two channels separately. For a solid line, the arrow indicates the recipient of the message, and the hollow circle indicates the sender. For users who both send and receive messages at a step, they are represented by solid circles.

the previous coins shifts are known, R could do the verification locally. It is symmetrical when coins are shifted in the opposite direction. When verification is passed, L and R reply to I with signatures on the updated bind state. In addition, they forward it to each other to inform that they acknowledged the update (Step (2)). To avoid that L or R send different messages to I and the other one, leading to an inconsistent view on whether the update is confirmed, I exchanges the received message with each other (Step (3)). Users proceed to the second phase when the new bind state is approved by all of them.

In the second phase, users update the state of each channel. Especially, the bind version and I 's balance recorded in the channel state need to be recalculated. I initiates the channel state update in the two channels separately (Step (4)). In the updated channel state, I 's balance is increased by c in γ while it is decreased by c in β , the channel version and the bind version are all increased by 1. Receiving the correctly updated channel state, the other user replies with the approval and completes the rebalancing (Step (5)).

Now we discuss cases where the process halts and what an honest user, denoted as P , should do. When it halts in the first phase, if P has not sent the signature on the updated bind state to others, the halt has no influence since the state is invalid yet. If P has sent the signature but not received the other's signature, then it is not sure whether the coin shift takes effect. At this point, P initiates a dispute on the blockchain, which informs the other users to submit the latest bind state in their view and decides the one with the highest version. From the state, P knows whether the coin shift is effective. When it halts in the second phase, it means that the bind state has been updated but the channel state update is refused by the other one in the channel. At that time, P sends the request of state update to the blockchain along with the latest bind state (if it has not been submitted) and channel state. The blockchain would verify that the bind version included in the bind state is 1 greater than that included in the channel state, indicating I 's balance has not been updated. Then the blockchain updates it according to the last coin shift recorded in the bind state, completing the rebalancing.




Fig. 4: Extension in payment channel network. Colored boxes represent coins of users and the length corresponds to the value. The dashed boxes upon a channel denote the bound amounts and \bowtie connects bound channels. B binds b_β coins in β to b_γ coins in γ , both values are higher than B 's balances in the channels, denoted as b_1 and b_2 , respectively. C binds c_γ coins in γ to c_η coins in η , the former value is lower than C 's balance c_1 in γ and the latter is higher than c_2 in η .

D. Extension in PCN

The basic structure introduced above has limited application in the payment channel network. On the one hand, among the three users of the two bound channels, only the intermediate can shift coins in the two channels, while the other two users cannot send or receive their coins to other channels. On the other hand, the intermediate can only shift coins between the two channels and cannot interact with a third channel. The above limitations are caused by the fact that one channel can only be bound to another channel. To relieve these limitations, one channel needs to be allowed to bind multiple channels, taking either end user as the intermediate.

Recall in the basic structure, it is required that the upper limit of coins that can be shifted out of a channel is its capacity. If the upper limit is set to be lower than the capacity, the remaining amount of coins can be shifted to other channels. In this way, one channel can be bound to multiple channels. When two channels are bound, the upper limit of funds that could be shifted to each other should be specified. In addition, when one channel is bound to multiple channels, the sum of coin shift limits to each channel should be no greater than the channel capacity to guarantee the coins can eventually be delivered. In the following, we demonstrate the extension in two kinds of payment in the off-chain network, payments via a path and a hub.

1) *Payment via a path*: Users without direct channels can conduct off-chain payments leveraging other users as intermediates. Figure 4 presents an example where A and D conduct payments through B and C . Currently, B has b_1 coins in β and b_2 coins in γ , and C has c_1 coins in γ and c_2 coins in η . Assuming that A has sufficient coins in β , the coins that can be paid by A to D is $\min(b_2, c_2)$.

The bind between the three channels are as follows: B binds b_β coins of β to b_γ coins of γ , C binds c_γ coins of γ to c_η coins in η . We first analyze the bind between channels, taking γ as an example. γ is bound to β taking B as the intermediate and b_γ as the upper limit of coins that can be shifted out. Although the upper limit is b_γ , B has only b_2 coins, then B can shift at most $\min(b_2, b_\gamma) = b_2$ coins out of γ . Here we assume there are no previous shifts. If there are, such as a total of e_β coins from β to γ , or e_γ coins from γ to β , the coins that can be shifted are $\min(b_2, b_\gamma + e_\beta)$ or $\min(b_2, b_\gamma - e_\gamma)$. It holds the same for coin shift from γ to η . Besides, due to the fact that the sum of the upper limits $b_\gamma + c_\gamma$ is less than the capacity $b_2 + c_1$, the remaining quota can be utilized in binding other channels.




Fig. 5: Extension in payment hub network. C escrows c coins in β , where c_1 coins are for payments to M_1 and c_2 coins are for M_2 . Out of the c coins, H binds c_1 coins to γ_1 , and c_2 coins to γ_2 . To deliver these payments, H escrows e_1 coins to γ_1 and e_2 coins to γ_2 , where e_1 and e_2 are the maximum amount of a single payment to M_1 and M_2 , respectively.

Since B can shift at most b_1 coins from β to γ and C can shift at most $\min(c_1, c_\gamma) = c_\gamma$ coins from γ to η , the amount that A could pay D is increased to $\min(b_2 + b_1, c_2 + c_\gamma)$. Correspondingly, the amount that D can pay A increase from $\min(c_1, b_1)$ to $\min(c_1 + c_2, b_1 + b_2)$. To sum up, with the same amount of deposits in channels, a path with bind facilitates more payments. In other words, users can bear the same amount of payments with fewer deposits.

2) *Payment via a hub*: Users could conduct payments through the payment hub. Here we focus on its typical and important scenarios [18] such as point-of-sale payments and retail shopping. Serving as one solution, the payment hub categorizes connected users into consumers and merchants. Figure 5 is a simplified situation, composed by consumer C , hub H , and merchants M_1 and M_2 . The payment flow is unidirectional, from C to M_1 and M_2 . C escrows c coins to β , where a total of c_1 coins are used to pay M_1 and c_2 coins are for M_2 . Each time C initiates a payment to the merchant, say M_1 , H receives the coins in β and sends the same amount of coins to M_1 in γ_1 . To accommodate these payments, H needs to escrow e_1 coins in γ_1 and e_2 coins in γ_2 . In summary, for each user in the network, H needs to escrow the same amount of coins in the network.

With binds between these channels, the required deposits of H can be sharply decreased. H binds c_1 coins of β to γ_1 and c_2 coins to γ_2 . At the same time, H escrows e_1 coins to γ_1 and e_2 coins to γ_2 , where e_1 and e_2 are the maximum amount of a single payment to M_1 and M_2 , respectively. Once a payment from C to a merchant, say M_1 , completes, H shifts the received coins in β to γ_1 . In this way, there are always e_1 coins in γ_1 , which are sufficient for further payments. The deposits H needs to escrow in the network are decreased from c to $e_1 + e_2$. The above solution can be naturally extended to cases with multiple customers and merchants connecting to the same hub. When the number of customers and merchants is large, such as thousands or tens of thousands, the bind is more effective in reducing H 's deposits in the network.

It is worth noting that the above analysis also demonstrates that single-channel users in Shaduf, e.g. A in Figure 4 and C in Figure 5, could help others, i.e. B and H , to achieve coin shift. This is attractive since the former constitutes a large portion of LN users, e.g. 45 percent in the snapshot.

IV. FORMAL DEFINITION OF SHADUF

A. The Security Model

We formalize Shaduf utilizing the universal composability (UC) [4] framework and deploy the version with global setup (GUC) [5]. The model is defined over a fixed set of parties $\mathcal{P} = \{P_1, \dots, P_n\}$, taking the underlying blockchain as a global ideal functionality \mathcal{L} with maximal blockchain delay Δ . It specifies two worlds, the real world and the ideal world. In the former world, the Shaduf protocol is executed by parties, facing an adversary \mathcal{A} and interacting with the contract functionality \mathcal{F} . In the latter one, the idealized functionality \mathcal{F} is executed through interactions with parties and the simulator \mathcal{S} which simulates behaviors of the adversary. All parties receive inputs from the environment \mathcal{Z} and send outputs to it. Security is defined as the computational indistinguishability between the outputs of \mathcal{Z} running in the real world and in the ideal world.

1) *Adversary model:* We consider an adversary \mathcal{A} who can corrupt arbitrary parties at the beginning of the protocol execution. Corruption means that \mathcal{A} gets all internal states of corrupted parties and takes full control of them.

2) *Security goals:* Motivated by [7], [16], our main goal is to guarantee the balance security of honest users, i.e., honest users would not lose money. Intuitively, it includes two parts. Firstly, the channel opening and state update are only valid when both channel users agree. The channel binding and bind state update also need the approval of three users in two bound channels. Secondly, it requires that the channel unbinding and closing could be completed within a reasonable time, and when they are settled, the shifted coins between bound channels and refunded balances from a channel is the latest in honest users' view. Properties to be achieved are listed as follows.

Consensus on channel opening and channel update. Within a channel, the opening and state update can only be executed when both users agree. Channel opening takes time $O(\Delta)$ to achieve the agreement. When both users are honest, time spent on channel state update is constant.

Consensus on channel binding and bind update. Between two channels, the binding and state update takes effect only when all involved three users agree. Channel binding takes time $O(\Delta)$. Bind state update takes constant time that is independent of the blockchain delay when all users behave honestly.

Guaranteed channel unbinding. For two bound channels, any of the three users can request to unbind and settle the coin shift between them. The process is guaranteed to be completed within time $O(\Delta)$ after the request.

Guaranteed channel closing. Any user in a channel can request to close the channel at any time. The channel is closed in $O(\Delta)$ time once the request is made.

Guaranteed balance payout. The channel unbinding and closing are settled following the latest states from the view of honest users.

3) *Communication network:* We assume a synchronous communication network and each party is aware of the current time, which is measured in rounds. The delivery of a message

between two parties takes 1 round. It means that one message sent in round r reaches the recipient in the beginning of round $r+1$. In addition, parties communicate via authenticated channels where the recipient can confirm the message source. The adversary can see the message contents but cannot modify or drop the messages. Other communications, e.g., with the environment, take 0 rounds. For simplicity, computations also take 0 rounds.

4) *Ledger and contract functionalities:* We take the formalization of the underlying blockchain from [7]. It is formalized as a global functionality \mathcal{L} with respect to the blockchain delay Δ , denoted as $\mathcal{L}(\Delta)$. The functionality is shown below,

Ledger Functionality
<p>Ledger initialization: Upon receiving $(x_1, \dots, x_n) \in \mathbb{R}_{\geq 0}^n$ from the environment \mathcal{Z}, store the tuple.</p>
<p>Adding coins: Upon receiving (add, P_i, y), if $P_i \in \mathcal{P}$ and $y \in \mathbb{R}_{\geq 0}$ then set $x_i := x_i + y$, else do nothing.</p>
<p>Removing coins: Upon receiving (remove, P_i, y), if $P_i \in \mathcal{P}$ and $y \in \mathbb{R}_{\geq 0}$ and $x_i \geq y$ then set $x_i := x_i - y$, else do nothing.</p>

The value x_i in tuple (x_1, \dots, x_n) recorded in the ledger indicates P_i 's balance and can be accessed by all parties, the adversary, and the environment. The “add” and “remove” operations are performed by the ideal functionality \mathcal{F} in the ideal world, and the contract functionality \mathcal{C} in the real world.

We formalize the contract functionality \mathcal{C} maintaining contract instances, leading to \mathcal{C} -hybrid real world. Each contract is identified by its corresponding channel and $\mathcal{C}(\beta)$ denotes the contract of channel β . Receiving messages from parties, \mathcal{C} accesses the ledger according to predefined procedures.

