General State Channel Networks

Stefan Dziembowski[1,*], Sebastian Faust[2,**], and Kristina Hostáková[3,**]

[1] stefan.dziembowski@crypto.edu.pl; University of Warsaw, Poland
[2] sebastian.faust@cs.tu-darmstadt.de; Technische Universität Darmstadt, Germany
[3] kristina.hostakova@crisp-da.de; Technische Universität Darmstadt, Germany

Abstract. One of the fundamental challenges that hinder further adaption of decentralized cryptocurrencies is scalability. Because current cryptocurrencies require that all transactions are processed and stored on a distributed ledger – the so-called blockchain – transaction throughput is inherently limited. An important proposal to significantly improve scalability are *off-chain protocols*, where the massive amount of transactions is executed without requiring the costly interaction with the blockchain. Examples of off-chain protocols include *payment channels* and networks, which are currently deployed by popular cryptocurrencies such as Bitcoin and Ethereum. A further extension of payment networks envisioned for cryptocurrencies are so-called *state channel* networks. In contrast to payment networks that only support off-chain payments between users, state channel networks allow execution of arbitrary complex smart contracts. The main contribution of this work is to give the first full specification for general state channel networks. Moreover, we provide formal security definitions and prove the security of our construction against powerful adversaries. An additional benefit of our construction is the use of channel virtualization, which further reduces latency and costs in complex channel networks.

1 Introduction

In recent years we have witnessed a growing popularity of distributed cryptocurrencies such as Bitcoin [23] or Ethereum [33]. The underlying main innovation of these currencies is a consensus mechanism that allows their users to maintain the so-called *blockchain* (or *ledger*). One of the most interesting potential applications of such currencies are the *microtransactions* [32, 28, 21, 26], i.e., transactions of very small values (typically less than 1 cent) that are executed instantaneously. Once implemented, they could enable many novel business models, e.g., fair sharing of WiFi connection, or devices paying to each other in the "Internet of Things".

Unfortunately, blockchain-based systems face inherent challenges that make it very hard, if not impossible, to use them directly for microtransactions. Firstly, each transaction that is processed via the network has to be stored on the blockchain. Moreover, consensus on the blockchain requires significant time to confirm transactions, e.g., in Bitcoin confirmation takes at least around 10 minutes. This imposes a fundamental limit on how many transactions can be processed per second (for instance, the Bitcoin network is currently limited to process up to 7 transactions per second [6]). Finally, the miners that process transactions, ask for fees. Once these fees surpass the actual value assigned to a transaction, micropayments become much less attractive.

A prominent tool for addressing the above challenges are *off-chain channels* [5, 27, 18, 4, 19, 29, 17] that allow two users to rapidly exchange money between each other without sending transactions to the blockchain. Channels are implemented using so-called *smart contracts*, which allow to transfer money according to complex program rules. Below we will first briefly describe this concept, and then give a short introduction to the state of the art in off-chain channels.

Smart contracts. Informally speaking, *smart contracts* (or simply: "contracts") are programmable money, described in form of self-enforcing programs that are published on the ledger. Technically, the term "smart contract" can have two meanings: (1) a contract *code* which is a static object written is some programming language, and (2) a contract *instance* (a dynamic object that executes this code and is running on a

[*] Supported by the NCN grant 2014/13/B/ST6/03540.
[**] Funded by the Emmy Noether Program FA 1320/1-1 of the German Research Foundation (DFG).

blockchain, or inside of a state channel, see below). In the sequel we will often use this distinction (which is similar to the distinction between "programs" and "processes" in operating systems). One can think of a smart contract instance as a trusted third party to which users can send coins, and that can distribute coins between the parties, according to conditions written in its code. Probably the best known currency that supports contracts of an arbitrary complexity is *Ethereum* [33], and its most popular contract language is *Solidity*. In this system, a contract instance never acts by itself, and its actions have to be triggered by the users (who pay the so-called *fees* for every contract execution). The users communicate with the contract instances using *functions calls* (from the contract code). An instance is deployed on the ledger by a call from a special function called *constructor*. For more details on smart contracts and their formal modeling we refer to Sec. 3.

Payment channels. Payment channels are one of the most promising proposals for addressing the scalability challenges in cryptocurrencies. The main idea behind this technology is to keep the massive bulk of transactions off-chain. To this end, the parties that want to *open* a channel deploy a special "channel contract" on the blockchain and lock a certain amount of coins in it. Afterwards they can freely update the channel's balance without touching the ledger. The blockchain is contacted only when parties involved in the payment channel want to *close* the channel, or if they disagree, in which case the channel contract handles fair settlement. In the normal case, when the two parties involved in the payment channel play honestly and off-chain transactions never hit the blockchain before the channel is closed, payment channels significantly improve on the shortcomings of standard blockchain-based payments mentioned above: they limit the load put on the blockchain, allow for instantaneous payments, and reduce transaction fees.

The idea of payment channels has been extended in several directions. One of the most important extensions are the so-called *payment networks*, which enable users to route transactions via intermediary hubs. To illustrate this concept, suppose that P_1 has a payment channel with P_2, and P_2 has a payment channel with P_3. A channel network allows P_1 to route payments to P_3 via the intermediary P_2 without the need for P_1 and P_3 to open a channel between each other on the ledger. This reduces the on-chain transaction load even further. The most well known example of such a system is the *Ligthning network* that has been designed and implemented by Poon and Dryja over Bitcoin [27]. It is based on a technique called *hash-locked transactions*, in which each transaction that is sent from P_1 to P_3 is routed explicitly via P_2 – meaning that P_2 confirms that this transaction can be carried out between P_1 and P_3. For further details on hash-locked transactions, we refer the reader to, e.g., the description of the Lightning network [27] and to Appx. A.

Virtual payment channels. An alternative technique for connecting channels has recently been proposed in [11] under the name "channel virtualization". Using this technique two parties can open a virtual channel over two "extended payment channels" running on the ledger.[4] Consider the example already mentioned above, where P_1 and P_3 are not connected by a payment channel, but each of them has an extended payment channel with an intermediary called P_2. In contrast to connecting payment channels via hash-locked transactions, virtual payment channels have the advantage that the intermediary P_2 does not need to confirm each transaction routed via him. As argued in [11], virtual channels can further reduce latency and fees, while at the same time improving availability.[5] To distinguish the standard channels from the virtual ones, the former ones are also called *ledger* channels. In [11] the authors present only a construction of virtual *payment channels* over a *single* intermediary hub, leaving the general construction as an open research problem. Addressing this shortcoming is one important contribution of our work.

State channels. A further generalization of payment channels are *state channels* [1], which radically enrich the functionality of payment channels. Concretely, the users of a state channel can, besides payments, execute complex smart contracts in an off-chain way. Alice and Bob who established a state channel between each

[4] Concretely, the contract representing the extended payment channel offers additional functionality to support connecting two payment channels.

[5] Availability is improved because payments via the virtual channel can be completed even if the intermediary is temporarily off-line.

other can maintain a "simulated ledger for contracts" and perform the execution of contracts on it "without registering them on the real blockchain". This happens as long as the parties do not enter into a conflict. The security of this solution comes from the fact that at any time parties can "register" the current off-chain state of the channel on the real blockchain, and let the channel contract fairly finish the execution of the contract. Examples of use cases for state channels are manifold and include contracts for digital content distribution, online gaming or fast decentralized token exchanges.

In contrast to payment channels, there has been only little work on general state channels.[6] One prominent project whose final goal is to implement general state channels over Ethereum is called *Raiden* [31], but currently it only supports simple payments, and a specification of protocols for full state channel networks has not been provided yet. The concept of an off-chain state maintained by parties was formalized in the work of Miller et al. [22], where it is used as a main building block for the payment channel construction. In contrast to [22], our general state channel construction allows two parties to have a virtual state channel whose opening does not require any interaction with the blockchain. This significantly improves the time complexity and the cost of a state channel creation. To our best knowledge, the only work considering longer general state channels is [9] recently published by Coleman et al. and developed independently from our work. The work of [9] lacks formal definitions and security proofs. On the other hand, it includes several features useful for practical implementation. We are in contact with the authors of [9] and planing collaboration to further improve our construction and move provably secure state channel networks closer to practice.

1.1 Our contribution

As described above, until now there has not been any satisfactory formal construction or security definition of general state channel networks. The main contribution of this work is to address this shortcoming by providing the *first* construction for building state channel networks of arbitrary complexity together with a formal definition and security analysis. Our construction (i) allows users to run arbitrary complex smart contracts off-chain, and (ii) permits to build channels over any number of intermediaries. Below we describe our core ideas in more detail.

Constructing state channel networks. In order to construct the general state channel networks, we follow a modular *recursive* approach where virtual state channels are built recursively on top of ledger or other – already constructed – virtual state channels. For a high-level description of our recursive approach see Sec. 2 (and Fig. 1 therein). As long as everybody is honest, the intermediaries in the virtual channel are contacted only when the channel is opened and when it is closed (and the ledger is never contacted). On the other hand, let us stress that no intermediary can lose its coins even if all other parties are dishonest and every user of a virtual state channel has the guarantee that he can execute a contract created in a virtual state channel even if all other parties collude.

Modeling state channel networks and security proofs. In addition to designing the first protocols for state channel networks, we develop a UC-style model for "state channel networks" – inspired by the universal composability framework introduced in the seminal work of Canetti [7]. To this end, similarly to [11], we model money via a global ledger ideal functionality $\widehat{\mathcal{L}}$ and describe a novel ideal functionality for state channel networks that provide an ideal specification of our protocols. Using our model, we formally prove that our protocols satisfy this ideal specification. Key challenges of our analysis are (i) a careful study of timings that are imposed by the processing of the ledger, and (ii) the need to guarantee that honest parties cannot be forced to lose money by the fact that the contracts are executed off-chain even if all other parties collude and are fully malicious.

We emphasize that in the context of cryptocurrencies, a sound security analysis is of particular importance because security flaws have a direct monetary value and hence, unlike in many other settings, are guaranteed

[6] A state channel that is not application specific and allows to run arbitrarily complex contracts, is called a *general state channel*. Since we consider only general state channels in this work, we usually omit the word "general" for brevity.

to be exploited. The later is, e.g., illustrated by the infamous attacks on the DAO [30]. Thus, we believe that before complex off-chain protocols are massively deployed and used by potentially millions of users, their specification must be analyzed using formal methods as done in our work using UC-style proofs.

Optimistic vs. pessimistic execution times. While constructing our protocols we will provide the "optimistic" and "pessimistic" execution times. The "optimistic" ones refer to the standard case when all parties behave honestly. In the optimistic case all our protocols allow for instantaneous off-chain contract execution, and a possible delay depends only on the latency of the network over which parties communicate. The "pessimistic" case corresponds to the situation when the corrupt parties try to delay the execution as much as they can by forcing contract execution on the blockchain. In our solution the pessimistic execution times grow linearly with the number of intermediaries ℓ involved. Notice that these pessimistic times can in reality happen only in the unlikely case when *all* but one party are corrupt. Since the main goal of this paper is to introduce the general framework, and not to fine-tune the parameters, we leave it as an important direction for future work to improve our construction and optimize these timings, possibly using the techniques of [22].

Further related work One of the first proposals for building payment channels is due to Decker [10], who in particular also introduced a construction for duplex payment channels. An alternative proposal for payment channel networks has been given by Miller et al. [22]. In this work, the authors show how to reduce the pessimistic timings to constant time (i.e., independent of the length of the channel path). It is an interesting question for future work to combine the techniques from [22] with the channel virtualization. Several works focus on privacy in channel networks, path finding or money re-balancing in payment channels [19, 29, 17]. In particular, [22, 19, 11] also provide a UC-based security analysis of their constructions. Channel constructions based on the sequence number maturity (that we also use in this paper) have been mentioned already in [27], and recently described in more detail (as "stateful duplex off-chain micropayment channels") by Bentov et al. in [4]. Another challenge in building and maintaining complex channel networks is the fact that parties have to continuously watch what happens on the blockchain regarding the state of their channels. This problem can be addressed using so-called watchtowers [25, 20], to which users can outsource the task of watching the blockchain.

1.2 Organization of the paper

We begin with an informal description of our state channel construction in Sec. 2, where we explain how state channels are created and how they can be used. The ideal specification of our construction is presented in Sec. 4 and the full description of the state channel protocols is given in Sec. 6 (for ledger state channels) and in Sec. 7 (for virtual state channels). We introduce the necessary formalism and present security and efficiency properties required from a general state channel in Sec. 3. Our modular approach of building state channels is discussed in Sec. 5. Finally, we conclude in Sec. 8.

2 State channel construction

Before we proceed to the more technical part of this work, let us give an intuitive explanation of our virtual state channel construction. We would like to emphasize that the description of our approach as presented in this section is very simplified and excludes many important technicalities. Formal definitions, detailed explanations of our protocols, and their full description are presented later in this paper (see Sections 3–5 and Appendices 6 and 7). As already mentioned in Sec. 1.1, we follow a recursive approach, which is shown for the case of 6 parties on Fig. 1 where we consider parties $P_1, \ldots, P_6$, with each P_i being connected with P_{i+1} via a ledger state channel $P_i \Leftrightarrow P_{i+1}$. To build a virtual state channel $\gamma_4 := P_1 \leftrightarrow P_6$, we first create a virtual state channel $\gamma_1 := P_1 \leftrightarrow P_3$ using ledger state channels $P_1 \Leftrightarrow P_2$ and $P_2 \Leftrightarrow P_3$. Then a virtual state channel $\gamma_2 := P_4 \leftrightarrow P_6$ is created using ledger state channels $P_4 \Leftrightarrow P_5$ and $P_5 \Leftrightarrow P_6$. The other virtual state channels are created recursively, as follows: first, channel $\gamma_3 := P_1 \leftrightarrow P_4$ is created using the virtual state channel γ_1 and the ledger state channel $P_3 \Leftrightarrow P_4$, and then channel γ_4 is created using the virtual state channels γ_3 and γ_2.