5) *Security definition:* Let λ be the security parameter. With respect to the global ledger $\mathcal{L}(\Delta)$, denote the output of the environment \mathcal{Z} when interacting with the adversary \mathcal{A} and the protocol π in the \mathcal{C} -hybrid world as $\text{EXEC}_{\pi, \mathcal{Z}, \mathcal{A}}^{\mathcal{L}(\Delta), \mathcal{C}}(\lambda)$, and \mathcal{Z} 's output when interacting with the ideal functionality \mathcal{F} and the simulator \mathcal{S} as $\text{EXEC}_{\mathcal{F}, \mathcal{Z}, \mathcal{S}}^{\mathcal{L}(\Delta)}(\lambda)$. The security is formally defined as follows.

Definition 1: Protocol π running in the \mathcal{C} -hybrid world UC-realizes the ideal functionality \mathcal{F} with respect to a global ledger $\mathcal{L}(\Delta)$, if for any PPT adversary \mathcal{A} there exists a simulator \mathcal{S} such that

$$\text{EXEC}_{\pi, \mathcal{Z}, \mathcal{A}}^{\mathcal{L}(\Delta), \mathcal{C}}(\lambda) \approx \text{EXEC}_{\mathcal{F}, \mathcal{Z}, \mathcal{S}}^{\mathcal{L}(\Delta)}(\lambda)$$

where \approx denotes the computational indistinguishability.

B. Notation

The following notations are used to define the attributes of the channel and the bind, and the communication between parties. They are also used in the real world protocol description.

1) *Channel:* A payment channel β is defined by the attribute tuple $(\beta.\text{id}, \beta.\text{users}, \beta.\text{balance}, \beta.\text{fund}, \beta.\text{bind}, \beta.\text{allot})$. Regarding the attributes, $\beta.\text{id} \in \{0, 1\}^*$ is the channel identifier, $\beta.\text{users} = \{\beta.\text{Alice}, \beta.\text{Bob}\}$ is the set of channel users, $\beta.\text{balance} : \beta.\text{users} \rightarrow \mathbb{R}_{\geq 0}$ denotes the balances of users, $\beta.\text{fund} \in \mathbb{R}_{\geq 0}$ is the deposits in the channel, $\beta.\text{bind}$

records all binds of the channel, and $\beta.\text{allot}$ is the remaining amount of coins for binding other channels. In addition, we define the function $\beta.\text{other-user} : \beta.\text{users} \rightarrow \beta.\text{users}$ which maps one user to another one. Specifically, $\beta.\text{other-user}(\beta.\text{Alice}) = \beta.\text{Bob}$ and $\beta.\text{other-user}(\beta.\text{Bob}) = \beta.\text{Alice}$.

A payment θ in the channel is defined as $\beta.\text{users} \rightarrow \mathbb{R}$ and $\theta(\beta.\text{Alice}) + \theta(\beta.\text{Bob}) = 0$. Applying payment θ to the current channel balance. i.e., $\beta.\text{balance} + \theta$, derives balances of users after the payment.

2) *Bind*: A bind Φ is defined by the attribute tuple $(\Phi.\text{id}, \Phi.\text{users}, \Phi.\text{channels}, \Phi.\text{shift}, \Phi.\text{reserve})$. About the attributes, $\Phi.\text{id} \in \{0, 1\}^*$ is the bind identifier, and $\Phi.\text{users} = (\Phi.\text{Lara}, \Phi.\text{Irene}, \Phi.\text{Rock})$ is the set of users, where $\Phi.\text{Irene}$ is the intermediate. $\Phi.\text{channels} = (\Phi.\text{c}_a, \Phi.\text{c}_p)$ is the set of two bound channels, and $\Phi.\text{c}_a$ consists of users $\Phi.\text{Lara}$ and $\Phi.\text{Irene}$, while $\Phi.\text{c}_p$ takes $\Phi.\text{Irene}$ and $\Phi.\text{Rock}$ as its users. $\Phi.\text{shift} : \Phi.\text{channels} \rightarrow \mathbb{R}$ denotes the total shifted coins between the channels and $\Phi.\text{shift}(\Phi.\text{c}_a) + \Phi.\text{shift}(\Phi.\text{c}_p) = 0$. For example, if $\Phi.\text{shift}(\Phi.\text{c}_a) = c > 0$, it indicates c coins are shifted to $\Phi.\text{c}_a$ from $\Phi.\text{c}_p$. $\Phi.\text{reserve} : \Phi.\text{channels} \rightarrow \mathbb{R}^{>0}$ indicates the upper limit of coin shift in each channel. When $\Phi.\text{reserve}(\Phi.\text{c}_a) = c$, it means that at most c coins can be shifted from $\Phi.\text{c}_a$ to $\Phi.\text{c}_p$. Furthermore, we define $\Phi.\text{end-users} = \{\Phi.\text{Lara}, \Phi.\text{Rock}\}$ as the set of end users, and the function $\Phi.\text{other-users}$ which maps one user to other two users. To cite an example, $\Phi.\text{other-users}(\Phi.\text{Lara}) = \{\Phi.\text{Irene}, \Phi.\text{Rock}\}$.

One coin shift $\tilde{\theta}$ between two channels is defined as $\Phi.\text{channels} \rightarrow \mathbb{R}$ and $\tilde{\theta}(\Phi.\text{c}_a) + \tilde{\theta}(\Phi.\text{c}_p) = 0$. When $\tilde{\theta}$ is executed, the resulting coin shift is updated to $\Phi.\text{shift} + \tilde{\theta}$.

3) *Communication*: We write the event that “send message m to P in round r ” as “ $m \xrightarrow{r} P$ ” and “receive message m from P in round r ” as “ $m \xleftarrow{r} P$ ”, where P could be a party or a contract. Especially, when P is a set of parties, the above presentations denote sending or receiving the message to or from each party in the set. Moreover, when the time is termed as $r \leq r_1$, it means that the event is happened in round r which is decided by the adversary but r_1 is the upper bound.

C. Ideal Functionality of Shaduf

The ideal functionality \mathcal{F} communicates with parties in the set \mathcal{P} , the simulator \mathcal{S} , and the ledger \mathcal{L} . All existing channels are recorded in the set Γ . Before describing the procedures of \mathcal{F} , we discuss the restrictions on instructions from the environment. We restrict that the environment would not give instructions that are apparently wrong. For example, the environment would never request a user to open a channel that already exists or when the user has insufficient money for the deposit. These restrictions are used to simplify the protocol description and can be easily eliminated by adding determinations on instructions from the environment. The full restrictions are listed in Appendix A.

There are 6 procedures in \mathcal{F} , as shown in Figure 6. Each procedure is triggered by one message from parties, including the procedure name and an identifier referring to the channel or the bind. Procedure (A) *Open* handles the opening of channel β when triggered by $\beta.\text{Alice}$. (B) *Channel Update* performs the state update, or in other words, the payment,

within the channel when receiving a request from one user and confirmed by another one. (C) *Bind* processes the bind of two channels and is triggered by the intermediate. (D) *Bind Update* executes the coin shift between two bound channels when receiving requests from the intermediate and agreement from the end users. (E) *Unbind* resolves the bind of two channels when triggered by any user in them and completes the final coin shift between them. (F) *Close* is triggered by any channel user and refunds balances on the blockchain.

Now we analyze the achievement of desired security goals.

Consensus on channel opening and channel update. A channel can only be opened when Alice applies and Bob approves it. This means that two users have achieved agreement on channel opening. From the time interval between the message (*opened*) and the opening request, it is easy to see the opening takes time $O(\Delta)$. Similarly, the channel state update is completed when both users agree. When both users are honest, it takes 2 rounds to confirm the result.

Consensus on channel binding and bind update. The bind of two channels is initiated by the intermediate and confirmed by the end users. Channels are bound only when they all agree, taking $O(\Delta)$ time. Likewise, the bind state is updated when all users confirm. It takes 3 rounds when all users behave honestly.

Guaranteed channel unbinding. The channel unbinding can be triggered by any user of the two bound channels, and it is guaranteed to be completed within 3Δ time.

Guaranteed channel closing. The channel closing can be initiated by any user of the channel. When the channel is not bound to other channels, the time needed to close the channel is at most 3Δ . Otherwise, the time is at most 5Δ , among which 3Δ is used to unbind other channels.

Guaranteed balance payout. When two channels are unbound, the coin shift is settled according to its latest bind state ($\Phi.\text{shift}$ in Figure 6 (E)). Similarly, when the channel is closed, the latest balance allocation between two end users is put on the blockchain ($\beta.\text{balance}$ in Figure 6 (F)).

V. DETAILED DESCRIPTION OF SHADUF

In this section, we give a detailed description of Shaduf protocol. Before the description, we first introduce the supplementary notations on bind and channel. The protocol is shown in Figure 7 and 8, and the contract functionality is shown in Figure 9. Finally, we provide the security theorem.

A. Notation

Recall in Section III-C, parties maintain states of the bind and the channel to perform off-chain coin shifts and payments. In the formal description, the states are presented as follows,

1) *Bind*: The state of bind Φ is defined as $\Phi.\text{state} = (\Phi.\text{version}, \Phi.\text{shift}, \Phi.\text{laShift}, \Phi.\text{sig})$, where $\Phi.\text{version}$ is the shift counter and increments by 1 each shift, $\Phi.\text{shift}$ is the total coin shift and $\Phi.\text{laShift}$ is the latest coin shift between the two bound channels, $\Phi.\text{sig}$ collects signatures from the three users on this state. The state is called valid and can be submitted to the blockchain when $\Phi.\text{sig}$ includes the signatures of all three users. We denote the latest valid state of bind Φ that party P is aware of as $\Phi^P.\text{state}$.

Assume the following messages concerning the channel β and the bind Φ , with their identifiers, $\beta.\text{id}$ and $\Phi.\text{id}$, denoted as id and \tilde{id} , respectively. In the procedures (B) and (D), we denote the requested payment and coin shift as θ and $\tilde{\theta}$, respectively.

(A) **Open**

Upon (open, β) \xleftrightarrow{t} $\beta.\text{Alice}$, (remove, $\beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})$) $\xleftrightarrow{t_1 \leq t + \Delta}$ \mathcal{L} , (opening, β) $\xleftrightarrow{t_1}$ $\beta.\text{Bob}$. If (open) $\xleftrightarrow{t_1}$ $\beta.\text{Bob}$, send (remove, $\beta.\text{Bob}, \beta.\text{balance}(\beta.\text{Bob})$) $\xleftrightarrow{t_2 \leq t_1 + \Delta}$ \mathcal{L} , (opened) $\xleftrightarrow{t_2}$ $\beta.\text{users}$, add β to Γ , and stop. Otherwise, upon (refund, β) $\xleftrightarrow{t_3 > t_1 + \Delta}$ $\beta.\text{Alice}$, send (add, $\beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})$) $\xleftrightarrow{t_4 \leq t_3 + \Delta}$ \mathcal{L} , and stop.

(B) **Channel Update**

Upon (chan-update, id, θ) \xleftrightarrow{t} $P \in \beta.\text{users}$, (chan-update-req, id, θ) $\xleftrightarrow{t+1}$ Q , where $Q := \beta.\text{other-user}(P)$. If (chan-update-ok) $\xleftrightarrow{t+1}$ Q , $\beta.\text{balance} := \beta.\text{balance} + \theta$, (chan-updated) $\xleftrightarrow{t+2}$ P , and stop. Else, if P is honest, denote θ as β 's pending update, execute procedure (F), and stop. Otherwise, stop.

(C) **Bind**

Upon (bind, Φ) \xleftrightarrow{t} $\Phi.\text{Irene}$, if $\Phi.\text{Irene}$ is honest, (bind-req, Φ) $\xleftrightarrow{t+1}$ end-users . Otherwise, for $P \in \Phi.\text{end-users}$: if (send-req, P) \xleftrightarrow{t} \mathcal{S} , (bind-req, Φ) $\xleftrightarrow{t+1}$ P . Do the following.

- 1) If (bind-ok) $\xleftrightarrow{t+1}$ $\Phi.\text{end-users}$, (bind-confirm) $\xleftrightarrow{t+2}$ $\Phi.\text{Irene}$. Else, stop.
- 2) If (bind-confirmed) $\xleftrightarrow{t+2}$ $\Phi.\text{Irene}$, (bound) $\xleftrightarrow{t_1 \leq t+2+\Delta}$ $\Phi.\text{users}$, for $\gamma \in \Phi.\text{channels}$: add Φ to $\gamma.\text{bind}$, $\gamma.\text{allot} = \gamma.\text{allot} - \Phi.\text{reserve}(\gamma)$, and stop. Else, (not-bound) $\xleftrightarrow{t+2+\Delta}$ $\Phi.\text{end-users}$, and stop.

(D) **Bind Update**

Upon (bind-update, $\tilde{id}, \tilde{\theta}$) \xleftrightarrow{t} $\Phi.\text{Irene}$, if $\Phi.\text{Irene}$ is honest, (bind-update-req, $\tilde{id}, \tilde{\theta}$) $\xleftrightarrow{t+1}$ $\Phi.\text{end-users}$. Otherwise, for $P \in \Phi.\text{end-users}$: if (send-req, P) \xleftrightarrow{t} \mathcal{S} , (bind-update-req, $\tilde{id}, \tilde{\theta}$) $\xleftrightarrow{t+1}$ P . Do the following.

- 1) If (bind-update-ok) $\xleftrightarrow{t_1 \leq t+2}$ $\Phi.\text{end-users}$, $\Phi.\text{shift} := \Phi.\text{shift} + \tilde{\theta}$, (bind-updated) $\xleftrightarrow{t+3}$ $\Phi.\text{users}$. For $\gamma \in \Phi.\text{channels}$: $\gamma.\text{balance}(\Phi.\text{Irene}) := \gamma.\text{balance}(\Phi.\text{Irene}) + \tilde{\theta}(\gamma)$. Else, if $\Phi.\text{Irene}$ is honest, or (bind-update-ok) $\xleftrightarrow{t+1}$ $P \in \Phi.\text{end-users}$ and P is honest, denote $\tilde{\theta}$ as Φ 's pending update, go to procedure (E), and stop. Otherwise, stop.
- 2) For $\gamma \in \Phi.\text{channels}$: if (bind-chan-not-updated, $\gamma.\text{id}$) $\xleftrightarrow{t_2 \leq t+4}$ $P \in \gamma.\text{users}$ and P is corrupted, (bind-chan-not-updated, $\gamma.\text{id}$) $\xleftrightarrow{t_2+1}$ Q where $Q := \gamma.\text{other-user}(P)$, execute procedure (F), and stop. Else, stop.

(E) **Unbind**

Upon (unbind, \tilde{id}) \xleftrightarrow{t} $P \in \Phi.\text{users}$, or be invoked by other procedures (let t be the current round), in round $t_1 \leq t + \Delta$, if Φ 's pending update $\tilde{\theta}$ is confirmed, $\Phi.\text{shift} := \Phi.\text{shift} + \tilde{\theta}$, and for $\gamma \in \Phi.\text{channels}$: $\gamma.\text{balance}(\Phi.\text{Irene}) := \gamma.\text{balance}(\Phi.\text{Irene}) + \tilde{\theta}(\gamma)$. (unbound) $\xleftrightarrow{t_2 \leq t_1 + 2\Delta}$ $\Phi.\text{users}$. For $\gamma \in \Phi.\text{channels}$, do the following.

- 1) $\gamma.\text{fund} := \gamma.\text{fund} + \Phi.\text{shift}(\gamma)$, $\gamma.\text{allot} = \gamma.\text{allot} + \Phi.\text{shift}(\gamma) + \Phi.\text{reserve}(\gamma)$, and remove Φ from $\gamma.\text{bind}$.
- 2) If (bind-chan-not-updated, $\gamma.\text{id}$) $\xleftrightarrow{t_3 \leq t_2 + 1}$ $P \in \gamma.\text{users}$ and P is corrupted, send (bind-chan-not-updated, $\gamma.\text{id}$) $\xleftrightarrow{t_3+1}$ Q where $Q := \gamma.\text{other-user}(P)$, execute procedure (F) and stop. Else, stop.

(F) **Close**

Upon (close, id) \xleftrightarrow{t} $P \in \beta.\text{users}$, or be invoked by other procedures (let t be the current round), in round $t_1 \leq t + \Delta$, if β 's pending update θ is confirmed, $\beta.\text{balance} := \beta.\text{balance} + \theta$. Do the following.

- 1) If $\beta.\text{bind} = \emptyset$, (add, $\beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})$) $\xleftrightarrow{t_2 \leq t_1 + 2\Delta}$ \mathcal{L} , (add, $\beta.\text{Bob}, \beta.\text{balance}(\beta.\text{Bob})$) $\xleftrightarrow{t_2}$ \mathcal{L} , (closed) $\xleftrightarrow{t_2}$ $\beta.\text{users}$, remove β from Γ , and stop.
- 2) Otherwise, for $\Phi \in \beta.\text{bind}$: execute procedure (E). Let t_3 be the current time. (add, $\beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})$) $\xleftrightarrow{t_4 \leq t_3 + \Delta}$ \mathcal{L} , (add, $\beta.\text{Bob}, \beta.\text{balance}(\beta.\text{Bob})$) $\xleftrightarrow{t_4}$ \mathcal{L} , (closed) $\xleftrightarrow{t_4}$ $\beta.\text{users}$, remove β from Γ , and stop.

Fig. 6: Ideal functionality \mathcal{F}

2) *Channel*: The state of channel β is represented as $\beta.\text{state} = (\beta.\text{version}, \beta.\text{balance}, \beta.\text{biList}, \beta.\text{sig})$. Among the attributes, $\beta.\text{version}$ is the counter of state update and increments by 1 each update, $\beta.\text{balance}$ is balance allocation between two users, $\beta.\text{biList} : \{0, 1\}^* \rightarrow \mathbb{N}^*$ maps the identifier of a bind to its version, and $\beta.\text{sig}$ records signatures of channel users on the state. Similarly, the state is valid when $\beta.\text{sig}$ contains signatures from both the users. And $\beta^P.\text{state}$ represents the latest valid channel state from P 's view.

B. *Procedures of Shaduf*

We start with describing the procedure of channel opening as shown in Figure 7(A), with Figure 9(A) presenting contract actions. $\beta.\text{Alice}$ sends the opening request to the contract instance $\mathcal{C}(\beta)$, which informs $\beta.\text{Bob}$ of the opening event. If $\beta.\text{Bob}$ confirms it within Δ time, the contract marks β as opened and informs (opened) to them (Step (1), Figure 9(A)). Otherwise, $\beta.\text{Alice}$ sends (timeout) to the contract, which returns back the coins (Step (2), Figure 9(A)).

The channel state update is performed off-chain between

Assume the following messages concerning the channel β and the bind Φ . We abbreviate $\Phi.\text{Irene}$ as I . In the procedure (B), we denote the channel identifier as id , the initiator of the payment as P , and the other user as Q .

(A) **Open**

- 1) Upon (open, β) \xleftrightarrow{t} \mathcal{Z} , $\beta.\text{Alice}$ sends (open, β) \xleftrightarrow{t} $\mathcal{C}(\beta)$, and goes to step 3.
- 2) Upon (opening, β) $\xleftrightarrow{\tau}$ $\mathcal{C}(\beta)$, $\beta.\text{Bob}$ sends (opening, β) $\xleftrightarrow{\tau}$ \mathcal{Z} . If (open) $\xleftrightarrow{\tau}$ \mathcal{Z} , he sends (open) $\xleftrightarrow{\tau}$ $\mathcal{C}(\beta)$, outputs (opened) upon receiving it from $\mathcal{C}(\beta)$, adds β to Γ^B , and waits for messages concerning β . Otherwise, he stops.
- 3) If (opened) $\xleftrightarrow{t_1 \leq t+2\Delta}$ $\mathcal{C}(\beta)$, $\beta.\text{Alice}$ adds β to Γ^A , outputs (opened), and waits for messages concerning β . Otherwise, upon receiving (refund, β) from \mathcal{Z} , she sends (timeout) to $\mathcal{C}(\beta)$, and stops.

(B) **Channel Update**

- 1) Upon (chan-update, id, θ) \xleftrightarrow{t} \mathcal{Z} , P generates signature σ_P on $msg = (\beta^P.\text{version} + 1, \beta^P.\text{balance} + \theta, \beta^P.\text{biList})$, sends (chan-update, msg, σ_P) \xleftrightarrow{t} Q , and goes to step 3.
- 2) Upon (chan-update, $msg = (ver, \alpha, biList), \sigma_P$) $\xleftrightarrow{\tau}$ P where σ_P is P 's signature on msg , Q proceeds as follows,
 - a) If $ver \neq \beta^Q.\text{version} + 1$, or $biList \neq \beta^Q.\text{biList}$, Q ignores the message and stops.
 - b) Otherwise, Q sends (chan-update-req, $id, \alpha - \beta^Q.\text{balance}$) $\xleftrightarrow{\tau}$ \mathcal{Z} . If (chan-update-ok) $\xleftrightarrow{\tau}$ \mathcal{Z} , Q generates signature σ_Q on msg , (chan-update, σ_Q) $\xleftrightarrow{\tau}$ P , and stops. Else, Q stops.
- 3) If (chan-update, σ_Q) $\xleftrightarrow{t+2}$ Q where σ_Q is Q 's signature on msg , P outputs (chan-updated) and stops. Else, P denotes θ as β 's pending update, initiates channel closing in procedure (F), and stops.

(C) **Bind**

- 1) Upon (bind, Φ) \xleftrightarrow{t} \mathcal{Z} , I generates signature σ_I on (Φ, t) , sends (bind, Φ, t, σ_I) \xleftrightarrow{t} $\Phi.\text{end-users}$, and goes to step 3.
- 2) Upon (bind, $\Phi, \tau - 1, \sigma_I$) $\xleftrightarrow{\tau}$ I where σ_I is I 's signature on $msg = (\Phi, \tau - 1)$, $P \in \Phi.\text{end-users}$ sends (bind-req, Φ) $\xleftrightarrow{\tau}$ \mathcal{Z} . If (bind-ok) $\xleftrightarrow{\tau}$ \mathcal{Z} , P generates signature σ_P on msg , (bind, σ_P) $\xleftrightarrow{\tau}$ I , and goes to step 4. Else, P stops.
- 3) If I receives the end users' signatures on (Φ, t) in $t + 2$, I sends (bind-confirm) $\xleftrightarrow{t+2}$ \mathcal{Z} and if (bind-confirmed) $\xleftrightarrow{t+2}$ \mathcal{Z} , I sends (Φ, t, Σ) $\xleftrightarrow{t+2}$ $\mathcal{C}(\Phi.\text{c}_a)$ where Σ is the set of signatures, outputs (bound) upon receiving it from $\mathcal{C}(\Phi.\text{c}_a)$, and for $\gamma \in \Phi.\text{channels}$, I adds Φ to $\gamma^I.\text{bind}$, sets $\gamma.\text{allot} = \gamma.\text{allot} - \Phi.\text{reserve}(\gamma)$, and stops. Otherwise, I stops.
- 4) Denote P in channel $\gamma \in \Phi.\text{channels}$. If (bound) $\xleftrightarrow{\tau_1 \leq \tau+1+\Delta}$ $\mathcal{C}(\Phi.\text{c}_a)$, P adds Φ to $\gamma^P.\text{bind}$, sets $\gamma.\text{allot} = \gamma.\text{allot} - \Phi.\text{reserve}(\gamma)$, outputs (bound), and stops. Otherwise, P outputs (not-bound) and stops.