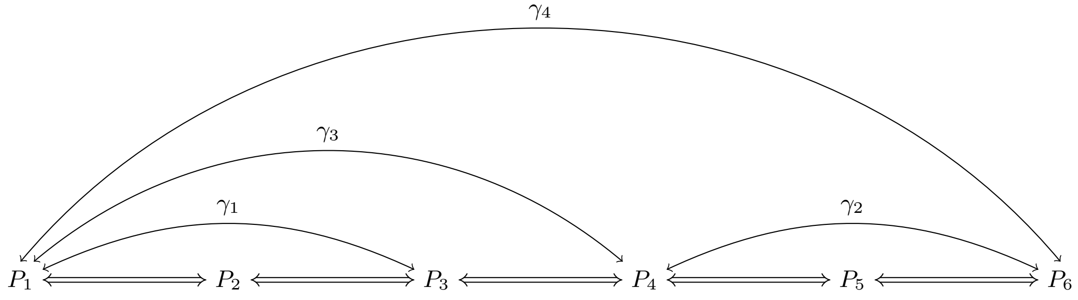

Fig. 1: Example of a recursive construction of a virtual state channel γ_4 (of length 5) between P_1 and P_6.

Ledger state channels – an overview. The terminology for ledger state channels is given in Sec. 3. and their construction can be found in Sec. 6. Below we explain only the main idea of the ledger state channel construction. A ledger state channel δ between Alice and Bob allows them to execute off-chain instances of some contract code $\texttt{C}$. An example could be a lottery game contract $\texttt{C}_{\texttt{lot}}$, where each user deposits 1 coin and then one user is randomly chosen to receive 2 coins. Technically, this is implemented using the standard cryptographic method based on commitment schemes (see, e.g., [2]), where the execution of the contract happens in the following steps: first the parties deposit their coins in the contract instance (call the resulting *initial state* of the game G_0)[7], then Alice sends to the contract her commitment to a random bit $r_A \in \{0, 1\}$ (which results in state G_1), afterwards Bob sends his random bit $r_B \in \{0, 1\}$ to the contract (denote the resulting state G_2). Then, Alice opens her commitment, the final state G_3 is computed, and 2 coins are given to Alice if $r_A \oplus r_B = 0$, or to Bob (otherwise). Finally, the contract instance terminates. Technically, the previous steps are implemented via function calls. For example: sending a bit r_B by Bob can be implemented as function call $\texttt{Reveal}(r_B)$ (where $\texttt{Reveal}$ is a function available in $\texttt{C}_{\texttt{lot}}$ that stores r_B in the storage of the contract).

As described in Sec. 1, two parties create a ledger state channel by deploying a *state channel contract*, ($\texttt{SCC}$), in which each party locks some amount of coins. Once the ledger state channel δ is established, parties can open instances of the contract code $\texttt{C}$ in the channel and execute them. For example the parties can open a channel in which each of them locks 10 coins and then run several instances of the lottery contract $\texttt{C}_{\texttt{lot}}$ in this channel. Every contract instance locks 1 coin of each party (from the coins that are locked in channel δ). A locked coin cannot be used for any other contract instance in δ. Once the contract instance terminates, the coins are unlocked and distributed back to the channel δ according to the rules of $\texttt{C}$. The state channel contract on the blockchain guarantees that if something goes wrong during the off-chain execution (parties disagree on a state of some contract instance, one of the parties stops communicating, etc.), they can always fairly resolve their disagreement and continue the execution via the state channel contract on the blockchain.

Off-chain contract execution in the ledger state channels. Let us now take a closer look how the off-chain contract execution is done via the ledger state channel. Let $\texttt{C}$ be a contract code, and let G denote the (dynamically changing) instance of $\texttt{C}$ that is executed in δ. To deploy G both parties agree on the initial state G_0 of G. The parties then exchanging signatures on $(G_0, 0)$. The rest of the execution is done by exchanging signatures on further states of G together with indices w that denote the *version numbers.* Assume that Alice wants to call a function f (with some parameters m) in the contract instance. Let (G_w, w) be the last state of the contract instance G on which the parties exchanged their signatures. She then (1) computes locally the new value G_{w+1} of the state, by calling $f(m)$ on G_w, and then (2) sends signed $(G_{w+1}, w + 1)$

[7] A reader familiar with Ethereum may object that "simultaneous" contract instance deployment is not allowed (as Ethereum does not support "multi-input" transactions). We stress that the example above illustrates a contract that is run "inside of a channel" (not on blockchain) and is compatible with our construction.

together with f and m to Bob. Bob checks if Alice's computation was correct, and if yes then he replies with his signature on $(G_{w+1}, w+1)$. When the instance G terminates, the coins resulting from this execution are distributed between the parties according to the outcome of the game.

For example if G is an instance of the lottery contract $\texttt{C}_{\texttt{lot}}$ described above then the states of the game are G_0, G_1, G_2 and G_3. Since the first move of the game is done by Alice, she locally computes the new state G_1 and sends it to Bob together with her commitment to r_A and her signature on $(G_1, 1)$. Then Bob replies with his signature on $(G_1, 1)$. Thereafter, Bob makes his move, i.e., he computes G_2, sends signed $(G_2, 2)$ together with his random bit r_B to Alice, and so on. Note that the interaction of the parties with the contract instance is always "local", i.e., the parties themselves compute the new states of G and then just exchange signatures.

As long as both Alice and Bob are honest, everything is done without any interaction with the blockchain. If, however, one party cheats (e.g. by refusing to communicate), the other party can always ask the SCC contract to finish the game. The version number w is used to make sure that SCC gets always the latest state of the game. More concretely: the contract is constructed in such a way that if a malicious party submits an old state, then the other party can always "overwrite" this state by providing a signed state of the contract instance with a higher version number. Once the SCC contract learns the latest state G_w, the game can be finished (starting from G_w) on-chain via SCC.

Virtual state channels – an overview. As described above, the virtual state channels are constructed recursively "on top" of the ledger state channels. Suppose that Alice and Bob want to run some contract code C (e.g. the lottery game) in an off-chain way in γ. This time, however, they do not have an open ledger state channel between each other. Instead, both Alice and Bob have a channel with a third party, which we call Ingrid. Denote these channels α and β respectively. With the help of Ingrid but *without interacting with the blockchain*, Alice and Bob can open a virtual state channel γ that has the same functionality and provides the same guarantees as if it would be a ledger state channel between them. In particular, Alice and Bob are allowed to create a contract instance of C in their channel γ and execute it just by communicating with each other (i.e. play their game without talking to any third party or the blockchain).

Recall that in case of the ledger state channels every dispute between Alice and Bob is resolved by the state channel contract, SCC. For the virtual state channel γ the role of such a "judge" is played by Ingrid. The main difference from the previous case is that, unlike SCC (that is executed on the ledger), *Ingrid cannot be trusted*, and in particular, she may even collude with a corrupt Alice or Bob. In order to prevent parties from cheating, we create special contracts in each of the ledger state channels α and β. Their code will be called "virtual state channel contract" (VSCC) and their instances will be denoted ν_α and ν_β, respectively. The instances ν_α provides security guarantees for Alice, and ν_β for Bob. In addition, both contract instances together provide guarantees for Ingrid. The contract code VSCC has to depend on the code C since it needs to interpret the code C in case the parties enter into a dispute (see below). Note that SCC depends on VSCC, and hence, indirectly, on C. This dependence is summarized in Fig. 2.

Creating the virtual state channel. Let us explain the virtual state channel creation in more detail. In the first step Alice and Bob inform Ingrid about their intention to use her as an *intermediary* for their virtual state channel γ. Alice does so by proposing to open an instance ν_α of VSCC in the channel α. This instance will contain all information about the virtual state channel γ (for example: how many coins each party wants to lock in the channel). In some sense ν_α can be viewed as a "copy" of the virtual state channel γ in which Ingrid plays the role of Bob — for example, if the initial balance in γ is 1 coin for Alice and 5 coins for Bob, then Alice would lock 1 coin and Ingrid 5 coins in ν_α. Symmetrically, Bob proposes a new instance ν_β of VSCC in the ledger state channel β that can be viewed as a "copy" of the virtual state channel γ in which Ingrid plays the role of Alice. In the example above, Ingrid would lock 1 coin and Bob 5 coins in ν_β. If Ingrid receives both proposals and she agrees to be the intermediary of the virtual state channel γ, she confirms both requests.

Contract execution in the virtual state channel γ. The *off-chain* contract execution in the virtual state channel is performed exactly in the same way as in case of the ledger state channels (see paragraph "Off-

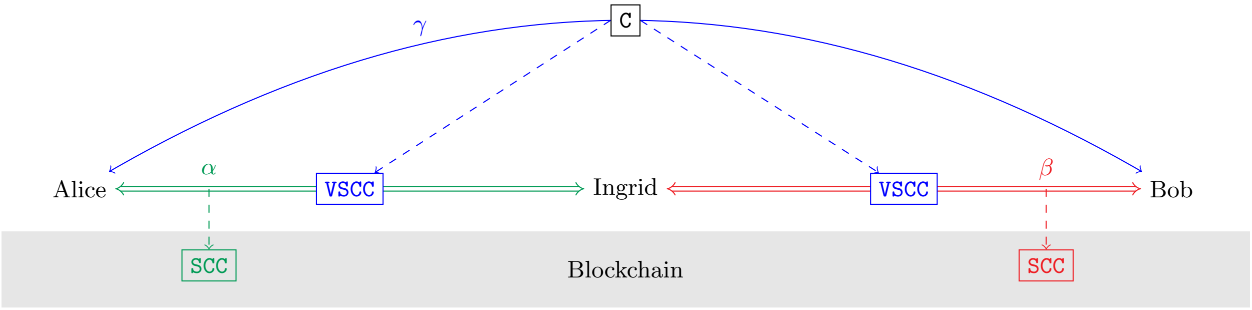

Fig. 2: Construction of a virtual state channel γ_1 of length 2 in which a contract instance of C is created.

chain contract execution in the ledger state channels" above). That is, as long as both Alice and Bob are honest, they execute a contract instance G by exchanging signatures on new versions of the game states without talking to Ingrid at all, and without updating ν_α and ν_β. The case when Alice and Bob disagree needs to be handled differently, since the parties cannot contact the blockchain contract, but have to resolve this situation using the channels α and β that they have with Ingrid. Consider, for example, the situation when, in the scenario described above, Bob is malicious and stops communicating with Alice, i.e. he does not send back his signature on $(G_{w+1}, w+1)$. In this situation, Alice has to make her move "forcefully" by using the channel α she has with Ingrid. More concretely, she will execute the contract instance ν_α. It is very important to stress that the virtual state channel construction uses this instance in a *black-box way*, i.e., when describing the protocols for virtual state channel execution this protocol uses the execution of ν_α in a black-box way via the interface of the underlying channel. Internally, of course this is done by a protocol between Alice and Ingrid realizing the off-chain execution of ν_α (as long as Alice and Ingrid are honest).

First, Alice starts the "state registration procedure". The goal is to let ν_α know that she has a disagreement with Bob, and to convince ν_α that G_w is the latest state of the contract instance G. To this end, she sends to ν_α the state (G_w, w, s_B), where s_B is Bob's signature on (G_w, w). She does it by calling a function "register" (see Step 1 on Fig. 3). Of course ν_α has no reason to believe Alice that this is really the latest state of G. Therefore ν_α forwards this message to Ingrid[8], that, in turn, calls a function "register(G_w, w, s_B)" of the contract instance ν_β in channel β (see Step 2). Bob now replies (in Step 3) to ν_β with his latest version of the contract instance (i.e. he calls "register$(G_{w'}, w', s_A)$", where s_A is Alice's signature). When Ingrid learns about Bob's version from ν_β, she forwards this information to ν_α (see Step 4). Suppose that $w > w'$, i.e., Alice is honest, and Bob is cheating by submitting and old version of the instance (the other case is handled analogously). Then, both ν_α and ν_β decide that (G_w, w) is the latest version of G (i.e. they "register G_w").

From the point of view of Ingrid, the most important security feature of this procedure is that there is a consensus among ν_α and ν_β about the latest state of G (even if Alice and Bob are *both* dishonest and playing against her). This consensus will be maintained during the entire execution of G in instances ν_α and ν_β. This is important, as otherwise she could lose coins.[9] This invariant will be maintained throughout the rest of the "forced execution procedure".

After the state registration is over, Alice calls (in Step 5, Fig. 3) a function "execute$(f(m))$" of ν_α, "asking" ν_α to execute $f(m)$ on the contract instance G starting from the registered state (G_w, w). Since we want to maintain the "consensus invariant" mentioned above, we cannot simply let ν_α perform this execution immediately after it receives this call. This is because some contracts may allow both parties to

[8] Recall again that this execution is realized by a protocol between Alice and Ingrid.

[9] Imagine, e.g, that the final state of G in ν_α is that Alice gets all the coins locked in G, and the final state of G in ν_β is that Bob gets all the coins locked in G. Then Ingrid loses these coins in *both* channels α and β.

call functions at the same time[10], and Bob could simultaneously call some other function execute$(f'(m'))$ of ν_β. This situation is especially subtle because function execution is generally not commutative, i.e., executing $f(m)$ and then $f'(m')$ can produce a different result than doing it in the different order. Consequently, this could result in ν_α and ν_β having different states of their local copies of γ. We solve this problem by delaying the execution of $f(m)$ until it is clear that no other function can be executed before $f(m)$. More precisely, the contract code `VSCC` is defined in such a way that $f(m)$ is only stored in the storage of the contract instance ν_A, resp. ν_B together with a time stamp of storage. The internal execution of $f(m)$ in ν_A, resp. ν_B, is performed only when the contract instance is being terminated (which happens when then virtual state channel γ is being closed).

Let us emphasize that the purpose of the description above is to explain the concepts and main ideas of our construction. The final protocol, however, works slightly differently due to several optimizations. For example, in order to decrease the pessimistic time complexity, the registration phase and the force execution phase for virtual state channels are run in parallel (i.e. Step 1 and Step 5 are happening in the same round). We refer the reader to Sec. 7 for the full description of our protocol.