Fig. 7: Procedures (A) Open, (B) Channel Update, and (C) Bind.

the channel users. When user P receives the instruction to update the channel state, P increments the update counter and the balance allocation, and sends the updated state to the other user, say Q , along with the signature (Step (1), Figure 7(B)). Q verifies the correctness of the counter and the signature and asks the environment whether the update is supported. If it is, Q replies with the signature on the updated state, which becomes valid then (Step (2), Figure 7(B)). If Q does not respond, P is not sure whether the state is updated. Then P marks the update as pending and initiates channel closing, from which the result of update is derived.

The procedure of bind is shown in Figure 7(C) and the contract operation appears in Steps (1-2) of Figure 9(B). To bind two channels, users need to declare it on the blockchain. It is required that the two contracts record the bind atomically. If one contract confirms the bind while the other one not, the eventual coin shift between the two channels cannot be settled since the contract can only send or receive coins with bound channels.¹ To guarantee the above cases would not happen, the setup of bind in one contract is triggered by another contract, which in turn records the bind only when the triggered contract returns confirmation. Specifically, users send

the bind request to the trigger channel contract, which verifies the request and informs the triggered contract if correct (Step (1a), Figure 9(B)). The triggered contract checks the bind and replies confirmation if correct (Step (2), Figure 9(B)). After that, the trigger contract records the bind and informs users (Step (1a), Figure 9(B)). To specify the character of “trigger” and “triggered” between the two contracts, we fix the contract with a lower channel identifier value as “trigger” and the other one as “triggered”. Their roles are the same when unbinding.

The request of bind Φ is submitted by the intermediate, who needs to collect signatures of the other two users on the request first (Step (1), Figure 7(C)). The request is in the form of (Φ, t) , where t is the timestamp to avoid the replay attack. Besides, the timestamp is used to avoid the long time waiting for end users to confirm the request's effectiveness, especially when the user has replied with the signature on it but the request has not appeared on the blockchain. The timestamp indicates that the request is only valid within time $2 + \Delta$, where 2 rounds are for the intermediate to collect signatures and Δ rounds are for submission. Receiving the request from the intermediate and confirmed by the environment, end users sign the bind request and send the signatures to the intermediate (Step (2), Figure 7(C)). They wait for at most $1 + \Delta$ rounds to confirm whether the bind takes effect (Step (4), Figure 7(C)).

After the bind Φ is recorded on the blockchain, the intermediate can shift coins between the two bound channels. The

¹It is also required that the contracts are secure, which can handle the unbind request and deliver shifted coins to bound channels. It can be implemented by checking whether the contract is of permitted implementations via opcode EXTCODEHASH, see <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1052.md>.

Assume the following messages concerning the channel β and the bind Φ , with their identifiers, $\beta.\text{id}$ and $\Phi.\text{id}$ denoted as id and \tilde{id} , respectively. We denote the requested coin shift in the procedure (D) as $\tilde{\theta}$ and abbreviate $I := \Phi.\text{Irene}$. In the procedures (E) and (F), we denote the initiator as P and the responder as Q .

(D) Bind Update

- 1) Upon $(\text{bind-update}, \tilde{id}, \tilde{\theta}) \xleftarrow{t} \mathcal{Z}$, I generates signature σ_I on $msg = (\Phi^I.\text{version} + 1, \Phi^I.\text{shift} + \tilde{\theta}, \tilde{\theta})$, sends $(\text{bind-update}, msg, \sigma_I) \xrightarrow{t} \Phi.\text{end-users}$, and goes to step 3.
- 2) Upon $(\text{bind-update}, msg = (ver, shift, \tilde{\theta}), \sigma_I) \xleftarrow{\tau} I$ where σ_I is I 's signature on msg , $ver = \Phi^P.\text{version} + 1$, and $shift = \Phi^P.\text{shift} + \tilde{\theta}$, $P \in \Phi.\text{end-users}$ sends $(\text{bind-update-req}, \tilde{id}, \tilde{\theta}) \xrightarrow{\tau} \mathcal{Z}$. If $(\text{bind-update-ok}) \xleftarrow{\tau} \mathcal{Z}$, P generates signature σ_P on msg , $(\text{bind-update}, \sigma_P) \xrightarrow{\tau} \Phi.\text{other-users}(P)$, and goes to step 4. Otherwise, P stops.
- 3) Upon $(\text{bind-update}, \sigma_P) \xleftarrow{t+2} P \in \Phi.\text{end-users}$ where σ_P is P 's signature on msg , I forwards the message to the other end user. In round $t+3$, if I has collected others' signatures, I outputs (bind-updated) , and for $\gamma \in \Phi.\text{channels}$: goes to step 5. Otherwise, I denotes $\tilde{\theta}$ as Φ 's pending update, initiates unbinding in the procedure (E) and stops.
- 4) If $(\text{bind-update}, \sigma_Q) \xleftarrow{\tau+1} Q$ where Q is the other end user and σ_Q is Q 's signature on msg , P sends $(\text{bind-update}, \sigma_Q) \xrightarrow{\tau+1} I$. In round $\tau+2$, if P has collected Q 's signature, P outputs (bind-updated) , and in P 's channel γ : goes to step 6. Otherwise, P denotes $\tilde{\theta}$ as Φ 's pending update, initiates unbinding in procedure (E), and stops.
- 5) Denote the current round as t_1 . I updates her balance by $\tilde{\theta}(\gamma)$ and denotes the updated channel balance as $\gamma^I.\text{balance}'$, adds 1 to $\gamma^I.\text{biList}(\tilde{id})$ (sets to 1 if \tilde{id} is not in the list) and denotes it as $\gamma^I.\text{biList}'$, generates signature σ_I on $msg = (\gamma^I.\text{version} + 1, \gamma^I.\text{balance}', \gamma^I.\text{biList}')$, and sends it to the other channel user. If I receives the other's signature, I stops. Otherwise, I outputs $(\text{bind-chan-not-updated})$, initiates channel closing in procedure (F), and stops.
- 6) Denote the current round as τ_1 . P updates I 's balance by $\tilde{\theta}(\gamma)$ and denotes the updated channel balance as $\gamma^P.\text{balance}'$, adds 1 to $\gamma^P.\text{biList}(\tilde{id})$ (sets to 1 if \tilde{id} is not in the list) and denotes it as $\gamma^P.\text{biList}'$. If P receives I 's signature on $msg = (\gamma^P.\text{version} + 1, \gamma^P.\text{balance}', \gamma^P.\text{biList}')$, P sends the signature on msg to I , and stops. Otherwise, P outputs $(\text{bind-chan-not-updated})$, initiates channel closing in procedure (F), and stops.

(E) Unbind

- 1) Upon $(\text{unbind}, \tilde{id}) \xleftarrow{t} \mathcal{Z}$, P sends $(\text{unbind}, \tilde{id}, \Phi^P.\text{state}) \xrightarrow{t} C(\Phi.\text{C}_a)$, waits for Δ time, and goes to step 3.
- 2) Upon $(\text{unbinding}, \tilde{id}) \xleftarrow{\tau} C(\Phi.\text{C}_a)$, Q sends $(\text{unbind}, \tilde{id}, \Phi^Q.\text{state}) \xrightarrow{\tau} C(\Phi.\text{C}_a)$, and goes to step 3.
- 3) For the channel $\gamma \in \Phi.\text{channels}$ where the party is in, proceed as follows,
 - a) If not received (unbound) from $C(\Phi.\text{C}_a)$ within Δ time, send $(\text{timeout}, \tilde{id})$ to $C(\Phi.\text{C}_a)$.
 - b) Upon receiving (unbound) from $C(\Phi.\text{C}_a)$, remove Φ from $\gamma.\text{bind}$, update $\gamma.\text{fund}$ and $\gamma.\text{allot}$ as the contract $C(\gamma)$, and output (unbound) .
 - c) If Φ 's pending update is confirmed and γ has not been initiated closing, go to step (5) of procedure (D) if the party is I and step (6) otherwise. Stop the procedure.

(F) Close

- 1) Upon $(\text{close}, id) \xleftarrow{t} \mathcal{Z}$, P sends $(\text{close}, \beta^P.\text{state}) \xrightarrow{t} C(\beta)$. If $\beta.\text{bind} = \emptyset$, P waits for Δ time and goes to step 3. Otherwise, for $\Phi \in \beta.\text{bind}$, P starts procedure (E). When all channels are unbound, P goes to step 4.
- 2) Upon $(\text{closing}) \xleftarrow{\tau} C(\beta)$, Q sends $(\text{close}, \beta^Q.\text{state}) \xrightarrow{\tau} C(\beta)$. If $\beta.\text{bind} = \emptyset$, Q goes to step 3. Otherwise, Q executes procedure (E) for each bind in $\beta.\text{bind}$. When all channels are unbound, Q goes to step 4.
- 3) If the party received (closed) from $C(\beta)$ within Δ time, then output (closed) and stop. Otherwise, send (close-resolve) to $C(\beta)$ after Δ time, output (closed) upon receiving it from $C(\beta)$, and stop.
- 4) Send (close-resolve) to $C(\beta)$, output (closed) upon receiving it from $C(\beta)$, and stop.

Fig. 8: Procedures (D) Bind Update, (E) Unbind, and (F) Close.

process of coin shift is the same as described in Section III, including two phases of bind state update (Steps (1-4), Figure 8(D)) and channel state update (Steps (5-6), Figure 8(D)). Here we distinguish the second phase from the procedure (B). The former updates the bind version and intermediate's balance recorded in the channel state, while the latter is to perform payments in the channel. Besides different contents to be updated, the former is initiated by the intermediate of bound channels, and the latter could be initiated by any channel user. When the first phase completes, this coin shift takes effect. If it halts, users stop performing bind updates and initiates the unbinding. If the second phase halts, users stop conducting channel updates and initiate the channel closing.

The process of unbind is presented in Figure 8(E). Behaviors of the trigger and triggered contract are shown in Steps (3) and (4) of Figure 9 respectively. Any user could initiate the unbind (Step (1), Figure 8(E)) and the other users respond with the latest bind state (Step (2), Figure 8(E)). They would trigger the contract that timeout if it is not handled in a timely manner (Step (3a), Figure 8(E)). From the aspect of the trigger contract, when triggered by one user to unbind Φ , the contract informs the other two users about the event and waits Δ for their responses (Step (3), Figure 9(B)). Finally, the contract decides the latest state, i.e., the valid state with the highest bind version, and triggers the other contract with the latest state. In addition, the contract sends or receives coins according

(A) The contract for channel opening

Upon $(\text{open}, \beta) \xrightarrow{\tau} \beta.\text{Alice}$, $(\text{remove}, \beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})) \xrightarrow{\tau} \mathcal{L}$, $(\text{opening}, \beta) \xrightarrow{\tau} \beta.\text{Bob}$. Wait for one of the following messages,

- 1) $(\text{open}) \xrightarrow{\tau_1 \leq \tau + \Delta} \beta.\text{Bob}$: $(\text{remove}, \beta.\text{Bob}, \beta.\text{balance}(\beta.\text{Bob})) \xrightarrow{\tau_1} \mathcal{L}$, $(\text{opened}) \xrightarrow{\tau_1} \beta.\text{users}$, and go to point (B).
- 2) $(\text{timeout}) \xrightarrow{\tau_2 > \tau + \Delta} \beta.\text{Alice}$: $(\text{add}, \beta.\text{Alice}, \beta.\text{balance}(\beta.\text{Alice})) \xrightarrow{\tau_2} \mathcal{L}$, and close the contract.

(B) The contract execution of channel β

Wait for the following messages:

- 1) $(\text{bind}, \Phi, t, \Sigma) \xrightarrow{\tau \leq t + 2 + \Delta} \Phi.\text{Irene}$ where β is $\Phi.\text{c}_a$ and Σ are signatures of $\Phi.\text{users}$ on (Φ, t) :
 - a) If $\Phi \notin \beta.\text{bind}$ and $\beta.\text{allot} \geq \Phi.\text{reserve}(\beta)$, trigger (bind, Φ) of $\mathcal{C}(\Phi.\text{c}_p)$. If $\mathcal{C}(\Phi.\text{c}_p)$ returns (bound) , add Φ to $\beta.\text{bind}$, $\beta.\text{allot} := \beta.\text{allot} - \Phi.\text{reserve}(\beta)$, $(\text{bound}) \xrightarrow{\tau} \Phi.\text{users}$. */*binding validity verification*/*
 - b) Otherwise, ignore the message.
- 2) $(\text{bind}, \Phi) \xrightarrow{\tau} \mathcal{C}(\Phi.\text{c}_a)$: If $\Phi \notin \beta.\text{bind}$ and $\beta.\text{allot} \geq \Phi.\text{reserve}(\beta)$, add Φ to $\beta.\text{bind}$, $\beta.\text{allot} := \beta.\text{allot} - \Phi.\text{reserve}(\beta)$, return (bound) . Otherwise, ignore the message. */*binding validity verification*/*
- 3) $(\text{unbind}, \tilde{id}, W) \xrightarrow{\tau} P$ where \tilde{id} is the identifier of $\Phi \in \beta.\text{bind}$ and β is $\Phi.\text{c}_a$: set $\Phi.\text{toResp} := \Phi.\text{other-users}(P)$, $(\text{unbinding}, \tilde{id}) \xrightarrow{\tau} \Phi.\text{toResp}$. If W is a valid state, set $\Phi.\text{state} := W$. Wait for the following messages,
 - a) $(\text{unbind}, \tilde{id}, W') \xrightarrow{\tau_1 \leq \tau + \Delta} P'$ where $P' \in \Phi.\text{toResp}$: Remove P' from $\Phi.\text{toResp}$. If W' is a valid state and $W'.\text{version} > \Phi.\text{version}$, set $\Phi.\text{state} := W'$. If $\Phi.\text{toResp} = \emptyset$, proceed procedure (C), remove Φ from $\beta.\text{bind}$, store $\Phi.\text{version}$ and $\Phi.\text{laShift}$, invoke $(\text{unbind}, \tilde{id}, \Phi.\text{state})$ of $\mathcal{C}(\Phi.\text{c}_p)$, and $(\text{unbound}) \xrightarrow{\tau} \Phi.\text{users}$.
 - b) $(\text{timeout}, \tilde{id})$ after Δ time : proceed procedure (C), remove Φ from $\beta.\text{bind}$, store $\Phi.\text{version}$ and $\Phi.\text{laShift}$, invoke $(\text{unbind}, \tilde{id}, \Phi.\text{state})$ of $\Phi.\text{c}_p$, and send (unbound) to $\Phi.\text{users}$.
- 4) $(\text{unbind}, \tilde{id}, \Phi.\text{state}) \xrightarrow{\tau} \mathcal{C}(\Phi.\text{c}_a)$: proceed procedure (C), remove Φ from $\beta.\text{bind}$, store $\Phi.\text{version}$ and $\Phi.\text{laShift}$.
- 5) $(\text{close}, M) \xrightarrow{\tau} P$: if M is a valid state, set $\beta.\text{state} := M$. Denote $P' := \beta.\text{other-user}(P)$, $(\text{closing}) \xrightarrow{\tau} P'$, and wait for the following messages,
 - a) $(\text{close}, M') \xrightarrow{\tau_1 \leq \tau + \Delta} P'$: if M' is a valid state and $M'.\text{version} > \beta.\text{version}$, set $\beta.\text{state} := M'$. If $\beta.\text{bind} = \emptyset$, go to procedure (D).
 - b) (close-resolve) after Δ time : If $\beta.\text{bind} = \emptyset$, go to procedure (D).

(C) Subroutine for settling Φ 's coin shift

Let γ be the channel bound to β . Distinguish the following cases according to $\Phi.\text{shift}(\beta)$, denoted as amt ,

- 1) If $\text{amt} > 0$ and $\text{amt} \leq \Phi.\text{reserve}(\gamma)$, or $\text{amt} < 0$ and $|\text{amt}| \leq \Phi.\text{reserve}(\beta)$, set $\beta.\text{fund} := \beta.\text{fund} + \text{amt}$, $\beta.\text{allot} := \beta.\text{allot} + \text{amt} + \Phi.\text{reserve}(\beta)$.
- 2) Otherwise, set $\beta.\text{allot} := \beta.\text{allot} + \Phi.\text{reserve}(\beta)$.

(D) Subroutine for channel closing

For each unbound Φ , if $\Phi.\text{version} = \beta.\text{biList}(\Phi.\text{id}) + 1$, or $\Phi.\text{version} = 1$ and $\Phi.\text{id} \notin \beta.\text{biList}$, set $\beta.\text{balance}(\Phi.\text{Irene}) := \beta.\text{balance}(\Phi.\text{Irene}) + \Phi.\text{laShift}(\beta)$. Refund $\beta.\text{balance}(\beta.\text{Alice})$ coins to $\beta.\text{Alice}$ and $\beta.\text{balance}(\beta.\text{Bob})$ coins to $\beta.\text{Bob}$, send (closed) to $\beta.\text{users}$ and close this contract.

Fig. 9: The contract functionality.

to $\Phi.\text{Shift}$ recorded in the state, and withdraws the reserved amounts in the bind (Figure 9(C)).

Procedure (F) in Figure 8 shows the procedure of channel closing, corresponding to the contract procedures of Step (5) in Figure 9. Users submit the closing request and the latest channel state to the contract. If there are unbound channels, users proceed procedure (E) for each bind. The contract decides the latest channel state according to states submitted by users. Before refunding coins to users, the contract checks whether all binds have been dismissed. Besides, it needs to be verified that for each bind Φ , whether the version recorded in the channel state is consistent with that recorded in the contract (Figure 9(D)). If not, the contract recalculates the balance of $\Phi.\text{Irene}$ according to $\Phi.\text{laShift}$. Finally, the contract refunds coins to users and closes the channel.

Theorem 1: The protocol Shaduf running in the \mathcal{C} -hybrid world UC-realizes the ideal functionality \mathcal{F} with respect to the

global ledger $\mathcal{L}(\Delta)$.

Proof: The proof is in Appendix B. ■

VI. IMPLEMENTATION AND EVALUATION

To illustrate the feasibility of Shaduf, we provide a proof of concept implementation in Ethereum. We also conduct an experimental evaluation of Shaduf's performance in the payment channel network, to compare it with Revive, in a simulation of the Lighting Network. The source code is open on Github².

A. Implementation

Upon Ethereum, we implement the contract functionality shown in Figure 9 using the Solidity programming language.

²<https://github.com/Lonely-Programmer/Shaduf>

TABLE II: On-chain gas costs for executing Shaduf protocol, where “opti” represents the optimistic case and “pess” represents the pessimistic case.

	open	bind	unbind		close	
			opti	pess	opti	pess
Initiator	72342	292916	198069	277453	178108	206113
Responder	57934	0	112958	243299	62991	135396

We focus on the on-chain costs running the Shaduf protocol, i.e., fees paid to miners when interacting with the contract. In Ethereum, it is calculated by the *gas*, which depends on the amount of data and the computation complexity of the contract call. We set the gas price as 20 Gwei (as of June 2021). Due to the recent Ether price fluctuations, we use the Nov 2020 exchange rate of 500 USD per Ether. Shaduf’s costs are classified under each procedure and distinguished by the role one user plays, the initiator and the responder. The contract deployment costs 3.1M gas (31 USD). Since it can be mitigated by technical ways [7], we only consider the Shaduf running costs. Except for the channel state update and bind state update which are executed off-chain, costs of the other 4 procedures are displayed in Table II.

Costs of Shaduf. From Table II, channel opening costs the initiator 72k gas (0.72 USD) and the responder 57k gas (0.57 USD). Initiator of a bind, i.e., intermediate of two bound channels who submits the bind request to the blockchain, bears the total bind costs of 292k gas (2.92 USD). The unbinding could be initiated by any of its three users, and we measure the cost of unbinding in two cases. In the optimistic case, after the initiator submits the latest state to the contract, the other two users respond accordingly. In the pessimistic case, either the other two users do not respond in time thus the initiator needs to trigger the contract to finish the unbinding, or an expired state is submitted by the initiator thus the responder needs to submit the latest state. Similarly, we measure the costs of channel closing in both cases. The results show that the costs of Shaduf are minimized when all users behave honestly. The costliest procedure for both the initiator and the responder will be less than 520k gas (5.20 USD) even in the pessimistic case.

Cost Comparison. Shaduf spends at most 813k gas (8.13 USD) for the bind and unbind procedures. Compared with Close-and-Reopen which costs 307k gas (3.07 USD) each time of channel refunding and LOOP which costs half the gas, costs of 3 times Close-and-Reopen and 6 times LOOP exceed Shaduf’s costs, while Shaduf enables unlimited times of rebalancing after binding. Moreover, 3 times Close-and-Reopen and 6 times LOOP need at least 6Δ time, while Shaduf requires only Δ for the one-time binding.

B. Evaluation

Generally, off-chain rebalancing methods have the potential to enhance the off-chain network’s performance by enabling payments that may fail in the payment channel networks, such as the LN. Concretely, when performing a multi-hop payment in the LN, the sender selects the shortest path to the receiver from the network topology, and completes this payment with the help of the path users. However, the payment may fail due to any user, along the path, lacking coins in its hop. Revive and Shaduf alleviate this problem by refunding

the depletion hop via coins in other channels, in global and local manners, respectively. To clarify their actual effects, we conduct the experimental evaluation and aim to answer the following questions,

- **Q1:** How does Shaduf perform in the off-chain network, in the respect of improving the network’s success ratio and volume, and reducing users’ deposits in channels, especially compared with LN and Revive?
- **Q2:** How does the binding strategy affect Shaduf’s performance?

1) *Evaluation setup:* We setup the evaluation from two aspects, including the network topology and off-chain payments.


Network topology. Although Shaduf is constructed upon Turing-complete blockchain platforms, we use the real and the largest off-chain network, LN topology, to evaluate the performance. We get the snapshot of LN topology on 2021-03-31 (from <https://ln.bigsun.xyz/>), which contains 10,529 nodes and 38,910 channels. Since the initial balance allocation within each channel is invisible, we evenly assign the fund in each channel to two channel users to bootstrap the network.

Off-chain payments simulation. Due to the fact that the off-chain payments (including the occurrence and amount) are invisible, to simulate the off-chain payments, we need to simulate both the sender-receiver pairs and the payment amounts between each pair.