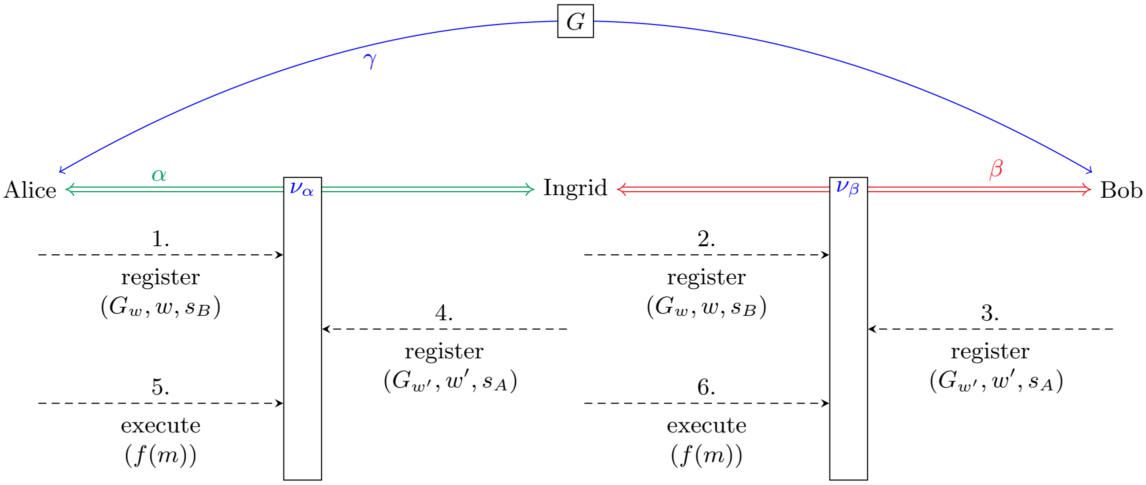

Fig. 3: Illustration of the forced execution process from our example in which Alice and Bob have a virtual state channel γ in which they opened a contract instance. Only the function calls are shown (the messages sent by the contracts are omitted).

Applying recursion. As already highlighted earlier, longer virtual state channels are constructed recursively. The key observation that enables this recursion is that the state channels α and β that are used to build γ are accessed in a "black-box" way. In other words, the only property of α and β needed in the construction of γ is that one can execute off-chain contracts in them. This "black-box" property guarantees that our virtual state channel construction works also if the channels α and β are *virtual* (not ledger), or in case one of them is virtual, and the other one is ledger.

Let us illustrate this by taking a look again at the situation depicted in Fig. 1. Consider first the virtual state channel γ_3 – a virtual state channel of length 3 build on top of a virtual state channel γ_1 of length 2 and the ledger state channel $P_3 \Leftrightarrow P_4$. Assume that `C` is the contract code whose instances can be opened in γ_3. Following the construction described earlier in this section, γ_3 can be created if both the underlying

[10] Note that it is *not* the case of the $\texttt{C}_{\texttt{lot}}$ contract, since there its always clear which party is expected to "make a move" in the game. However, in general, we do not want to have such restrictions on contracts in this paper.

state channels γ_1 and $P_3 \Leftrightarrow P_4$ support contract instances of the virtual state channel contract $\mathtt{VSCC}$ which depends on $\mathtt{C}$. This, in particular, implies that the ledger state channels $P_1 \Leftrightarrow P_2$ and $P_2 \Leftrightarrow P_3$, on top of which the virtual state channel γ_1 is created, must support contract instances of the virtual state channel contract $\mathtt{VSCC}'$ which depends on $\mathtt{VSCC}$ (thus indirectly also on $\mathtt{C}$).

This reasoning can be repeated for longer channels. For example, if $\mathtt{C}$ is a contract code whose instances can be opened in the virtual state channel γ_4, then contract instances of $\mathtt{VSCC}$ must be supported by both γ_2 and γ_3, contract instances of $\mathtt{VSCC}'$ must be supported by γ_1, $P_3 \Leftrightarrow P_4$, $P_4 \Leftrightarrow P_5$ and $P_5 \Leftrightarrow P_6$. Finally, contract instances of the virtual state channel contract $\mathtt{VSCC}''$, which depends on $\mathtt{VSCC}'$, must be supported by the ledger state channels $P_1 \Leftrightarrow P_2$, $P_2 \Leftrightarrow P_3$. More details of this recursion, including the analysis of pessimistic and optimistic timing, are provided in further sections. Let us just mention here that in order to achieve linear pessimistic time complexity (in the channel length), our construction assumes that virtual state channels are built in a balanced way as in Fig. 1 (i.e. the two state channels used to build a virtual state channel have approximately the same length).

The notion of time In the description above we ignored the notion of time. This was done to simplify this informal description. We define this notion in the technical part of the paper (see Sec. 3.3). In our construction parties are always aware of the current time, and they pass the time information to the contract functions in the state channels. Time is modeled as a natural number, and the time unit is called a *round* (think of it as 1 second, say).

Key features of our construction. An important property of our construction and our model is that we support full concurrency. That is, we allow several virtual state channels to be created simultaneously over the same ledger state channels, and allow parties to be involved in several concurrent executions of (possibly complex) contracts. This is possible because our ledger state channels can store and execute several contracts "independently" (i.e., these are multi-contract state channels).

Another important feature of our modular construction is that it naturally allows for building channels via multiple (possible incompatible) cryptocurrencies as long as they have a sufficiently complex scripting language (in particular, they allow to deploy a state channel contract). For illustration, consider Alice having a ledger state channel with Ingrid in cryptocurrency called "A-coin", and Bob having a ledger state channel with Ingrid in cryptocurrency called "B-coin". Now, Alice and Bob can build a virtual state channel over Ingrid, where Alice (resp. Bob) is oblivious of the details of B-coin (resp. A-coin). This makes sense as long as the exchange rate between the currencies does not change too much during the lifetime of the virtual channel. Note that, since the virtual channel opening and closing does not require interacting with the ledger, the lifetime of a virtual state channel can be made very short (minutes or hours). In addition, virtual state channels also improve on privacy. This is the case because channel updates are fully P2P and do not require involvement of intermediaries.

Finally, we point out that our concept of higher-level channel virtualization has the key feature that it adds further "layers of defense" against malicious parties before honest users need to communicate with the blockchain. Consider, for example, the situation shown in Fig. 1. Even if P_6 and the intermediary P_4 in the virtual state channel γ_4 are corrupt, then P_1 can resolve possible conflicts via the intermediary P_3 using the virtual state channel γ_1, i.e. P_1 does not need to communicate with the ledger.

3 Definitions and Security model

In the sequel, following [11], we present tuples of values using the following convention. The individual values in a tuple T are identified using keywords called *attributes*: $\mathsf{attr1}, \mathsf{attr2}, \ldots$. Strictly speaking an *attribute tuple* is a function from its set of attributes to $\{0,1\}^*$. The *value of an attribute* attr in a tuple T (i.e. $T(\mathsf{attr})$) will be referred to as $T.\mathsf{attr}$. This convention will allow us to easily handle tuples that have dynamically changing sets of attributes. We assume that $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$ is a signature scheme that is existentially unforgeable against a chosen message attack (see, e.g., [15]). The ECDSA scheme used in Ethereum is believed to satisfy this definition.

3.1 Definitions of contracts and channels

We now present our syntax for describing contracts and channels. The notation presented in this section can be viewed as an extension of the one used in [11]. In the rest of this paper we assume that the set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of parties that use the system is fixed.

Contracts We consider only contracts executed between two parties. A *contract storage* is an attribute tuple σ that contains at least the following attributes: (1) $\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R \in \mathcal{P}$ that denote the users involved in the contract, (2) $\sigma.\mathsf{locked} \in \mathbb{R}_{\geq 0}$ that denotes the total amount of coins that is locked in the contract and (3) $\sigma.\mathsf{cash}\colon \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\} \to \mathbb{R}$ that denotes the amount of coins that the users have freely available. It must hold that $\sigma.\mathsf{locked} \geq \sigma.\mathsf{cash}(\sigma.\mathsf{user}_L) + \sigma.\mathsf{cash}(\sigma.\mathsf{user}_R)$. Let us explain the difference between locked coins and freely available coins as well as the above inequality on a concrete example. Assume that parties are playing a game where each party initially invests 5 coins. During the game, parties make a bet, where each party puts 1 coin in the "pot". Now the amount of coins *locked* in the game did not change, it is still equal to 10 coins; however, the amount of *freely available* coins decreased (each party has only 4 freely available coins). In addition to the attributes mentioned above, a contract storage may contain other application-specific data.

We will now define formally the notion of *contract code* that was already described informally in Sec. 1. Formally a contract code consists of some functions (in Ethereum they are written in Solidity) that operate on contract storage. The set of possible contract storages is usually restricted (e.g. the functions expect that it has certain attributes defined). We call the set of restricted storages the *admissible contract storages* and typically denote it Λ.

Formally, we define a *contract code* as a tuple $\mathtt{C} = (\Lambda, g_1, \ldots, g_r, f_1, \ldots, f_s)$, where Λ are the admissible contract storages and $g_1, \ldots, g_r$ are functions called *contract constructors*, and $f_1, \ldots, f_s$ are called *contract functions*. Each contract constructor g_i is function that takes as input a tuple (P, τ, z), with $P \in \mathcal{P}, \tau \in \mathbb{N}$, and $z \in \{0,1\}^*$, and produces as output an admissible contract storage σ or a special symbol $\perp$ (in which case we say that the contract construction *failed*). The meaning of these parameters is as follows: P is the identity of the party that called the function, τ is the current round (see Sec. 3.3 for more on how we model time and rounds), and z is used to pass additional parameters to g_i. The constructors are used to create a new instance of the contract. If the contract construction did not fail, then $g_i(P, \tau, z)$ is the initial storage of a new contract instance.

Each contract function f_i takes as input a tuple (σ, P, τ, z), with $\sigma \in \Lambda$ being an admissible contract storage, $P \in \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\}$, $\tau \in \mathbb{N}$ and $z \in \{0,1\}^*$ (the meaning of this parameters is as before). It outputs a tuple $(\tilde{\sigma}, add_L, add_R, m)$, where $\tilde{\sigma}$ is the new contract storage (that replaces contract storage σ in the contract instance), values $add_L, add_R \in \mathbb{R}_{\geq 0}$ correspond to the amount of coins that were *unlocked* from the contract storage to each user (as a result of the execution of f_i), and $m \in \{0,1\}^* \cup \{\perp\}$ is an *output message*. If the output message is $\perp$, we say that the execution *failed* (we assume that the execution always fails if a function is executed on input that does not satisfy the constraints described above, e.g., it is applied to σ that is not admissible). If the output message $m \neq \perp$, then we require that $\tilde{\sigma}$ is an admissible contract storage and the attributes user_L and user_R in $\tilde{\sigma}$ are identical to those in σ. In addition, it must hold that $add_L + add_R = \sigma.\mathsf{locked} - \tilde{\sigma}.\mathsf{locked}$. Intuitively, this condition guarantees that executions of a contract functions can never result in unlocking more coins than what was originally locked in the contract storage.

As described in Sec. 1 a *contract instance* represents an instantiation of a contract code. Formally, a contract instance is an attribute tuple ν with a contract $\mathsf{storage}$ and code, where $\nu.\mathsf{code} = (\Lambda, g_1, \ldots, g_r, f_1, \ldots, f_s)$ is a contract code, and $\nu.\mathsf{storage} \in \Lambda$ is a contract storage.

Ledger state channel. We next present our terminology for ledger state channels, which is inspired by the notation for payment channels used in [11]. Formally, a ledger state channel γ is defined as an attribute tuple $\gamma := (\gamma.\mathsf{id}, \gamma.\mathsf{Alice}, \gamma.\mathsf{Bob}, \gamma.\mathsf{cash}, \gamma.\mathsf{cspace})$. We call the attribute $\gamma.\mathsf{id} \in \{0,1\}^*$ the identifier of the ledger state channel. Attributes $\gamma.\mathsf{Alice} \in \mathcal{P}$ and $\gamma.\mathsf{Bob} \in \mathcal{P}$ are the identities of parties using the ledger state channel γ. For convenience, we also define the set $\gamma.\mathsf{end\text{–}users} := \{\gamma.\mathsf{Alice}, \gamma.\mathsf{Bob}\}$ and the function $\gamma.\mathsf{other\text{–}party}$ as $\gamma.\mathsf{other\text{–}party}(\gamma.\mathsf{Alice}) := \gamma.\mathsf{Bob}$ and $\gamma.\mathsf{other\text{–}party}(\gamma.\mathsf{Bob}) := \gamma.\mathsf{Alice}$. The attribute $\gamma.\mathsf{cash}$ is a function mapping the set $\gamma.\mathsf{end\text{–}users}$ to $\mathbb{R}_{\geq 0}$ such that $\gamma.\mathsf{cash}(T)$ is the amount of coins the party $T \in \gamma.\mathsf{end\text{–}users}$ has locked in the ledger state channel γ. Finally, the attribute $\gamma.\mathsf{cspace}$ is a partial function that is used to describe the set of all contract instances that are currently open in this channel. It takes as input a *contract instance identifier cid* $\in \{0,1\}^*$ and outputs a contract instance ν such that $\{\nu.\mathsf{storage.user}_L, \nu.\mathsf{storage.user}_R\} = \gamma.\mathsf{end\text{–}users}$. We will refer to $\gamma.\mathsf{cspace}(cid)$ as the *contract instance with identifier cid in the ledger state channel* γ.