- For choosing sender-receiver pairs, we consider two typical cases, uniform and skewed ones, separately. In the uniform case, a user performs payments with each other at the same probability. Therefore, senders and receivers are uniformly sampled from the network. In the skewed case, users such as customers pay more while merchants receive more. To simulate it, inspired by [23], we sample the sender and receiver from independent exponential distributions. The skewness of distribution is represented by the skewness factor, and a higher value indicates more skewed payments.
- As for off-chain transaction values, we simulate them by randomly sampling from the Bitcoin trace from 2021-03-01 to 2021-03-31. Furthermore, since LN is mainly used for high-frequency and low-amount transactions (or so-called micro-payments), to make the evaluation more compatible with the real off-chain network, we abandon the values higher than the median of the channel capacity. Finally, 2.65M transaction values are available.

2) *Evaluation methodology:* We design methods of evaluation for Revive and Shaduf. For the former, we implement the original Revive as suggested by [13], and optimize it to promote the performance. For the latter, we propose three different binding strategies.

Revive and its optimization. In Revive, the rebalancing demand is resolved in a global manner, i.e., the user completes the refund with other users who also have the refunding demand, under the help of a global fair leader. We simulate the rebalancing process in the following way. Firstly, each round of Revive starts by collecting the failed payments and regarding each user whose coins are insufficient for the payment to have the refunding demand, i.e., shifting coins in other channels to the current channel to continue the payment. Secondly,



(a) Comparing succ. ratio of 4 schemes under uniform payment demands, with channel capacity factor from 1 to 25. (b) Comparing succ. ratio of 4 schemes under skewed payment demands, with capacity factor 10 and payment skewness factor from 1 to 8. (c) Comparing succ. ratio of 4 schemes under skewed payment demands, with skewness factor 8 and channel capacity factor from 1 to 25. (d) Comparing succ. ratio of LN and 3 binding strategies of Shaduf under uniform payment demands, with channel capacity factor from 1 to 25.

Fig. 10: Effect of Shaduf varying the channel capacity, payment skewness and the binding strategy.

since more demands mean a higher possibility of directed cycles, to see the best performance of Revive, we assume that the fair leader starts this round of rebalancing when the number of demands reaches the maximum, i.e., 300 users or 1000 channels as suggested by [13]. After the generation of rebalancing transactions by the fair leader, we assume that all cycle users would cooperate with the rebalancing, after which all the failed payments are retried.

The performance of Revive can be improved by relaxing the refunding restrictions. A natural optimization of Revive could be achieved by exploiting the shortest cycles in the payment channel network. Instead of waiting for the other refunding demands to form directed cycles, a channel with rebalancing demand in the optimized Revive (OPT-Revive for short) gets help from users in its shortest cycles. In this way, the user is allowed to move coins in each other channel to the required channel, until sufficient coins are collected.

Shaduf. The refunding demand is resolved in a local manner, i.e., a user can shift coins within bound channels, with the cooperation of the other two users in each pair of bound channels. To simulate the case, we need to initiate the bindings for each user. For binding amounts, funds in a channel are evenly assigned to both end users to bind other channels, i.e., the maximum of funds that can be shifted out. When the channel is bound to multiple channels, the amount is equally distributed among them. As for binding strategy, a user with x channels can choose any subset from the C_x^2 pairs of bindings as its strategy. Three reasonable choices, including “high-to-low”, “random-bind” and “all-bind”, are described below.

In the first strategy, each user binds the channel with the highest balance to the channel with the lowest balance, the channel with the second-highest balance to the channel with the second-lowest balance, and so on. In the second strategy, a user with x channels chooses bindings randomly, with the limitation that exactly $\lfloor 0.5 \cdot x \rfloor$ bindings are permitted. That is to say, the first and second strategies share the same number of bindings. In the third strategy, each pair of channels are bound. We denote the three strategies as “HL-Shaduf”, “RB-Shaduf” and “AB-Shaduf” respectively.

Ways of evaluation and comparison. Following the methodology and with the aforementioned evaluation setup, we carry out the evaluation. More concretely,

- To answer the first question (Q1), we test Shaduf’s performance with the “high-to-low” strategy, under the uniform (labeled with ‘-uni’) and skewed (labeled with

‘-skew’) payment demands separately. Performances of LN and (OPT-)Revive under the same setting are also displayed for comparison. To derive Shaduf’s effect on deposit reduction, we scale the channel capacity by a factor and see at which scale would (OPT-)Revive and Shaduf achieve the same performance as LN.

- To answer the second question (Q2), we vary Shaduf’s binding strategy and test the performance under the uniform payments.

It should be noted that in each evaluation, 50,000 sender-receiver pairs with corresponding payment values are chosen. Since randomness is introduced in the simulation of off-chain payments, we execute each evaluation 10 times and calculate the average. Both the success ratio and volume are taken as the metrics. Since the success volume is proportional to the success ratio, we put the volume data in Appendix C.

3) Evaluation results:

Success ratio in uniform payments. We vary the channel capacity from 1x to 25x and show the success ratio in uniform payments in Figure 10a. While Revive’s performance is slightly better than LN, OPT-Revive achieves around 7% enhancement, and Shaduf achieves around 10% - 22% enhancement.

Success ratio in skewed payments. We test Shaduf’s performance in skewed payments, under the fixed channel capacity factor 10. We increase the skewness factor until LN’s success ratio drops to 50%, which we regard as the minimum-usable case. From the results displayed in Figure 10b, OPT-Revive achieves 7.5% enhancement and Shaduf achieves around doubled enhancement on average.

Deposit comparison in uniform payments. From Figure 10a, we see at which scale would OPT-Revive and Shaduf achieves the same success ratio of LN. For example, when the success ratio is 70%, the capacity factors of LN, OPT-Revive, and Shaduf are 15, 10, and 6, respectively. It means that OPT-Revive and Shaduf achieve the same performance with 0.67x and 0.4x capacity of LN. Correspondingly, they save 33% and 60% deposit respectively. On average, OPT-Revive and Shaduf saves 31% and 60% of the capacity, respectively.

Deposit comparison in skewed payments. Facing skewed payments, we consider the performances of OPT-Revive and Shaduf under the skewness factor 8. We vary the channel capacity and the results are displayed in Figure 10c. From the results, Shaduf improves the success ratio of LN and Revive by

12% on average. Correspondingly, 0.66x and 0.48x capacity is required for OPT-Revive and Shaduf to perform as effectively as LN, with 34% and 52% capacity saved. For example, when the success ratio is around 62%, the capacity factors of LN, OPT-Revive, and Shaduf are 20, 13, and 10, respectively.

Binding effect. Figure 10d shows the effect of the binding strategy. Shaduf under the “high-to-low” strategy performs better than the “random-bind” strategy, even under the same number of bindings, while the “all-bind” strategy performs best. In terms of the success ratio, the three strategies’ enhancements vary from 6% to 28%. Under the same success ratio of LN, they achieve 37% - 75% deposit savings.

In summary, Shaduf’s performance fluctuates under different payment demands and binding strategies. In general, Shaduf outperforms LN and Revive by 6% - 28% in success ratio, and reduces users’ deposits by 37% - 75% under the same success ratio. A more dedicated binding strategy would further improve Shaduf’s effect.

VII. DISCUSSION

A. Choosing the Binding Strategy

Apart from the three binding strategies described above, there are many other strategies for users to bind the channels. While more bindings cost higher on-chain fees, i.e., \$2.92 times the number of bindings, in general, more payments would be enabled, from which the user would save the costs for on-chain channel refunding and charge fees for facilitating these payments. From another aspect, the deposit in channels could be reduced while maintaining the same performance. Therefore, users could trade off the costs and benefits of bindings, then decide and adjust the bindings according to the practical and dynamic demands.

B. Privacy Issue

In the Lightning Network, the channel privacy means that the concrete fund allocation within one channel is invisible to anyone outside the channel. In other words, from their point of view, each channel user’s balance is random in the range of 0 to the channel capacity. Recall that in Section III-B, in order to guarantee the balance security, the on-chain binding is required before performing off-chain coin shifts. Furthermore, to enable one channel to be bound to multiple channels, the upper limit of coin shifts is declared on the blockchain. Therefore, the relation of binding and the upper limit are public. Nevertheless, the privacy is not compromised in Shaduf. By dividing users outside the channel into two categories, i.e., blockchain observers and the user in its bound channel, we provide the non-formal privacy analysis from their perspectives, as shown below.

From the aspect of blockchain observers, no further information about the channel state is obtained since the coin shifts between bound channels are also off-chain performed among bound users. Moreover, the potential flowing-in coins from bound channels actually increase the channel capacity. Therefore, the range of each user’s balance is expanded.

Compared with blockchain observers, the user in the bound channel knows the extra information, i.e., the concrete number of shifted coins between them. When coins are shifted out, the

channel’s capacity seems to be reduced, shrinking the channel users’ balance range. However, the channel privacy is retained for two reasons. Firstly, the concrete balance allocation is still unknown. Secondly, since one channel could be bound to multiple channels, the user’s balance range is affected by coins shifted from all bound channels.

In short, even though the binding relation and the upper limit of coin shift are public, Shaduf does not compromise the privacy. We leave the formal proof as the future work.

VIII. RELATED WORK

The issue of channel fund depletion is first proposed by Revive [13]. With the atomic multi-channel update protocol [9], Revive is Bitcoin-compatible and can be deployed to, for example, LN. Spider [23] solves the problem in a different way, which proposes a balanced routing scheme to avoid fund depletion in channels facing skewed payment flows. Similarly, to avoid payment failure due to depleted funds in channels, Li *et al.* [14] plan the deposition in channels to satisfy the payment demands. Both approaches aim to avoid fund depletion, which are orthogonal to Revive and our work since we provide solutions for channels that have already been depleted. The two kinds of solutions can be combined to achieve a better performance in maintaining balanced channels.

1) *Payment channel:* Many variants of the payment channel have been extensively studied. Bolt [10] constructs anonymous channels where payments in one channel are unlinkable. Payment channels are built upon multiple blockchain platforms [20], [26], and extended to state channels where arbitrary state transitions are enabled [1], [6], [8], [19].

2) *Payment channel network:* The routing strategies have been studied in [15], [22], [23], [25]. Tumblebit [12] constructs payment unlinkable PCH. Maintaining the privacy feature, A²L [24] enhances Tumblebit’s efficiency and interoperability. Malavolta *et al.* [17] report the wormhole attack and propose the secure and privacy-preserving PCN construction. Sprites [19] reduces the collateral costs of multi-hop payments. Egger *et al.* [9] achieve the collateral reduction on Bitcoin-compatible platforms. Aumayr *et al.* [3] build the secure and collateral-reduced multi-hop payments in a single round. Perun [7] proposes the concept of the virtual channel, via which users with one-hop distance can perform off-chain payments directly. The virtual channel is also enabled on Bitcoin-compatible platforms [2]. Dziembowski *et al.* [8] connect arbitrary users directly via virtual channels and build the state channel network. The virtual channel is also extended to multi-party where multiple parties can perform state transitions [6]. We refer readers to [11] for a comprehensive understanding of the payment channel and the network.

IX. CONCLUSION

In this paper, we have proposed a non-cycle off-chain rebalancing scheme named Shaduf. It allows users to rebalance depleted channels via exploiting coins from other channels in more general cases. We formalize the protocol in the UC framework and prove its security. Experimental evaluation shows that Shaduf outperforms LN and Revive in raising the performance of the off-chain payment network. We also remark

that Shaduf is an orthogonal and compatible work to the well-studied off-chain routing methods, therefore we believe the joint applications of Shaduf with them will perform even better.