Virtual state channel. Formally, a virtual state channel γ is a tuple $\gamma := (\gamma.\mathsf{id}, \gamma.\mathsf{Alice}, \gamma.\mathsf{Bob}, \gamma.\mathsf{Ingrid}, \gamma.\mathsf{subchan}, \gamma.\mathsf{cash}, \gamma.\mathsf{cspace}, \gamma.\mathsf{length}, \gamma.\mathsf{validity})$. The attributes $\gamma.\mathsf{id}$, $\gamma.\mathsf{Alice}$, $\gamma.\mathsf{Bob}$, $\gamma.\mathsf{cash}$ and $\gamma.\mathsf{cspace}$, are defined as in the case of a ledger state channel. The same holds for the set $\gamma.\mathsf{end\text{–}users}$ and the function $\gamma.\mathsf{other\text{–}party}$. The new attribute $\gamma.\mathsf{Ingrid} \in \mathcal{P}$ denotes the identity of the intermediary of the virtual state channel. For technical reasons (see Sec. 7.2 on Page 31), we restrict $\gamma.\mathsf{cspace}$ for virtual state channels to contain only a single contract instance. We emphasize that this is not a restrictions of the functionality since ledger state channels support an arbitrary number of contract instances, and hence we can build any number of virtual state channels.

The attribute $\gamma.\mathsf{subchan}$ is a function mapping the set $\gamma.\mathsf{end\text{–}users}$ to $\{0,1\}^*$. The value of $\gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$ equals the identifier of the ledger/virtual state channel between $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Ingrid}$. Analogously for the value of $\gamma.\mathsf{subchan}(\gamma.\mathsf{Bob})$. We often call these channels the *subchannels* of the virtual state channel γ. The attribute $\gamma.\mathsf{validity}$ denotes the round in which the virtual state channel γ will be closed (see Sec. 3.3 for more on the notion of rounds). The reason to have this parameter is to ensure that the channel γ will not remained open forever. Otherwise $\gamma.\mathsf{Ingrid}$ could have her money blocked forever, as (unlike $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$) she cannot herself request the channel closing. Finally, the attribute $\gamma.\mathsf{length} \in \mathbb{N}_{>1}$ refers to the length of the virtual state channel, i.e., the number of ledger state channels over which it is built. For example in Fig. 1 (see Page 5) we have: $\gamma_1.\mathsf{length} = 2$, $\gamma_2.\mathsf{length} = 2$, $\gamma_3.\mathsf{length} = 3$, $\gamma_4.\mathsf{length} = 5$. Sometimes it will be convenient to say that ledger state channels have length one.

3.2 Security and efficiency goals

Before presenting our formal security model in Sec. 3.3, let us start by listing some security guarantees that are desirable for a state channel network. In the following description, if it is not important whether γ is ledger state channel or a virtual state channel, and hence we will refer to γ as a *state channel*.

(1) **Consensus on creation:** A state channel γ can be successfully created only if all users of γ agree with its creation.
(2) **Consensus on updates:** A contract instance in a state channel γ can be successfully updated (this includes also creation of the contract instance) only if both end-users of γ agree with the update.
(3) **Guarantee of execution:** An honest end-user of a ledger state channel γ can execute a contract function f of a created contract instance in any round τ_0 on input value z even if the other end-user of γ is corrupt. This property holds also for virtual state channels with the restriction that $\tau_0 < \gamma.\mathsf{validity}$.
(4) **Balance security:** The intermediary of a virtual state channel γ never loses coins even if both end-users of γ are corrupt.

While property (4) provides a strong monetary security guarantee to the intermediary of a virtual state channel, the guarantees for the end-users given by properties (2) and (3) only ensure that party can not be forced to create a contract instance and that contract instances can be executed at any time. We emphasize that this is similar to what is guaranteed by the ledger to on-chain contracts. Concretely, this means that if the contract rules allow that a certain end-user may lose money (e.g., by losing the lottery as described in the example from Sec. 2), then this is not in violation with the security properties guaranteed by a state channel network.

In addition to the security properties, we identify the following two efficiency goals. Below, by constant number of rounds we mean that the required rounds for executing the procedure is independent of the channel length and the ledger delay Δ (looking ahead, the parameter Δ models the fact that changes on a blockchain come with a certain delay, see Sec. 3.3 for more details).

(1) **Constant round optimistic update/execute:** In the optimistic case when both end-users of a state channel γ are honest, they can update/execute a contract instance in γ within a constant number of rounds.
(2) **Constant round virtual state channel creation:** Successful creation of a virtual state channel γ takes a constant number of rounds.

3.3 Our model

To formally model the security of our construction, we use a UC-style model following the works of [3, 11] that consider protocols that operate with *coins*.[11] In particular, our model uses a synchronous version of the global UC framework (GUC) [8] which extends the standard UC framework [7] by allowing for a global setup.

Protocols and adversarial model. We consider an *n-party protocol* π that runs between parties from the set $\mathcal{P} = \{P_1, \ldots, P_n\}$ which are connected by authentic communication channels. A protocol is executed in the presence of an *adversary* Adv that takes as input a security parameter 1^λ (with $\lambda \in \mathbb{N}$) and an auxiliary input $z \in \{0,1\}^*$, and who can *corrupt* any party P_i at the beginning of the protocol execution (so-called static corruption). By corruption we mean that Adv takes full control over P_i including learning its internal state. Parties and the adversary Adv receive their inputs from a special party – called the *environment* $\mathcal{Z}$ – which represents anything "external" to the current protocol execution. The environment also observes all outputs returned by the parties of the protocol. In addition to the above entities, the parties can have access to ideal functionalities $\mathcal{G}_1, \ldots, \mathcal{G}_m$. In this case we say that the protocol *works in the* $(\mathcal{G}_1, \ldots, \mathcal{G}_m)$*-hybrid model*.

Modeling communication and time. We assume a synchronous communication network, which means that the execution of the protocol happens in rounds. Let us emphasize that the notion of rounds is just an abstraction which simplifies our model (see, e.g, [13, 14, 24, 16] for a formalization of this model and its relation to the model with real time). Whenever we say that some operation (e.g. sending a message or simply staying in idle state) *takes at most* $\tau \in \mathbb{N} \cup \{\infty\}$ *rounds* we mean that it is up to the adversary to decide how long this operation takes (as long as it takes at most τ rounds). Let us now discuss the amount of time it takes for different entities to communicate with each other. The communication between two parties P_i takes exactly one round. All other communication – for example, between the adversary Adv and the environment $\mathcal{Z}$ – takes zero rounds. For simplicity we assume that any computation made by any entity takes zero *rounds* as well.

Handling coins. We follow [11] and model the money mechanics offered by crypotcurrencies such as Bitcoin or Ethereum via a global ideal functionality $\widehat{\mathcal{L}}$ using the *global UC (GUC)* model [8]. The state of the ideal functionality $\widehat{\mathcal{L}}$ is public and can be accessed by all parties of the protocol π, the adversary Adv and the environment $\mathcal{Z}$. It keeps track on how much money the parties have in their accounts by maintaining a vector of non-negative (finite precision) real numbers $(x_1, \ldots, x_n)$, where each x_i is the amount of coins that P_i owns.[12]

The functionality $\widehat{\mathcal{L}}$ is initiated by the environment $\mathcal{Z}$ that can also freely add and remove money in user's accounts, via the operations add and remove. While parties $P_1, \ldots, P_n$ *cannot* directly perform operations on $\widehat{\mathcal{L}}$, the ideal functionalities can carry out add and remove operations on the $\widehat{\mathcal{L}}$ (and hence, indirectly, P_i's can also modify $\widehat{\mathcal{L}}$, in a way that is "controlled" by the functionalities). Every time an ideal functionality issues an add or remove command, this command is sent to $\widehat{\mathcal{L}}$ within Δ rounds, for some parameter $\Delta \in \mathbb{N}$. The exact round when the command is sent is determined by the adversary Adv. The parameter Δ models the fact that in cryptocurrencies updates on the ledger are not immediate. We denote a ledger functionality $\widehat{\mathcal{L}}$ with maximal delay Δ by $\widehat{\mathcal{L}}(\Delta)$ and an ideal functionality $\mathcal{G}$ with access to $\widehat{\mathcal{L}}(\Delta)$ by $\mathcal{G}^{\widehat{\mathcal{L}}(\Delta)}$. The ledger functionality $\widehat{\mathcal{L}}$ is formally defined in Fig. 4.

[11] Throughout this work, the word *coin* refers to a monetary unit.
[12] This is similar to the concept of a *safe* of [3].

Functionality $\widehat{\mathcal{L}}$

Functionality $\widehat{\mathcal{L}}$, running with parties $P_1, \ldots, P_n$ and the environment $\mathcal{Z}$, gets as input $(x_1, \ldots, x_n) \in \mathbb{R}^n_{\geq 0}$ (where $\mathbb{R}_{\geq 0}$ are finite-precision non-negative reals). It stores the vector $(x_1, \ldots, x_n)$ and accepts queries of following types:

Adding money

Upon receiving a message $(\mathsf{add}, sid, P_i, y)$ from $\mathcal{Z}$ (for $y \in \mathbb{R}_{\geq 0}$) set $x_i := x_i + y$. We say that *y coins are added to P_i's account in $\widehat{\mathcal{L}}$.*

Removing money

Upon receiving a message $\left(\mathsf{remove}, sid, \{(P_{i_j}, y_{i_j})\}_{j=1}^{t}\right)$ (for some $t \in \{1, \ldots, n\}$) and $y_{i_j} \in \mathbb{R}_{\geq 0}$):

- Check if *for every $j \in \{1, \ldots, t\}$* we have that $x_{i_j} \geq y_{i_j}$; if not then reply with a message $(\mathsf{nofunds}, sid)$ and stop.
- Otherwise for $j \in \{1, \ldots, t\}$ let $x_{i_j} := x_{i_j} - y_{i_j}$. We say that $y_{i_1}, \ldots, y_{i_t}$ *coins were removed from the accounts of $P_{i_1}, \ldots, P_{i_t}$ (resp.) in $\widehat{\mathcal{L}}$.*

Fig. 4: The ledger functionality $\widehat{\mathcal{L}}$.

The GUC-security definition. Let π be a protocol working in the $\mathcal{G}$-hybrid model with access to the global ledger $\widehat{\mathcal{L}}(\Delta)$. The output of an environment $\mathcal{Z}$ interacting with a protocol π and an adversary Adv on input 1^λ and auxiliary input z is denoted as $\text{EXEC}^{\widehat{\mathcal{L}}(\Delta),\mathcal{G}}_{\pi,\mathsf{Adv},\mathcal{Z}}(\lambda, z)$. If π is a trivial protocol in which the parties simply forward their inputs to an ideal functionality $\mathcal{F}$, then we call the parties *dummy parties*, the adversary a *simulator* Sim, and we denote the above output as $\text{IDEAL}^{\widehat{\mathcal{L}}(\Delta)}_{\mathcal{F},\mathsf{Sim},\mathcal{Z}}(\lambda, z)$.

To simplify the description of our protocols and the ideal functionalities, we consider a class of restricted environments which we denote $\mathcal{E}_{res}$. These restrictions typically disallow the environment to carry out certain actions, e.g., we forbid $\mathcal{Z}$ to instruct one party to start a protocol without instructing the other party to start the protocol as well.[13] We emphasize that these restrictions can easily be eliminated by integrating additional checks into the protocols and functionalities. The restrictions defining $\mathcal{E}_{res}$ are informally introduced in Sec. 4 and their complete list can be found in Appx. B. We are now ready to state our main security definition.

Definition 1. *Let $\mathcal{E}$ be some set of restricted environments. We say that a protocol π working in a $\mathcal{G}$-hybrid model* emulates an ideal functionality $\mathcal{F}$ with respect to a global ledger $\widehat{\mathcal{L}}(\Delta)$ against environments from class $\mathcal{E}$ *if for every adversary* Adv *there exists a simulator* Sim *such that for every environment $\mathcal{Z} \in \mathcal{E}$ we have*

$$\left\{\text{EXEC}^{\widehat{\mathcal{L}}(\Delta),\mathcal{G}}_{\pi,\mathsf{Adv},\mathcal{Z}}(\lambda, z)\right\}_{\substack{\lambda\in\mathbb{N},\\ z\in\{0,1\}^*}} \overset{c}{\approx} \left\{\text{IDEAL}^{\widehat{\mathcal{L}}(\Delta)}_{\mathcal{F},\mathsf{Sim},\mathcal{Z}}(\lambda, z)\right\}_{\substack{\lambda\in\mathbb{N},\\ z\in\{0,1\}^*}}$$

(where "$\approx^c$" denotes computational indistinguishability of distribution ensembles, see, e.g., [12]).

Informally, the above definition says that any attack that can be carried out against the real-world protocol π can also be carried out against the ideal functionality $\mathcal{F}$. Since the ideal functionality is secure by design (see Sec. 4.2 and Appx. C), also the protocol offers the same level of security. In Sec. 5 we will discuss in more detail the roles of $\mathcal{F}$ and $\mathcal{G}$.

[13] For readers familiar with UC, we notice that general UC composition of course requires arbitrary environments. In Appx. F we prove that for our particular set of restrictions composition of our sub-protocols is preserved. An alternative approach would be to use a wrapper. However, due to the complexity of our protocol the description, of the wrapper would be highly convoluted.

Simplifying assumptions To simplify exposition, we omit the session identifiers *sid* and the sub-session identifiers *ssid*. Instead, we will use expressions like “message m is a reply to message m'”. We believe that this approach improves readability. Another simplifying assumption we make is that before the protocol starts the following public-key infrastructure is setup by some trusted party: (1) For every $i = 1, \ldots, n$ let $(pk_{P_i}, sk_{P_i}) \leftarrow_{\$} \mathsf{KGen}(1^{\lambda})$, (2) For every $i = 1, \ldots, n$ send the message $(sk_{P_i}, (pk_{P_1}, \ldots, pk_{P_n}))$ to P_i. We emphasize that the use of a PKI is only an abstraction, and can easily be realized using the blockchain.

4 State channels ideal functionality

In this section, we describe the ideal functionality that defines how ledger state channels and virtual state channels are created, maintained and closed. Before we do so, let us establish several conventions which simplify the description of the ideal functionality.