ACKNOWLEDGMENT

The authors thank Yanxue Jia, Shi-Feng Sun, Zhen Liu, Zhi-Qiang Liu, and the anonymous NDSS reviewers for their constructive suggestions. This work is supported by the Key (Keygrant) Project of Chinese Ministry of Education (No. 2020KJ010201), the National Natural Science Foundation of China (NSFC) (No. 61872142), and the Key Research and Development Plan of Shandong Province (No. 2021CXGC010105).

REFERENCES


- [1] L. Aumayr, O. Ersoy, A. Erwig, S. Faust, K. Hostáková, M. Maffei, P. Moreno-Sanchez, and S. Riahi, “Generalized channels from limited blockchain scripts and adaptor signatures,” in *Advances in Cryptology - ASIACRYPT*, vol. 13091. Springer, 2021, pp. 635–664.
- [2] L. Aumayr, M. Maffei, O. Ersoy, A. Erwig, S. Faust, S. Riahi, K. Hostáková, and P. Moreno-Sanchez, “Bitcoin-compatible virtual channels,” in *42nd IEEE Symposium on Security and Privacy, SP*. IEEE, 2021, pp. 901–918.
- [3] L. Aumayr, P. Moreno-Sanchez, A. Kate, and M. Maffei, “Blitz: Secure multi-hop payments without two-phase commits,” in *30th USENIX Security Symposium*. USENIX Association, 2021, pp. 4043–4060.
- [4] R. Canetti, “Universally composable security: A new paradigm for cryptographic protocols,” in *42nd Annual Symposium on Foundations of Computer Science, FOCS*. IEEE Computer Society, 2001, pp. 136–145.
- [5] R. Canetti, Y. Dodis, R. Pass, and S. Walfish, “Universally composable security with global setup,” in *Theory of Cryptography, 4th Theory of Cryptography Conference, TCC*, vol. 4392. Springer, 2007, pp. 61–85.
- [6] S. Dziembowski, L. Eckey, S. Faust, J. Hesse, and K. Hostáková, “Multi-party virtual state channels,” in *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, vol. 11476. Springer, 2019, pp. 625–656.
- [7] S. Dziembowski, L. Eckey, S. Faust, and D. Malinowski, “Perun: Virtual payment hubs over cryptocurrencies,” in *IEEE Symposium on Security and Privacy, SP*. IEEE, 2019, pp. 106–123.
- [8] S. Dziembowski, S. Faust, and K. Hostáková, “General state channel networks,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, 2018, pp. 949–966.
- [9] C. Egger, P. Moreno-Sanchez, and M. Maffei, “Atomic multi-channel updates with constant collateral in bitcoin-compatible payment-channel networks,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, 2019, pp. 801–815.
- [10] M. Green and I. Miers, “Bolt: Anonymous payment channels for decentralized currencies,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, 2017, pp. 473–489.
- [11] L. Gudgeon, P. Moreno-Sanchez, S. Roos, P. McCorry, and A. Gervais, “Sok: Layer-two blockchain protocols,” in *Financial Cryptography and Data Security - 24th International Conference, FC*, ser. Lecture Notes in Computer Science, vol. 12059. Springer, 2020, pp. 201–226.
- [12] E. Heilman, L. Alshenibr, F. Baldimtsi, A. Scafuro, and S. Goldberg, “Tumblebit: An untrusted bitcoin-compatible anonymous payment hub,” in *24th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2017.
- [13] R. Khalil and A. Gervais, “Revive: Rebalancing off-blockchain payment networks,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, 2017, pp. 439–453.
- [14] P. Li, T. Miyazaki, and W. Zhou, “Secure balance planning of off-blockchain payment channel networks,” in *39th IEEE Conference on Computer Communications, INFOCOM*. IEEE, 2020, pp. 1728–1737.
- [15] G. Malavolta, P. Moreno-Sanchez, A. Kate, and M. Maffei, “Silentwhispers: Enforcing security and privacy in decentralized credit networks,” in *24th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2017.
- [16] G. Malavolta, P. Moreno-Sanchez, A. Kate, M. Maffei, and S. Ravi, “Concurrency and privacy with payment-channel networks,” in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, 2017, pp. 455–471.
- [17] G. Malavolta, P. Moreno-Sanchez, C. Schneidewind, A. Kate, and M. Maffei, “Anonymous multi-hop locks for blockchain scalability and interoperability,” in *26th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2019.
- [18] V. Mavroudis, K. Wüst, A. Dhar, K. Kostianen, and S. Capkun, “Snappy: Fast on-chain payments with practical collaterals,” in *27th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2020.
- [19] A. Miller, I. Bentov, S. Bakshi, R. Kumaresan, and P. McCorry, “Sprites and state channels: Payment networks that go faster than lightning,” in *Financial Cryptography and Data Security - 23rd International Conference, FC*, vol. 11598. Springer, 2019, pp. 508–526.
- [20] P. Moreno-Sanchez, A. Blue, D. V. Le, S. Noether, B. Goodell, and A. Kate, “DLSAG: non-interactive refund transactions for interoperable payment channels in monero,” in *Financial Cryptography and Data Security - 24th International Conference, FC*, vol. 12059. Springer, 2020, pp. 325–345.
- [21] J. Poon and T. Dryja, “The bitcoin lightning network: Scalable off-chain instant payments,” 2016.
- [22] S. Roos, P. Moreno-Sanchez, A. Kate, and I. Goldberg, “Settling payments fast and private: Efficient decentralized routing for path-based transactions,” in *25th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2018.
- [23] V. Sivaraman, S. B. Venkatakrishnan, K. Ruan, P. Negi, L. Yang, R. Mittal, G. C. Fanti, and M. Alizadeh, “High throughput cryptocurrency routing in payment channel networks,” in *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI*. USENIX Association, 2020, pp. 777–796.
- [24] E. Tairi, P. Moreno-Sanchez, and M. Maffei, “A²I: Anonymous atomic locks for scalability in payment channel hubs,” in *42nd IEEE Symposium on Security and Privacy, SP*. IEEE, 2021, pp. 1834–1851.
- [25] P. Wang, H. Xu, X. Jin, and T. Wang, “Flash: efficient dynamic routing for offchain networks,” in *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies, CoNEXT*. ACM, 2019, pp. 370–381.
- [26] Y. Zhang, Y. Long, Z. Liu, Z. Liu, and D. Gu, “Z-channel: Scalable and efficient scheme in zerocash,” in *Information Security and Privacy - 23rd Australasian Conference, ACISP*, vol. 10946. Springer, 2018, pp. 687–705.

APPENDIX A

RESTRICTIONS ON THE ENVIRONMENT

Restrictions on the environment are listed in the following.

- The environment would never ask a user to open a channel that already exists, or when the user has insufficient coins for the deposit.
- The environment would never ask a user to perform payments in the channel when the user has insufficient balance to afford this payment.
- The environment would never request two already bound channels to bind or two not-bound channels to unbind.
- The environment would never ask a user to shift coins between channels that are not bound, or the coins requested to shift exceed the upper limit. More concretely, for the bind Φ between channels β and γ , the maximum



(a) Comparing succ. volume of 4 schemes under uniform payment demands, with channel capacity factor from 1 to 25. (b) Comparing succ. volume of 4 schemes under skewed payment demands, with capacity factor 10 and payment skewness factor from 1 to 8. (c) Comparing succ. volume of 4 schemes under skewed payment demands, with skewness factor 8 and channel capacity factor from 1 to 25. (d) Comparing succ. volume of LN and 3 binding strategies of Shaduf under uniform payment demands, with channel capacity factor from 1 to 25.

Fig. 14: Success volume of the evaluation.

number of coins that could be shifted from β to γ is $\min(\Phi.\text{reserve}(\beta) + \Phi.\text{shift}(\beta), \beta.\text{balance}(\Phi.\text{Irene}))$.

- The environment would never request a user to update the channel state or bind state when the previous update is not completed.
- The environment would never request a user to close a channel that does not exist.

APPENDIX B PROOF OF THEOREM 1

To achieve the indistinguishability between the ideal world and the real world, the simulator \mathcal{S} sends inputs to the ideal functionality \mathcal{F} observing behaviors of the corrupted parties and sends messages to the corrupted parties on behalf of the contract functionality \mathcal{C} and the honest parties.

At the beginning, \mathcal{S} corrupts the parties that \mathcal{A} corrupts. \mathcal{S} also generates the public-private key pairs for all parties and passes all public keys to the corrupted parties. For each corrupted party, its private key is sent along. Simulation on each procedure is shown in Figure 12 and 13. In each procedure, we ignore cases where involved parties are all honest or corrupted, since \mathcal{S} simulates behaviors of the corrupted parties and we aim to protect the balance security of the honest parties.

APPENDIX C SUCCESS VOLUME OF EVALUATION

The total payment volumes in the evaluation is 94 BTC. The data of success volume displayed in Figure 14a, 14b, 14c, and 14d correspond to the four evaluation that is shown in Figure 10a, 10b, 10c, and 10d, respectively.

Assume the following messages concerning the channel β and the bind Φ , with their identifiers, $\beta.\text{id}$ and $\Phi.\text{id}$, denoted as id and \tilde{id} , respectively. In procedures (B) and (D), we denote the requested payment and coin shift as θ and $\tilde{\theta}$, respectively. We abbreviate $A := \beta.\text{Alice}$, $B := \beta.\text{Bob}$, $I := \Phi.\text{Irene}$, and denote the initiator of channel update as P and the responder as Q .

(A) Open

Case A is honest and B is corrupt

Upon A sending $(\text{open}, \beta) \xrightarrow{t} \mathcal{F}$, send $(\text{open}, \beta) \xrightarrow{t} \mathcal{C}(\beta)$ on behalf of A. Assume it reaches \mathcal{L} in round $\tau \leq t + \Delta$. If B sends $(\text{open}) \xrightarrow{\tau} \mathcal{C}(\beta)$, send $(\text{open}) \xrightarrow{\tau} \mathcal{F}$ on behalf of B. Otherwise, upon A sending (refund, β) to \mathcal{F} , send (timeout) to $\mathcal{C}(\beta)$ on behalf of A.

Case A is corrupt and B is honest

Upon A sending $(\text{open}, \beta) \xrightarrow{t} \mathcal{C}(\beta)$, send $(\text{open}, \beta) \xrightarrow{t} \mathcal{F}$ on behalf of A. Assume it reaches \mathcal{L} in round $\tau \leq t + \Delta$. If B sends $(\text{open}) \xrightarrow{\tau} \mathcal{F}$, send $(\text{open}) \xrightarrow{\tau} \mathcal{C}(\beta)$ on behalf of B. Otherwise, upon B sending (timeout) to $\mathcal{C}(\beta)$, send (refund, β) to \mathcal{F} on behalf of B.

(B) Channel Update

Case P is honest and Q is corrupt

Upon P sending $(\text{chan-update}, id, \theta) \xrightarrow{t} \mathcal{F}$, generate signature σ_P on $msg = (id, \beta^P.\text{version} + 1, \beta^P.\text{balance} + \theta, \beta^P.\text{biList})$, and send $(\text{chan-update}, msg, \sigma_P) \xrightarrow{t} Q$ on behalf of P. If Q sends his signature on msg to P in round $t + 1$, send (chan-update-ok) to \mathcal{F} on behalf of Q. Otherwise, initiate channel closing on behalf of P.