4.1 Abbreviated notation

When it is clear from the context which state channel γ we are talking about, we will denote the parties of γ as $A := \gamma.\mathsf{Alice}$, $B := \gamma.\mathsf{Bob}$ and $I := \gamma.\mathsf{Ingrid}$. We also introduce symbolic notation for sending and receiving messages. Instead of the instruction “Send the message *msg* to party P in round τ”, we write *msg* $\overset{\tau}{\hookrightarrow} P$. Instead of the instruction “Send the message *msg* to all parties in the set $\gamma.\mathsf{end\text{–}users}$ in round τ”, we write *msg* $\overset{\tau}{\hookrightarrow} \gamma.\mathsf{end\text{–}users}$. By *msg* $\overset{\tau}{\hookleftarrow} P$ we mean that an entity (i.e. the ideal functionality) receives a message *msg* from party P in round τ, and we use *msg* $\overset{\tau \leq \tau_1}{\hookleftarrow} P$ when an entity receives *msg* from party P latest in round τ_1. In the description of the ideal functionality we use two “timing functions”: TimeExe Req(i) that represents the maximal number of rounds it takes to inform a party that execution of a contract instance in a state channel of length $i > 0$ was requested by the other party, and TimeExe(i) that represents the maximal number of rounds it takes to execute of a contract instance in a state channel of length $i > 0$. Both of these functions are of the order $O(\Delta \cdot i)$ (see Sec. 7.3 for formal definition of these function and their relationship).

Each entity stores and maintains a set of all state channels it is aware of. Following [11] this set will be called *channel space* and denoted Γ. Sometimes we will abuse notation and interpret the channel space as a function which on input $id \in \{0,1\}^*$ returns a state channel with identifier id if such state channel exist and otherwise $\perp$. Every time a new contract instance in some of the state channels stored in Γ is successfully created (or an existing one is executed), the channels space Γ must be updated accordingly. To this end we define an auxiliary procedure `UpdateChanSpace`. The procedure takes as input a channel space Γ, a channel identifier id, a contract instance identifier cid, a new contract instance ν and two values add_A and add_B representing the required change in the cash values of the state channel with identifier id. The procedure sets $\Gamma(id).\mathsf{cspace}(cid) := \nu$, adds add_A coins to $\Gamma(id).\mathsf{cash}(A)$ and adds add_B coins to $\Gamma(id).\mathsf{cash}(B)$. Finally, it outputs the updated channel space Γ. In Fig. 5 we define the procedure formally.

```
UpdateChanSpace($\Gamma, id, cid, \tilde{\sigma}, \mathtt{C}, add_A, add_B$)

Let $\gamma := \Gamma(id)$ and $\sigma := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$. Make the following updates:
  1. Add $add_A$ coins to $\gamma.\mathsf{cash}(\gamma.\mathsf{Alice})$
  2. Add $add_B$ coins to $\gamma.\mathsf{cash}(\gamma.\mathsf{Bob})$
  3. Set $\gamma.\mathsf{cspace}(cid)$ equal to the tuple $(\tilde{\sigma}, \mathtt{C})$.
Output $\Gamma$ with the updated contract instance $cid$ in the state channel $\gamma$.
```

Fig. 5: Auxiliary procedure for updating the channel space.

4.2 The ideal functionality

We denote the state channel ideal functionality by $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$, where $i \in \mathbb{N}$ is the maximal length of a state channel that can be opened via the functionality, and $\mathcal{C}$ denotes the set of contract codes whose instances can be created in the state channels. The ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ communicates with parties from the set $\mathcal{P}$, and has access to the global ideal functionality $\widehat{\mathcal{L}}$ (the ledger). $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ maintains a channel space Γ containing all the open state channels. The set Γ is initially empty.

Since inputs of parties and the messages they send to the ideal functionality do not contain any private information, we implicitly assume that the ideal functionality forwards all messages it receives to the simulator Sim. More precisely, upon receiving the message m from party P the ideal functionality sends the message (P, m) to the simulator. The task of the simulator is to instruct the ideal functionality to make changes on the ledger and to output messages to the parties in the correct round (both depends on the choice made by the adversary Adv in the real world). In the description of the ideal functionality, we do not explicitly mention these instructions of Sim, but instead use the following abbreviation. By saying "wait for at most Δ rounds to remove/add x coins from P's account on the ledger" we mean that the ideal functionality waits until it is instructed by the simulator, which will happen within at most Δ rounds, and then request changes of P's account on the ledger. Let us emphasize this abbreviated notation does not affect the reactive nature of the ideal functionality (meaning that every action of the functionality has to be triggered by some other entity).

We present the formal definition of the $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ functionality in Fig. 6 (the constants that appear in the formal description follow from the technical details of our protocols, see Sec. 6.3 and Sec. 7.2. Here we provide some intuitions behind this definition, informally introduce the restrictions on the environment (see Sec. 3.3) whose full list is given in Appx. B, and argue why the ideal functionality satisfies all the security and efficiency properties stated in Sec. 3.2.

State channel creation The $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ functionality consists of two "state channel creation" procedures: one for ledger and one for virtual state channels. The ledger state channel creation procedure starts with a "create" message from A (without loss of generality we assume that A always initiates the creation process). The functionality removes the coins that A wants to deposit in the ledger state channel from A's account on the ledger, and waits for B to declare that he wants to create the ledger state channel as well. If this happens within Δ rounds, then B's coins are removed from the ledger and the ledger state channel is created which is communicated to the parties with the "created" message. Otherwise A can get her money back by sending a "refund" message. Since both parties have to send the message "create", the *consensus on creation* security property is clearly satisfied for ledger state channels.

The creation procedure for a virtual state channel γ works slightly differently since its effects are visible on the subchannels of γ. The intention to create γ is expressed by $P \in \gamma.\mathsf{end\text{–}users} \cup \{I\}$ by sending a "create" message to the functionality. Once such a message is received from P, the coins that are needed to create γ are locked immediately in the corresponding subchannel of γ (if $P = I$, then coins are locked in both subchannels of γ). If the functionality receives the "create" messages from all three parties within three rounds, then the virtual state channel is created, which is communicated to $\gamma.\mathsf{end\text{–}users}$ by the "created" message.[14] Thus, the *consensus on creation* security property is satisfied also for virtual state channels and since the successful creation takes three rounds, the *constant round virtual state channel creation* holds as well.

After the virtual state channel is created, $\gamma.\mathsf{end\text{–}users}$ can use it until round $\gamma.\mathsf{validity}$. When this round comes, the parties initiate the closing procedure. The functionality then distributes the coins of γ back to its subchannels according to the balance in γ's last version. In case there exists cid such that $\gamma.\mathsf{cspace}(cid)$ is a contract instance with locked coins, then all of these coins go back to I in *both* subchannels of γ. This

[14] Note that the intermediary I is not informed whether the virtual channel has been created. This choice is made to keep the protocol as simple as possible. Note also that I does not need this information, as she is not allowed to update this virtual channel.

Functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$

This functionality accepts messages from parties in $\mathcal{P}$. We use the abbreviated notation defined in Sec. 4.1.

Ledger state channel creation

Upon $(\text{create}, \gamma) \xhookleftarrow{\tau_0} A$ where γ is a ledger state channel:

1. Within Δ rounds remove $\gamma.\mathsf{cash}(A)$ coins from A's account on $\widehat{\mathcal{L}}$.
2. If $(\text{create}, \gamma) \xhookleftarrow{\tau_1 \leq \tau_0 + \Delta} B$, remove within 2Δ rounds $\gamma.\mathsf{cash}(B)$ coins from B's account on $\widehat{\mathcal{L}}$ and then set $\Gamma(\gamma.\mathsf{id}) := \gamma$, send $(\text{created}, \gamma) \hookrightarrow \gamma.\mathsf{end\text{–}users}$ and stop.
3. Otherwise upon $(\text{refund}, \gamma) \xhookleftarrow{> \tau_0 + 2\Delta} A$, within Δ rounds add $\gamma.\mathsf{cash}(A)$ coins to A's account on $\widehat{\mathcal{L}}$.

Virtual state channel creation

1. Upon $(\text{create}, \gamma) \hookleftarrow P$, where $P \in \gamma.\mathsf{end\text{–}users} \cup \{I\}$, record the message and proceed as follows:
   - If $P \in \gamma.\mathsf{end\text{–}users}$ proceed as follows: If you have not yet received (create, γ) from I, then remove $\gamma.\mathsf{cash}(P)$ coins from P's balance in $\gamma.\mathsf{subchan}(P)$ and $\gamma.\mathsf{cash}(\gamma.\mathsf{other\text{-}party}(P))$ coins from I's balance in $\gamma.\mathsf{subchan}(P)$.
   - If $P = I$, then for both $P \in \gamma.\mathsf{end\text{–}users}$ proceed as follows: If you have not yet received (create, γ) from P then remove $\gamma.\mathsf{cash}(P)$ coins from P's balance in $\gamma.\mathsf{subchan}(P)$, and $\gamma.\mathsf{cash}(\gamma.\mathsf{other\text{-}party}(P))$ coins from I's balance in $\gamma.\mathsf{subchan}(P)$.
2. If within 3 rounds you record (create, γ) from all users in $\gamma.\mathsf{end\text{–}users} \cup \{\gamma.\mathsf{Ingrid}\}$, then define $\Gamma(\gamma.\mathsf{id}) := \gamma$, send $(\text{created}, \gamma) \hookrightarrow \gamma.\mathsf{end\text{–}users}$ and wait for channel closing in Step 4 (in the meanwhile accepting the update and execute messages concerning γ).
3. Otherwise wait until round $\gamma.\mathsf{validity}$. Then within $2 \cdot (\text{TimeExeReq}(\lceil j/2 \rceil) + \text{TimeExe}(\lceil j/2 \rceil))$ rounds, where $j := \gamma.\mathsf{length}$, refund the coins that you removed from the subchannels in Step 1.

Automatic closure of virtual state channel γ when round $\gamma.\mathsf{validity}$ comes:

4. Let $j := \gamma.\mathsf{length}$. Within $2 \cdot (\text{TimeExeReq}(\lceil j/2 \rceil) + \text{TimeExe}(\lceil j/2 \rceil))$ rounds proceed as follows. Let $\hat{\gamma}$ be the current version of the virtual state channel, i.e. $\hat{\gamma} := \Gamma(\gamma.\mathsf{id})$, and let $\hat{c}_A := \hat{\gamma}.\mathsf{cash}(A)$ and $\hat{c}_B := \hat{\gamma}.\mathsf{cash}(B)$.
5. Add $\hat{c}_A$ coins to A's balance and $\hat{c}_B$ coins to I's balance in $\gamma.\mathsf{subchan}(A)$. Add $\hat{c}_A$ coins to I's balance and $\hat{c}_B$ coins to B's balance in $\gamma.\mathsf{subchan}(B)$. If there exists $cid \in \{0,1\}^*$ such that $\sigma_{cid} := \hat{\gamma}.\mathsf{cspace}(cid).\mathsf{storage} \neq \bot$ and $\hat{c} := \sigma_{cid}.\mathsf{locked} > 0$, then add $\hat{c}$ coins to I's balance in both $\gamma.\mathsf{subchan}(A)$ and $\gamma.\mathsf{subchan}(B)$. Erase $\hat{\gamma}$ from Γ and $(\text{closed}, \gamma.\mathsf{id}) \hookrightarrow \gamma.\mathsf{end\text{–}users}$.

Contract instance update

Upon $(\text{update}, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookleftarrow{\tau_0} P$, let $\gamma := \Gamma(id)$, $j = \gamma.\mathsf{length}$. If $P \notin \gamma.\mathsf{end\text{–}users}$ then stop. Else proceed as follows:

1. Send $(\text{update–requested}, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookrightarrow{\tau_0 + 1} \gamma.\mathsf{other\text{–}party}(P)$ and set $T := \tau_0 + 1$ in optimistic case when both parties in $\gamma.\mathsf{end\text{–}users}$ are honest. Else if $j = 1$, set $T := \tau_0 + 3\Delta + 1$ and if $j > 1$, set $T := \tau_0 + 4 \cdot \text{TimeExeReq}(\lceil j/2 \rceil) + 1$.
2. If $(\text{update–reply}, ok, id, cid) \xhookleftarrow{\tau_1 \leq T} \gamma.\mathsf{other\text{–}party}(P)$, then set $\Gamma := \mathtt{UpdateChanSpace}(\Gamma, id, cid, \tilde{\sigma}, \mathtt{C}, add_A, add_B)$, where $add_A := -\tilde{\sigma}.\mathsf{cash}(A)$ if $\gamma.\mathsf{cspace}(cid) = \bot$ and $add_A := \sigma.\mathsf{cash}(A) - \tilde{\sigma}.\mathsf{cash}(A)$ otherwise for $\sigma := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$. The value add_B is defined analogously. Then send $(\text{updated}, id, cid) \xhookrightarrow{\tau_1 + 1} \gamma.\mathsf{end\text{–}users}$ and stop.

Contract instance execution

Upon $(\text{execute}, id, cid, f, z) \xhookleftarrow{\tau_0} P$, let $\gamma := \Gamma(id)$ and $j = \gamma.\mathsf{length}$. If $P \notin \gamma.\mathsf{end\text{–}users}$ then stop. Else set T_1 and T_2 as:

- In the optimistic case when both parties in $\gamma.\mathsf{end\text{–}users}$ are honest, set $T_1 := \tau_0 + 4$ and $T_2 := \tau_0 + 5$.
- In the pessimistic case when at least one party in $\gamma.\mathsf{end\text{–}users}$ is corrupt, set $T_1, T_2 := \tau_0 + 4\Delta + 5$ if $j = 1$ and set $T_1 := \tau_0 + 2 \cdot \text{TimeExeReq}(\lceil j/2 \rceil) + 5$, $T_2 := \tau_0 + 4 \cdot \text{TimeExeReq}(\lceil j/2 \rceil) + 5$ if $j > 1$.