Case P is corrupt and Q is honest

Upon P sending $(\text{chan-update}, msg = (ver, \alpha, biList), \sigma_P) \xrightarrow{t} Q$ where σ_P is P's signature on msg , if $ver = \beta^Q.\text{version} + 1$ and $biList = \beta^Q.\text{biList}$, send $(\text{chan-update}, id, \alpha - \beta^Q.\text{balance}) \xrightarrow{t} \mathcal{F}$ on behalf of P. Else, stop. If Q sends $(\text{chan-update-ok}) \xrightarrow{t+1} \mathcal{F}$, generate Q's signature σ_Q on msg and send it to P on behalf of Q.

(C) Bind

Case I is honest

Upon I sending $(\text{bind}, \Phi) \xrightarrow{t} \mathcal{F}$, generate I's signature σ_I on (Φ, t) , send $(\text{bind}, \Phi, t, \sigma_I)$ to $\Phi.\text{end-users}$ on behalf of I. Proceed as follows,

- 1) For $P \in \Phi.\text{end-users}$, if P is corrupt and sends $(\text{bind}, \sigma_P) \xrightarrow{t+1} I$ where σ_P is P's signature on (Φ, t) , send $(\text{bind}, \Phi) \xrightarrow{t+1} \mathcal{F}$ on behalf of P. If P is honest and sends $(\text{bind-ok}) \xrightarrow{t} \mathcal{F}$, generate P's signature σ_P on (Φ, t) , and send it to I on behalf of P.
- 2) In round $t + 2$, if I sends (bind-confirmed) to \mathcal{F} , send $(\text{bind}, \Phi, t, \sigma_P, \sigma_I, \sigma_Q) \xrightarrow{t+2} \mathcal{C}(\Phi.\text{ca})$ on behalf of I.

Case I is corrupt

Upon I sending $(\text{bind}, \Phi, t, \sigma_I) \xrightarrow{t} P \in \Phi.\text{end-users}$ and P is honest, send $(\text{bind}, \Phi) \xrightarrow{t} \mathcal{F}$ on behalf of I, and send $(\text{send-req}, P)$ to \mathcal{F} . Proceed as follows,

- 1) For $P \in \Phi.\text{end-users}$, if P is corrupt, send (bind-ok) to \mathcal{F} in round $t + 1$ on behalf of P. Else, if P is honest and sends (bind-ok) to \mathcal{F} , generate P's signature σ_P on (Φ, t) and send (bind, σ_P) to I on behalf of P.
- 2) In round $t + 2$, if I sends $(\text{bind}, \Phi, t, \Sigma)$ to $\mathcal{C}(\Phi.\text{ca})$ where Σ is the set of users' signatures on (Φ, t) , $(\text{bind-confirmed}) \xrightarrow{t+2} \mathcal{F}$ on behalf of I.

(D) Bind Update

Case I is honest

Upon I sending $(\text{bind-update}, \tilde{id}, \tilde{\theta}) \xrightarrow{t} \mathcal{F}$, generate I's signature σ_I on $msg = (\Phi^I.\text{version} + 1, \Phi^I.\text{shift} + \tilde{\theta}, \tilde{\theta})$, $(\text{bind-update}, msg, \sigma_I) \xrightarrow{t} \Phi.\text{end-users}$ on behalf of I, and proceed as follows,

- 1) For $P \in \Phi.\text{end-users}$, distinguish the following cases,
 - a) If P is honest and sends $(\text{bind-update-ok}) \xrightarrow{t} \mathcal{F}$, generate P's signature σ_P on msg send $(\text{bind-update}, \sigma_P) \xrightarrow{t+1} \Phi.\text{other-users}(P)$ on behalf of P, and forward it to the other end user on behalf of I. In round $t + 2$, if P receives the other end user's signature, forward it to I on behalf of P.
 - b) If P is corrupt and sends her signature on msg to I in round $t + 1$, forward it to the other end user on behalf of I. If P's signature is sent to I before round $t + 2$, send (bind-update-ok) to \mathcal{F} on behalf of P.
- 2) In round $t + 3$, if I does not collect others' signatures, initiate the unbinding in procedure (E) on behalf of I. Otherwise, for $\gamma \in \Phi.\text{channels}$:
 - a) Update I's balance by $\tilde{\theta}$ and denote the updated channel balance allocation as $\gamma^I.\text{balance}'$, add 1 to $\gamma^I.\text{biList}(\tilde{id})$ (set to 1 if \tilde{id} is not in the list) and denote it as $\gamma^I.\text{biList}'$. Generate I's signature σ_I on $msg = (\gamma^I.\text{version} + 1, \gamma^I.\text{balance}', \gamma^I.\text{biList}')$, and send $(\text{chan-update}, \sigma_I)$ to the other channel user, denoted as P, on behalf of I.
 - b) In round $t + 4$, if P is honest, generate P's signature on msg and send it to I on behalf of P. If P is corrupt and does not send P's signature on msg to I, send $(\text{bind-chan-not-updated}, \gamma.\text{id})$ to \mathcal{F} on behalf of P, and initiate channel closing in procedure (E) on behalf of I.

Fig. 12: Simulation: (A) Open, (B) Channel Update, (C) Bind, and (D) Bind Update.

Assume the following messages concerning the channel β and the bind Φ , with their identifiers, $\beta.\text{id}$, and $\Phi.\text{id}$, denoted as id and $\tilde{\text{id}}$, respectively. In the procedures (E) Unbind and (F) Close, we denote the initiator as P and the responder as Q .

(D) Bind Update

Case I is corrupt

Upon I sending $(\text{bind-update}, \text{msg} = (\text{ver}, \text{shift}, \tilde{\theta}), \sigma_I) \xrightarrow{t} P \in \Phi.\text{end-users}$, where $\text{ver} = \Phi.\text{version} + 1$, $\text{shift} = \Phi.\text{shift} + \tilde{\theta}$, and σ_I is I 's signature on msg , $(\text{bind-update}, \tilde{\text{id}}, \tilde{\theta}) \xrightarrow{t} \mathcal{F}$ on behalf of I , and $(\text{send-req}, P) \xrightarrow{t} \mathcal{F}$. Proceed as follows,

1) For $P \in \Phi.\text{end-users}$, distinguish the following cases,

- a) If P is honest and sends $(\text{bind-update-ok}) \xrightarrow{t+1} \mathcal{F}$, generate P 's signature σ_P on msg and send $(\text{bind-update}, \sigma_P) \xrightarrow{t+1} \Phi.\text{other-users}(P)$ on behalf of P . If P receives the other end user's signature on msg in round $t+2$, forward the message to I on behalf of P .
- b) If P is corrupt and P 's signature is sent to the honest end user before round $t+2$, send (bind-update-ok) to \mathcal{F} on behalf of P .

2) If the honest users do not collect others' signatures, initiate unbinding on behalf of the honest users in procedure (E). Otherwise, for each honest user P , in his channel $\gamma \in \Phi.\text{channels}$, proceed as follows,

- a) Update I 's balance by $\tilde{\theta}$ and denote the updated channel balance allocation as $\gamma^I.\text{balance}'$, add 1 to $\gamma^I.\text{biList}(\tilde{\text{id}})$ (set to 1 if $\tilde{\text{id}}$ is not in the list) and denote it as $\gamma^I.\text{biList}'$.
- b) If I sends $(\text{chan-update}, \sigma_I) \xrightarrow{t+3} P$ where σ_I is I 's signature on msg , generate P 's signature σ_P on msg , and send $(\text{chan-update}, \sigma_P) \xrightarrow{t+4} I$ on behalf of P . Otherwise, $(\text{bind-chan-not-updated}, \gamma.\text{id}) \xrightarrow{t+3} \mathcal{F}$ on behalf of I , and initiates channel closing in procedure (F) on behalf of P .

(E) Unbind

1) Distinguish the following cases,

- a) If P is honest, upon P sending $(\text{unbind}, \tilde{\text{id}}) \xrightarrow{t} \mathcal{F}$, or be invoked by other procedures (let t be the current round), $(\text{unbind}, \tilde{\text{id}}, \Phi^P.\text{state}) \xrightarrow{t} \mathcal{C}(\Phi.\text{c}_a)$ on behalf of P . Denote it reaches \mathcal{L} in round $\tau \leq t + \Delta$.
- b) If P is corrupt, upon P sending $(\text{unbind}, \tilde{\text{id}}, W) \xrightarrow{t} \mathcal{C}(\Phi.\text{c}_a)$, $(\text{unbind}, \tilde{\text{id}}) \xrightarrow{t} \mathcal{F}$ on behalf of P . Denote it reaches \mathcal{L} in round $\tau \leq t + \Delta$. If the pending update is confirmed in W , send confirmation to \mathcal{F} on behalf of P .

2) For $Q \in \Phi.\text{other-users}(P)$, distinguish the following cases,

- a) If Q is honest, send $(\text{unbind}, \tilde{\text{id}}, \Phi^Q.\text{state}) \xrightarrow{\tau} \mathcal{C}(\Phi.\text{c}_a)$ on behalf of Q .
- b) If Q is corrupt and sends a state that confirms the pending state, send confirmation to \mathcal{F} on behalf of Q .

3) If any corrupt user does not send message to $\mathcal{C}(\Phi.\text{c}_a)$, send $(\text{timeout}, \tilde{\text{id}}) \xrightarrow{\tau+\Delta} \mathcal{C}(\Phi.\text{c}_a)$ on behalf of honest users.

4) If the pending update is confirmed, for channel $\gamma \in \Phi.\text{channels}$ that has not been initiated closing: go to step (2) of the simulation on procedure (D).

(F) Close

Case P is honest and Q is corrupt

Upon P sending $(\text{close}, \text{id}) \xrightarrow{t} \mathcal{F}$, or be invoked by other procedures (let t be the current round), $(\text{close}, \beta^P.\text{state}) \xrightarrow{t} \mathcal{C}(\beta)$ on behalf of P . Assume it reaches \mathcal{L} in round $\tau \leq t + \Delta$. If Q sends a state with the pending update confirmed to $\mathcal{C}(\beta)$ in round τ , send the confirmation to \mathcal{F} on behalf of Q . Distinguish the following cases,

- 1) If $\beta.\text{bind} = \emptyset$ and Q sends no message to $\mathcal{C}(\beta)$, send (close-resolve) to $\mathcal{C}(\beta)$ on behalf of P in round $\tau + \Delta$.
- 2) Otherwise, proceed as follows,

- a) For each bind in $\beta.\text{bind}$, proceed procedure (E) in round t taking P as the initiator.
- b) Let $\tau_1 \leq t + 3\Delta$ denote the round when all channels are unbound, send $(\text{close-resolve}) \xrightarrow{\tau_1} \mathcal{C}(\beta)$ on behalf of P .

Case P is corrupt and Q is honest

Upon P sending $(\text{close}, W) \xrightarrow{t} \mathcal{C}(\beta)$, send $(\text{close}, \text{id}) \xrightarrow{t} \mathcal{F}$ on behalf of P . Assume it reaches \mathcal{L} in round $\tau \leq t + \Delta$. If W confirms the pending channel update, send confirmation to \mathcal{F} on behalf of P . Distinguish the following cases,

- 1) If $\beta.\text{bind} = \emptyset$, send $(\text{close}, \beta^Q.\text{state}) \xrightarrow{\tau} \mathcal{C}(\beta)$ on behalf of Q .
- 2) Otherwise, proceed as follows,
 - a) For each bind in $\beta.\text{bind}$, if the unbinding has not been initiated in round τ , go to procedure (E) taking Q as the initiator.
 - b) Let $\tau_1 \leq \tau + 3\Delta$ be the round when all channels are unbound, send $(\text{close-resolve}) \xrightarrow{\tau_1} \beta$ on behalf of Q .

Fig. 13: Simulation: (D) Bind Update, (E) Unbind, and (F) Close.