1. In round $\tau_1 \leq T_1$, send $(\text{execute–requested}, id, cid, f, z) \xhookrightarrow{\tau_1} \gamma.\mathsf{other\text{–}party}(P)$.
2. In round $\tau_2 \leq T_2$, let $\gamma := \Gamma(id)$, $\nu := \gamma.\mathsf{cspace}(cid)$, $\sigma := \nu.\mathsf{storage}$, and $\tau := \tau_0$ if P is honest and else τ is set by the simulator. Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma, P, \tau, z)$. If $m = \bot$, then stop. Else set $\Gamma := \mathtt{UpdateChanSpace}(\Gamma, id, cid, \tilde{\sigma}, \nu.\mathsf{code}, add_L, add_R)$ and send $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_3} \gamma.\mathsf{end\text{–}users}$.

Ledger state channel closure

Upon $(\text{close}, id) \xhookleftarrow{\tau_0} P$, let $\gamma = \Gamma(id)$. If $P \notin \gamma.\mathsf{end\text{–}users}$ then stop. Else wait at most 7Δ rounds and distinguish the following two cases:

1. If there exists $cid \in \{0,1\}^*$ such that $\sigma_{cid} := \gamma.\mathsf{cspace}(cid).\mathsf{storage} \neq \bot$ and $\sigma_{cid}.\mathsf{locked} \neq 0$, then stop.
2. Otherwise wait up to Δ rounds to add $\gamma.\mathsf{cash}(A)$ coins to A's account and $\gamma.\mathsf{cash}(B)$ coins to B's account on the ledger $\mathcal{L}$. Then set $\Gamma(id) := \bot$, send $(\text{closed}, id) \xhookrightarrow{\tau_2 \leq \tau_0 + 8\Delta} \gamma.\mathsf{end\text{–}users}$ and stop.

Fig. 6: The state channel ideal functionality.

is to guarantee that I never loses coins even if end-users of γ do not terminate their contract instance in γ before $\gamma.\mathsf{validity}$. See Appx. C for a formal proof of the *balance security* property.

In both cases ("ledger" and "virtual") we assume that all the honest parties involved in channel creation initiate the procedure in the same round and that they have enough funds for the new state channel. In case of a virtual state channel, we additionally assume that the length of its two subchannels differ at most by one.[15]

Contract instance update The procedure for updating a contract instance is identical for ledger and virtual state channels (this procedure is also used for creating new contract instances). It is initiated by a party $P \in \gamma.\mathsf{end\text{--}users}$ that sends an "update" message to the ideal functionality. This message has parameters *id* and *cid* that identify a state channel γ and a contract instance in this state channel (respectively). The other parameters, $\tilde{\sigma}$ and C, denote the new storage and code of the contract instance. The party $Q := \gamma.\mathsf{other\text{--}party}(P)$ is asked to confirm the update via an "update-requested" message. If Q replies with an "update-reply" message within 1 round if both parties are honest and within T rounds otherwise (where T is a function of state channel length, see Step 2), the contract instance with identifier *cid* in γ gets replaced with a contract instance determined by the tuple $(\tilde{\sigma}, \mathsf{C})$. In the next round, both parties in $\gamma.\mathsf{end\text{--}users}$ get notified via an "updated" message. Note that Q always has to confirm the update which implies the *consensus on update* security property. The *constant round optimistic update* efficiency property holds as well since the update takes exactly 2 rounds if both parties are honest.

We assume that the environment never asks the parties to do obviously illegal things, like updating a contract instance in a state channel that does not exits, or creating a contract instance when there are not enough coins in the subchannels. Moreover, we assume that the environment never asks to update a contract instance when it is already being updated or executed.[16]

Contract instance execution The procedure for executing a contract instance is initiated by one of the parties $P \in \gamma.\mathsf{end\text{--}users}$ that sends an "execute" message to the ideal functionality in round τ_0. This message has parameters *id* and *cid* whose meaning is as in the update procedure. Other parameters are: f denoting the contract function to be executed, and z which is an additional input parameter to the function f. The execution results in updating the contract instance with identifier *cid* according to the result of computing $f(\sigma, P, \tau, z)$, where σ is the current storage of the contract instance and $\tau := \tau_0$ in case P is honest and determined by the simulator otherwise. The other party of the state channel is notified about the execution request before round $\tau_0 + 5$ in the optimistic case and before round $\tau_0 + T_1$ otherwise. Both parties from the set $\gamma.\mathsf{end\text{--}users}$ learn the result of the execution before round $\tau_0 + 5$ in the optimistic case (which implies the *constant round optimistic execute*) and before round $\tau_0 + T_2$ otherwise. The values T_1 and T_2 are functions of state channel length, see the formal description in Fig. 6. Observe that contract instance execution initiated by party P does *not* require approval of the other party of the channel (although the other party is informed about the execution request). This implies that the *guarantee of execution* security property is satisfied.

We would like to emphasize that if two different execute messages are received by the ideal functionality at the same time (or not too many rounds from each other), then it is up to the adversary to decide which function is executed first.[17] Designers of contract codes and users of the protocols should be aware of this possible asynchronicity.

Ledger state channel closure The procedure for closing a ledger state channel γ starts when a party $P \in \gamma.\mathsf{end\text{--}users}$ sends to the ideal functionality a message (close, id), where id is the identifier of ledger state channel γ to be closed. The functionality checks (in Step 1) if there are no contract instances that are open over γ. If not, then in Step 2 the functionality distributes the coins from γ to the ledger accounts of the parties

[15] As discussed in Sec. 2, we make this assumption to achieve pessimistic time complexity which is linear in the state channel length.

[16] Although we forbid parallel updates of the *same* contract instance, we do not make any restrictions about parallel updates of two different contract instances even if they are in the same ledger state channel. This in particular means that we allow concurrent creation of virtual state channels.

[17] Note that this is the case also for execution of smart contracts on the blockchain.

according to γ's latest balance, and notifies the parties about a successful closure. The restrictions on the environment in case of the contract instance execution and ledger state channel closure are straightforward (see Appx. B).

4.3 Using the state channel ideal functionality

Let us now demonstrate how to use our ideal functionality for generalized state channel networks in practice. We do it on a concrete example of the two party lottery (already discussed in Sec. 2). The first step is to define a contract code $\mathtt{C}_{\mathtt{lot}}(i)$ which allows two parties to play the lottery in a state channel of length at most i. A contract storage σ of $\mathtt{C}_{\mathtt{lot}}(i)$ has, in addition to the mandatory attributes $\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R, \sigma.\mathsf{cash}$ and $\sigma.\mathsf{locked}$ (see Sec. 4.1), the attribute $\sigma.\mathsf{start} \in \mathbb{N}$, whose purpose it to store the construction round, the attribute $\sigma.\mathsf{com} \in \{0,1\}^*$, to store the commit value when submitted by $\sigma.\mathsf{user}_L$, and the attribute $\sigma.\mathsf{bit} \in \{0,1\}$ to store the secret bit when provided by $\sigma.\mathsf{user}_R$.

The contract code has one constructor $\mathtt{Init}_{\mathtt{lot}}$ which generates the initial contract storage σ such that both $\sigma.\mathsf{cash}(\sigma.\mathsf{user}_L)$ and $\sigma.\mathsf{cash}(\sigma.\mathsf{user}_R)$ are equal to 1 (each user deposits 1 coin). The contract functions are: (i) $\mathtt{Com}$ which, if executed by $\sigma.\mathsf{user}_L$ on input c, stores c in $\sigma.\mathsf{com}$, (ii) $\mathtt{Reveal}$ which, if executed by $\sigma.\mathsf{user}_R$ on input r_B, stores r_B in $\sigma.\mathsf{bit}$, (iii) $\mathtt{Open}$ which allows $\sigma.\mathsf{user}_L$ to open the commitment stored in $\sigma.\mathsf{com}$ and pays out 2 coins to the winner, and (iv) $\mathtt{Punish}$ which allows a party to unlock coins from the contract instance in case the other party misbehaves. See Appx. D for a formal definition of $\mathtt{C}_{\mathtt{lot}}(i)$.

Assume now that parties Alice and Bob have a virtual state channel γ created via the ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$, where $\mathtt{C}_{\mathtt{lot}}(i) \in \mathcal{C}$. If Alice wants to play the lottery using γ, she first locally executes the constructor $\mathtt{Init}_{\mathtt{lot}}$ to obtain the initial contract storage σ. Then she sends the message $(\text{update}, \gamma.\mathsf{id}, \mathit{cid}, \sigma, \mathtt{C}_{\mathtt{lot}}(i))$ to $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ for some contact instance identifier cid never used before. The ideal functionality informs Bob about Alice's intention to play by sending the message $(\text{update–requested}, \gamma.\mathsf{id}, \mathit{cid}, \sigma, \mathtt{C}_{\mathtt{lot}}(i))$ to him. If Bob agrees with playing the game, he sends the reply $(\text{update–reply}, \mathit{ok}, \gamma.\mathsf{id}, \mathit{cid})$. Alice and Bob can now start playing in a way we describe below (let τ_0 be the current round).

1. **Commit:** In round τ_0 Alice locally chooses a random bit $r_A \in \{0,1\}$ and a random string $s \in \{0,1\}^\lambda$, where λ is the security parameter, locally computes the commit value c using the randomness s. Then she submits c by sending the message $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Com}, c)$ to the ideal functionality.
2. **Reveal:** If before round $\tau_0 + \text{TimeExe}(i)$ Bob receives a message from the ideal functionality that Alice committed to her secret bit, Bob locally chooses a random bit $r_B \in \{0,1\}$ which he submits by sending the message $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Reveal}, r_B)$ to the ideal functionality. Otherwise, in round $\tau_0 + \text{Time}$ $\text{Exe}(i)$, he sends the message $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Punish}, \perp)$ to the ideal functionality to unlock all coins from the lottery contract by which he punishes Alice for her misbehavior.
3. **Open:** If before round $\tau_0 + 2 \cdot \text{TimeExe}(i)$ Alice receives a message from the ideal functionality that Bob reveled his secret bit r_B, she opens her commitment by sending $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Open}, (r_A, s))$. Otherwise, in round $\tau_0 + 2 \cdot \text{TimeExe}(i)$, she sends the message $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Punish}, \perp)$ to unlock all coins from the lottery contract by which she punishes Bob for his misbehavior.
4. **Finalize:** If until round $\tau_0 + 3 \cdot \text{TimeExe}(i)$ Bob did not receive a message from the ideal functionality that Alice opened her commitment, Bob sends the message $(\text{execute}, \gamma.\mathsf{id}, \mathit{cid}, \mathtt{Punish}, \perp)$ to the ideal functionality to unlock all coins from the lottery contract and finalize the game.

5 An overview of our approach

In this section we provide a high level idea of the modular design of our protocol realizing the state channel ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ (the main ideas behind our construction were already presented in Sec. 2).

Ledger state channels Our first step is to define an ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ which models the behavior of a concrete smart contract, which we call state channel contract. This contract allows two parties to open, maintain and close a ledger state channel on the blockchain. The ideal functionality is parametrized by the set of contract codes $\mathcal{C}$ whose instances can be opened in the ledger state channels created via this ideal functionality. The ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ together with the ledger functionality $\widehat{\mathcal{L}}$ can be implemented by a cryptocurrency which supports such state channel contracts on its blockchain (a candidate cryptocurrency would be, e.g., Ethereum). We use this contract ideal functionality to design a protocol $\Pi(1, \mathcal{C})$ which realizes the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ (i.e. the protocol for ledger state channels). We describe the protocol as well as the ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ in Sec. 6.3. Furthermore, in Appx. E we prove that the protocol $\Pi(1, \mathcal{C})$ indeed emulates the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ in the $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ hybrid world. This statement is formalized by the following theorem.

Theorem 1. *Suppose the underlying signature scheme is existentially unforgeable against chosen message attacks. The protocol $\Pi(1, \mathcal{C})$ working in $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$-hybrid model emulates the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ against environments from class $\mathcal{E}_{res}$ for every set of contract codes $\mathcal{C}$ and every $\Delta \in \mathbb{N}$.*

Virtual state channels As already mentioned in Sec. 2, our technique allows to create virtual state channels of arbitrary length, via using the state channel functionality recursively. By this we mean that a protocol for constructing state channels of length up to i will work in a model with access to an ideal functionality for constructing state channels of length up to $i - 1$. More formally, for every $i > 1$ we construct (in Sec. 7) a protocol $\Pi(i, \mathcal{C})$ realizing the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ in the $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i - 1, \mathcal{C}')$-hybrid world. Here $\mathcal{C}'$ is a set of contract codes defined as $\mathcal{C}' := \mathcal{C} \cup \mathtt{VSCC}_i(\mathcal{C})$, where $\mathtt{VSCC}_i(\mathcal{C})$ is a contract code (also presented in Sec. 7), which we call the virtual state channel contract, that allows to create a virtual state channel of length i in which contract instance with code from the set $\mathcal{C}$ can be opened. Thus importantly, the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i - 1, \mathcal{C}')$ allows to create state channels that can serve as subchannels of a virtual channel of length i.

Very briefly, the hybrid ideal functionality is used by parties of the protocol $\Pi(i, \mathcal{C})$ as follows. If a party receives a message regarding a state channel of length $j < i$, then it simply forwards this message to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i - 1, \mathcal{C}')$. The more interesting case is when a party receives a message regarding a virtual state channel γ of length exactly i. Then it uses the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i - 1, \mathcal{C}')$ to make changes in the subchannels of the virtual state channels γ. In Appx. F we prove the following theorem.

Theorem 2. *Suppose the underlying signature scheme is existentially unforgeable against chosen message attacks. The protocol $\Pi(i, \mathcal{C})$ working in $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i - 1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$-hybrid model emulates the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ against environments from class $\mathcal{E}_{res}$ for every set of contract codes $\mathcal{C}$, every $i > 1$ and every $\Delta \in \mathbb{N}$.*

By applying the composition recursively, we get a construction of a protocol realizing $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ in the $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\widehat{\mathcal{C}})$-hybrid model, where $\widehat{\mathcal{C}}$ is a result of applying the "$\mathcal{C} := \mathcal{C} \cup \mathtt{VSCC}_i(\mathcal{C})$" equation i times recursively. See Fig. 7 for an example for $i = 3$.

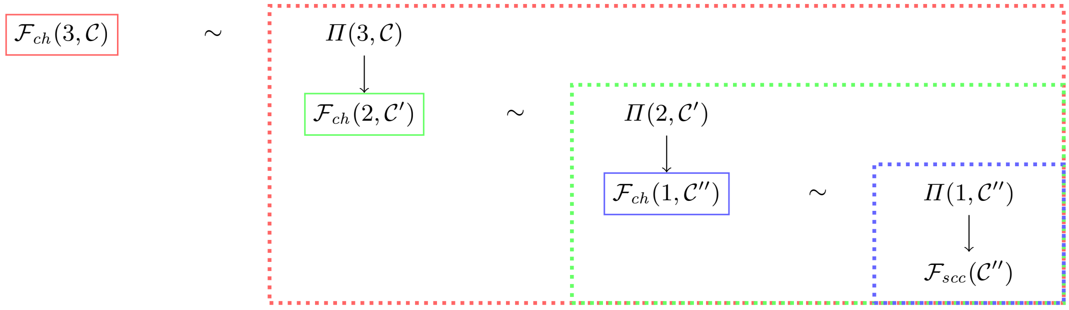
$\mathcal{F}_{ch}(3, \mathcal{C})$
$\sim$
$\Pi(3, \mathcal{C})$
$\mathcal{F}_{ch}(2, \mathcal{C}')$
$\sim$
$\Pi(2, \mathcal{C}')$
$\mathcal{F}_{ch}(1, \mathcal{C}'')$
$\sim$
$\Pi(1, \mathcal{C}'')$
$\mathcal{F}_{scc}(\mathcal{C}'')$

Fig. 7: Our modular approach. Above $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}$, $\mathcal{F}_{scc} := \mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}$, $\mathcal{C}' := \mathcal{C} \cup \mathtt{VSCC}_3(\mathcal{C})$ and $\mathcal{C}'' := \mathcal{C}' \cup \mathtt{VSCC}_2(\mathcal{C}')$.

6 Ledger State Channels

In this section, we define an ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ which represents the smart contract allowing two parties to open, maintain and close a ledger state channel. We call such a smart contract a state channel contract. Then we describe the protocol $\Pi(1, \mathcal{C})$ that realizes the state channels ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ (see Sec. 4) in the hybrid world $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ for any set of contract codes $\mathcal{C}$. But before we do so, we introduce several additional terms, auxiliary procedures and notation which will be used the protocol descriptions.

6.1 Maintaining a local channel space

In order to update the contract instances off-chain, the users of the state channel γ will store some additional information in their local copies of γ. To this end, we introduce the following terminology. A *contract instance version* is an attribute tuple ν that in addition to the attributes of contract instance (which are $\nu.\mathsf{code}, \nu.\mathsf{storage}$) has an attribute $\nu.\mathsf{version} \in \mathbb{N}$. As the name suggests, the purpose of $\nu.\mathsf{version}$ is to indicate the version of the contract instance. When a contract instance is created, each user locally sets the version number of the contract instance to 0. Each time users want to update the contract instance off-chain, they increase the value of the version attribute by one. A *signed contract instance version* additionally contains an attribute $\nu.\mathsf{sign}\colon \{\nu.\mathsf{user}_L, \nu.\mathsf{user}_R\} \to \{0,1\}^*$, where $\nu.\mathsf{sign}(\nu.\mathsf{user}_L)$ is a signature of $\nu.\mathsf{user}_L$ on the contract instance version and $\nu.\mathsf{sign}(\nu.\mathsf{user}_R)$ is a signature of $\nu.\mathsf{user}_R$ on the contract instance version. An attribute tuple γ is *state channel's extended version* if it is defined as the normal state channel, except that every $\gamma.\mathsf{cspace}(cid)$ is a signed contact instance version. To shorten the description of our protocols, we define an auxiliary function which on input a signed contract instance version checks both signatures and outputs 1 if both of them are valid and outputs 0 otherwise.

$\mathtt{VerifyInstance}(id, cid, \nu)$

Let $\sigma_n := \nu.\mathsf{storage}$, $L := \sigma_n.\mathsf{user}_L$, $R := \sigma_n.\mathsf{user}_R$, $\mathsf{C} := \nu.\mathsf{code}$, $w := \nu.\mathsf{version}$, $s_L := \nu.\mathsf{sign}(L)$, $s_R := \nu.\mathsf{sign}(R)$.
If $\mathtt{Vfy}_{pk_R}(id, cid, \sigma_n, \mathsf{C}, w; s_R) \neq 1$ or $\mathtt{Vfy}_{pk_L}(id, cid, \sigma_n, \mathsf{C}, w; s_L) \neq 1$, then return 0. Else return 1.

Fig. 8: Auxiliary function for validation of a signed contract instance version.

Recall from Sec. 4.1 that each entity (ideal functionality or party in a protocol), stores and maintains a set of all state channels it is aware of. This set is called *channel space* and denoted Γ. When we want to emphasize that we are referring to a local version of a state channel stored by some entity T, we add T to the superscript. So for instance, $\gamma^T := \Gamma^T(id)$ denotes T's local version of the state channel γ with identifier id as stored in T's channel space Γ^T. In Sec. 4.1, we described an auxiliary `UpdateChanSpace` whose purpose is to update a channels space stored by some entity T. See Fig. 5 on page 14 for the formal definition of the procedure. In our protocols, it will be often the case that the values of the input parameters add_L and add_R will correspond to the difference between the amount of coins locked in the contract instance before the update and the amount of coins in the new contract instance (recall that this was the case in the update part of the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$). To simplify the protocol description even further and avoid repetition of protocol code, we define another auxiliary procedure $\mathtt{UpdateChanSpace}^*$ which will derive the values add_L and add_R automatically from the new contract storage $\tilde{\sigma}$. See the formal description in Fig. 9.

We define both $\mathtt{UpdateChanSpace}^*$ and `UpdateChanSpace` in case a party wants to update the private extended version of the contract instance. Notice that in this case procedures will take additional two parameters: the new version number and the signatures created by the parties.

In addition to the channel space, each party P maintains a set Γ_{aux}^P containing additional information about the contract instances in the open state channels of the party. The tuple $aux := \Gamma_{aux}^P(id, cid)$ has the following attributes: $aux.\mathsf{next\text{-}version} \in \mathbb{N}$ denoting the version number to be used during the next update

$\texttt{UpdateChanSpace}^*(\varGamma, id, cid, \tilde{\sigma}, \mathtt{C})$

Let $\gamma := \varGamma(id)$ and $\sigma := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$. If $\sigma = \perp$, the set $(x_A, x_B) := (0, 0)$. Else set $(x_A, x_B) := (\sigma.\mathsf{cash}(\gamma.\mathsf{Alice}), \sigma.\mathsf{cash}(\gamma.\mathsf{Bob}))$. Make the following updates:

1. Add $x_A - \tilde{\sigma}.\mathsf{cash}(\gamma.\mathsf{Alice})$ coins to $\gamma.\mathsf{cash}(\gamma.\mathsf{Alice})$
2. Add $x_B - \tilde{\sigma}.\mathsf{cash}(\gamma.\mathsf{Bob})$ coins to $\gamma.\mathsf{cash}(\gamma.\mathsf{Bob})$
3. Set $\gamma.\mathsf{cspace}(cid)$ equal to the tuple $(\tilde{\sigma}, \mathtt{C})$.

Output $\varGamma$ with the updated contract instance cid in the state channel γ.

Fig. 9: Modification of the auxiliary procedure for updating the channel space.

of the contract instance (id, cid);[18]; $aux.\mathsf{corrupt} \in \{0, 1\}$ which is set to 1 the first time parties run into disagreement about the contract instance (id, cid); $aux.\mathsf{registered} \in \{0, 1\}$ which is set to 1 the if the contract instance (id, cid) is registered (on the blockchain in case of ledger state channel and in the subchannels in case of virtual state channel); and if $\varGamma^P(id)$ is a virtual state channel, then aux has an addition attribute $aux.\mathsf{toExecute}$ which is a set containing all functions that party P requested to "forcefully" execute via the subchannels in case of a virtual state channel.

For better readability of the protocol descriptions, we write "Mark (id, cid) as corrupt" instead of the instruction "Set $\varGamma_{aux}(id, cid).\mathsf{corrupt} := 1$". Similarly, we write "Mark (id, cid) as registered" instead of the instruction "Set $\varGamma_{aux}(id, cid).\mathsf{registered} := 1$".

6.2 Ideal functionality for the State Channel Contract

The ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ is parametrized by a set $\mathcal{C}$ defining the contract codes whose instance can be constructed in a ledger state channel. The ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ has access to the global ideal functionality $\widehat{\mathcal{L}}$ (the ledger). The ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ accepts messages from parties $\mathcal{P} := \{P_1, \ldots, P_n\}$. Let us emphasize that since the ideal functionality models a concrete smart contracts on the ledger, each communication session (party sends a message to the ideal functionality which potentially makes some modifications on the ledger and replies) comes with a delay up to Δ rounds. The exact timing (and if applicable, the exact round when transaction on the ledger takes place), is determined by the adversary. In order to shorten the description of the ideal functionality, we do not mention the transact instructions explicitly.

The functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ maintains a space $\varGamma$ containing all open ledger state channels. The set $\varGamma$ is initially empty. The functionality consists of four parts: "Create a ledger state channel", "Contract instance registration", "Contract instance execution" and "Close a ledger state channel". These parts will be described and formally defined together with the protocol for ledger state channels in Sec. 6.3.

6.3 Protocol for Ledger State Channels

Create a ledger state channel. In order to create a new ledger state channel γ, the environment sends the message (create, γ) to both parties in $\gamma.\mathsf{end\text{–}users}$. The protocol for creating a ledger state channel works at a high level as follows.

The initiating party $\gamma.\mathsf{Alice}$ requests construction of the state channel contract by sending the message (construct, γ) to the ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$. The ideal functionality locks the required amount of coins in her account on the ledger and sends the message (initializing, γ) to both parties. If party $\gamma.\mathsf{Bob}$ confirms the initialization by sending the message (confirm, γ), the ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ outputs (created, γ). In case $\gamma.\mathsf{Bob}$ does not confirm, the ledger state channel cannot be created and the initiating party $\gamma.\mathsf{Alice}$ has the option to refund the coins that were locked in her account on the ledger during the first step.

[18] For technical reasons (see Appx. 6.3) it is not always the case that $\varGamma(id).\mathsf{cspace}(cid).\mathsf{version}{+}1 = \varGamma_{aux}(id, cid).\mathsf{next\text{-}version}$.

Creation of a ledger state channels takes up to $2\varDelta$ rounds since it requires two interactions with the hybrid ideal functionality modeling a smart contract on the ledger. In case the ledger state channel is not created but $\gamma.\mathsf{Alice}$'s coins were locked in the first phase of the ledger state channel creation, she can receive them back latest after $3\varDelta$ rounds. Formal description of the protocol for ledger state channel creation and the corresponding part of the $\mathcal{F}_{scc}$ functionality can be found below.

Protocol $\varPi(1, \mathcal{C})$: Create a ledger state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, let $\mathcal{F}_{scc} := \mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\varDelta)}(\mathcal{C})$.

$\boxed{\text{Party } A \text{ upon } (\text{create}, \gamma) \overset{\tau_0}{\hookleftarrow} \mathcal{Z}}$

1. Send $(\text{construct}, \gamma) \overset{\tau_0}{\hookrightarrow} \mathcal{F}_{scc}$ and wait.

$\boxed{\text{Party } B \text{ upon } (\text{create}, \gamma) \overset{\tau_0}{\hookleftarrow} \mathcal{Z}}$

2. If $(\text{initializing}, \gamma) \xhookleftarrow{\tau_1 \leq \tau_0 + \varDelta} \mathcal{F}_{scc}$, then $(\text{confirm}, \gamma) \overset{\tau_1}{\hookrightarrow} \mathcal{F}_{scc}$ and wait. Else stop.
3. If $(\text{initialized}, \gamma) \xhookleftarrow{\tau_2 \leq \tau_0 + 2 \cdot \varDelta} \mathcal{F}_{scc}$, then set $\varGamma^B(\gamma.\mathsf{id}) := \gamma$, output $(\text{created}, \gamma) \overset{\tau_2}{\hookrightarrow} \mathcal{Z}$ and stop. Else stop.

$\boxed{\text{Back to party } A}$

4. If $(\text{initialized}, \gamma) \xhookleftarrow{\tau_2 \leq \tau_0 + 2 \cdot \varDelta} \mathcal{F}_{scc}$, then set $\varGamma^A(\gamma.\mathsf{id}) := \gamma$, output $(\text{created}, \gamma) \overset{\tau_2}{\hookrightarrow} \mathcal{Z}$ and stop. Else go to next step.
5. If $(\text{refund}, \gamma) \xhookleftarrow{\tau_3 > \tau_0 + 2 \cdot \varDelta} \mathcal{Z}$, then $(\text{refund}, \gamma) \overset{\tau_3}{\hookrightarrow} \mathcal{F}_{scc}$ and stop.

Functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\varDelta)}(\mathcal{C})$: Create a ledger state channel

We use the abbreviated notation from Sec. 4.1.

Upon $(\text{construct}, \gamma) \overset{\tau_0}{\hookleftarrow} P$**:**

1. If $P \neq A$, there already exists a state channel γ' such that $\gamma.\mathsf{id} = \gamma'.\mathsf{id}$, $\gamma.\mathsf{cspace}(cid) \neq \perp$ for some $cid \in \{0,1\}^*$ or $\gamma.\mathsf{cash}(A) < 0$ or $\gamma.\mathsf{cash}(B) < 0$, then stop.
2. Within $\varDelta$ rounds remove $\gamma.\mathsf{cash}(A)$ coins from A's account on the ledger $\widehat{\mathcal{L}}$. If it is impossible due to insufficient funds, then stop. Else $(\text{initializing}, \gamma) \hookrightarrow B$ and store the pair $tamp := (\tau_0, \gamma)$.

Upon $(\text{confirm}, \gamma) \overset{\tau_1}{\hookleftarrow} P$**:**

3. If there is no pair $tamp = (\tau_0, \gamma)$ in the storage, $(\tau_1 - \tau_0) > \varDelta$ or $P \neq B$, then stop.
4. Within $\varDelta$ rounds remove $\gamma.\mathsf{cash}(B)$ coins from B's account on the ledger $\widehat{\mathcal{L}}$. If it is impossible due to insufficient funds, then stop. Else set $\varGamma(\gamma.\mathsf{id}) := \gamma$ and delete $tamp$ from the memory. Thereafter send $(\text{initialized}, \gamma) \hookrightarrow \gamma.\mathsf{end\text{–}users}$.

Upon $(\text{refund}, \gamma) \overset{\tau_2}{\hookleftarrow} P$**:**

5. If there is no pair $tamp = (\tau_0, \gamma)$ in the storage, $(\tau_2 - \tau_0) \leq 2\varDelta$ or $P \neq A$, then stop.
6. Else within $\varDelta$ rounds add $\gamma.\mathsf{cash}(A)$ coins to A's account in $\widehat{\mathcal{L}}$ and delete $tamp$ from the storage.

Register a contract instance in a ledger state channel. As long as both end-users of a ledger state channel behave honestly, they can update, execute and close contract instances running in the ledger state channel off-chain; i.e. without communicating with the ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\varDelta)}(\mathcal{C})$. However, once the parties run into dispute (e.g., one party does not communicate, sends an invalid message, etc.), parties have to resolve their disagreement on the ledger. We call this process "registration of a contract instance", and will describe its basic functionality below.

The registration of a contract instance might be necessary either when the contract instance is being updated, executed or when a ledger state channel is being closed. To prevent repeating the same part of the protocol multiple times in each of the protocols, we state the registration process as a separate procedure $\texttt{Register}(P, id, cid)$ which can be called by parties running one of the sub-protocols mentioned above. The procedure takes as input party P which initiates the registration and the identifiers defining the contract instance to be registered, i.e. identifier of the ledger state channel id and the contract instance identifier cid.

At a high level, the initiating party (assume for now that it is $\gamma.\mathsf{Alice}$) sends her contract instance version to the ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ which first checks the validity of the received version, and then within Δ rounds the hybrid ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ informs both users that the contract instance is being registered. Party $\gamma.\mathsf{Bob}$ then reacts by sending his own version of the contract instance to $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$. The ideal functionality compares the two received versions, registers the one with higher version number and within Δ rounds informs both users which version was registered. In case $\gamma.\mathsf{Bob}$ did not send in his version, $\gamma.\mathsf{Alice}$ can finalize the registration by sending the message "finalize–register" to the ideal functionality.

In the optimistic case when $\gamma.\mathsf{Bob}$ submits a valid version of the contract instance, the registration procedure takes up to 2Δ rounds since it requires two interactions with the ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$. In the pessimistic case when $\gamma.\mathsf{Bob}$ does not react or submits an invalid version, the procedure takes up to 3Δ. Formal description of this procedure and the corresponding part of the $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ functionality follow.

Procedure $\texttt{Register}(P, id, cid)$

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, we denote $\mathcal{F}_{scc} := \mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$.

Party P:

1. Let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, and let τ_0 be the current round. Send (instance–register, id, cid, ν^P) $\overset{\tau_0}{\hookrightarrow} \mathcal{F}_{scc}$.

Party Q upon (instance–registering, id, cid) $\overset{\tau_1}{\hookleftarrow} \mathcal{F}_{scc}$

2. Let $\gamma^Q := \Gamma^Q(id)$ and $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$. Then send (instance–register, id, cid, ν^Q) $\overset{\tau_1}{\hookrightarrow} \mathcal{F}_{scc}$ and goto step 4.

Back to party P:

3. If not (instance–registered, $id, cid, \tilde{\nu}$) $\xhookleftarrow{\tau_2 \leq \tau_0 + 2\Delta} \mathcal{F}_{scc}$, then send (finalize–register, id, cid) $\xhookrightarrow{\tau_3 = \tau_1 + \Delta} \mathcal{F}_{scc}$.

End for both $T = A$ and $T = B$

4. Upon (instance–registered, $id, cid, \tilde{\nu}$) $\hookleftarrow \mathcal{F}_{scc}$, mark (id, cid) as registered and set $\Gamma^T := \texttt{UpdateChanSpace}^*(\Gamma^T, id, cid, \tilde{\nu})$.

Functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$: Contract instance registration

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1.

Upon (instance–register, id, cid, ν) $\overset{\tau_0}{\hookleftarrow} P$**,** let $\gamma := \Gamma(id)$ and do:

1. If $P \in \gamma.\mathsf{end\text{–}users}$ and $\nu = \bot$, then goto step 3.

2. If party $P \notin \gamma.\mathsf{end\text{–}users}$, $\mathtt{VerifyInstance}(id, cid, \nu) \neq 1$, $\nu.\mathsf{code} \notin \mathcal{C}$, $\gamma.\mathsf{cspace}(cid) \neq \bot$ or $\nu.\mathsf{storage} \notin \nu.\mathsf{code}.\Lambda$, then stop.
3. Let $Q := \gamma.\mathsf{other\text{–}party}(P)$ and consider the following four cases:
   - If your memory contains a tuple $(P, id, cid, \widehat{\nu}, \widehat{\tau}_0)$, then stop.
   - If your memory contains a tuple $(Q, id, cid, \widehat{\nu}, \widehat{\tau}_0)$ and $\nu = \bot$, then stop.
   - If your memory contains a tuple $(Q, id, cid, \widehat{\nu}, \widehat{\tau}_0)$ and $\nu \neq \bot$, then first compare the version number, i.e. if $\widehat{\nu} = \bot$ or $\nu.\mathsf{storage.version} > \widehat{\nu}.\mathsf{storage.version}$, then set $\tilde{\nu} := (\nu.\mathsf{storage}, \nu.\mathsf{code})$ and otherwise set $\tilde{\nu} := (\widehat{\nu}.\mathsf{storage}, \widehat{\nu}.\mathsf{code})$. Thereafter wait for at most Δ rounds to send (instance–registered, $id, cid, \tilde{\nu}$) $\xhookrightarrow{\tau_1 \leq \tau_0 + \Delta}$ $\gamma.\mathsf{end\text{–}users}$, update $\Gamma := \mathtt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$ and erase $(Q, id, cid, \widehat{\nu}, \widehat{\tau}_0)$ from your memory.
   - Else save $(P, id, cid, \nu, \tau_0)$ to your memory and send (instance–registering, id, cid) $\xhookrightarrow{\tau_1 \leq \tau_0 + \Delta}$ $\gamma.\mathsf{end\text{–}users}$.

Upon (finalize–register, id, cid) $\overset{\tau_2}{\hookleftarrow}$ P**,** let $\gamma := \Gamma(id)$ and do

- If $P \in \gamma.\mathsf{end\text{–}users}$ and your memory contains a value $(P, id, cid, \widehat{\nu}, \widehat{\tau}_0)$ such that $\tau_2 - \widehat{\tau}_0 \geq 2\Delta$, then set $\tilde{\nu} := (\widehat{\nu}.\mathsf{storage}, \widehat{\nu}.\mathsf{code})$, send (instance–registered, $id, cid, \tilde{\nu}$) $\xhookrightarrow{\tau_3 \leq \tau_2 + \Delta}$ $\gamma.\mathsf{end\text{–}users}$, set $\Gamma := \mathtt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$ and erase $(P, id, cid, \widehat{\nu}, \widehat{\tau}_0)$ from your memory.
- Else ignore this call.

Update a contract instance in a ledger state channel. An update of the storage of a contract instance in a ledger state channel starts when the environment sends the message (update, $id, cid, \tilde{\sigma}, \mathtt{C}$) to the initiating party $P \in \gamma.\mathsf{end\text{–}users}$ and works as follows. The initiating party P signs the new contract instance with increased version number (i.e. if ν is the contract instance version stored by P until now, then the new contract instance version ν' will be such that $\nu'.\mathsf{version} = \nu.\mathsf{version} + 1$). Party P then sends her signature on this value to the party $Q := \gamma.\mathsf{other\text{–}party}(P)$. The other party verifies the signature and informs the environment that the update was requested. If the environment confirms the update, the party Q signs the updated contract version and sends the signature to P. In this optimistic case, the update takes 2 rounds.

Let us discuss how parties behave in case the environment does not confirm the update. If Q simply aborts in this situation, P does not know if the update failed because Q is malicious or because the environment did not confirm the update. Therefore, Q has to inform P about the failure. This is, however, still not sufficient. Note that Q holds P's signature of the updated contract instance version. If Q is corrupt, he can register the updated contract instance on the ledger at any later point. Thus, party Q in order to convince P that he is not malicious, signs the *original* contract instance ν but with version number increased by 2 (i.e. the contract instance ν^* signed by Q is such that $\nu^*.\mathsf{storage} = \nu.\mathsf{storage}$, $\nu^*.\mathsf{code} = \nu.\mathsf{code}$ but $\nu^*.\mathsf{version} = \nu.\mathsf{version} + 2$). Party Q then sends the signature to party P. Note that since $\nu^*.\mathsf{storage} = \nu.\mathsf{storage}$, party P does not need to send her signature on ν^* back to Q.

If P does not receive a valid signature on either the updated contract instance version or the original contract instance with increased version number from Q, it is clear that Q is malicious and therefore P initiates the registration of the contract instance on the ledger by calling the procedure $\mathtt{Register}(P, id, cid)$. Note that Q can still register the updated contract instance (the one that was signed by P). But importantly, after at most $2 + 3\Delta$ rounds it will be clear to both parties what the current contract instance version is. Formal description of the protocol for updating a contract instance in a ledger state channel can be found below.

Execute a contract instance in a ledger state channel. In order to execute a contact instance stored in a ledger state channel γ, the environment sends the message (execute, $\gamma.\mathsf{id}, cid, f, z$) to the initiating party $P \in \gamma.\mathsf{end\text{–}users}$. The parameter cid points to the contract instance, f is the contact function and z are additional input values for f. For $P = \gamma.\mathsf{Alice}$ the protocol works as follows. If the parties never registered the contract instance with identifier cid, then $\gamma.\mathsf{Alice}$ first tries to execute the contract instance "peacefully". This means that she locally executes f on the contract version she stores in $\Gamma^{\gamma.\mathsf{Alice}}$, signs the new contract instance and sends the signature to $\gamma.\mathsf{Bob}$. Party $\gamma.\mathsf{Bob}$ also executes f locally on his own version of the contract instance stored in $\Gamma^{\gamma.\mathsf{Bob}}$ and thereafter verifies $\gamma.\mathsf{Alice}$'s signature. If the signature is valid, $\gamma.\mathsf{Bob}$ immediately confirms the execution by sending his signature on the new contract instance to party $\gamma.\mathsf{Alice}$.

A technical challenge occurs when both parties want to peacefully execute the same contract instance in the same round τ since it becomes unclear what is the new contract instance. This can be resolved be having designated rounds for each party.

In case the contract instance with identifier cid has already been registered on the ledger or the peaceful execution fails, the initiating party executes the contract instance "forcefully". By this we mean that $\gamma.\mathsf{Alice}$ first initiates registration of the contract instance by calling the procedure $\texttt{Register}(\gamma.\mathsf{Alice}, id, cid)$ if it was not done before, and then instructs the hybrid ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ to execute the contract instance. The $\texttt{Register}$ procedure can take up to 3Δ rounds and the contract instance execution on the ledger can take up to Δ rounds. Thus, pessimistic time complexity of the execution protocol is equal to $4\Delta + 5$ rounds. Formal description of the protocol for executing a contract instance in a ledger state channel and the corresponding part of the functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$ follow.

Protocol $\Pi(1, \mathcal{C})$: Contract instance execution

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, let $\mathcal{F}_{scc} := \mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$.

Party P upon $(\text{execute}, id, cid, f, z) \xhookleftarrow{\tau_0} \mathcal{Z}$

1. Let $\gamma^P := \Gamma^P(id), \nu^P := \gamma^P.\mathsf{cspace}(cid)$, $\sigma^P := \nu^P.\mathsf{storage}$, $\mathtt{C}^P := \nu^P.\mathsf{code}$ and set $w^P := \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}$.
2. Set $\tau_1 := \tau_0 + x$, where x is the smallest offset such that $\tau_1 = 1 \mod 4$ if $P = \gamma^P.\mathsf{Alice}$ and $\tau_1 = 3 \mod 4$ if $P = \gamma^P.\mathsf{Bob}$.
3. For round $\tau \in [\tau_0, \tau_1]$ proceed as follows: If (id, cid) is marked as corrupt in Γ^P_{aux}, goto step step 5.
4. Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^P, P, \tau_0, z)$. If $m = \bot$, then stop. Otherwise compute $s_P := \mathtt{Sign}_{sk_P}(id, cid, \tilde{\sigma}, \mathtt{C}^P, w^P)$, send $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookrightarrow{\tau_1} Q$ and goto step 12.
5. If (id, cid) is marked as corrupt but not registered, then execute $\texttt{Register}(id, cid, \nu^P)$.
6. Goto step 13.

Party Q upon $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookleftarrow{\tau_Q} P$

7. Let $\gamma^Q := \Gamma^Q(id), \nu^Q := \gamma^Q.\mathsf{cspace}(cid), \sigma^Q := \nu^Q.\mathsf{storage}, \mathtt{C}^Q := \nu^Q.\mathsf{code}, w^Q := \Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version}$. If $\gamma^Q = \bot$, $P, Q \notin \gamma^Q.\mathsf{end\text{–}users}$, $\nu^Q = \bot$ or $f \notin \mathtt{C}^Q$, then goto step 11.
8. If $P = \gamma^Q.\mathsf{Alice}$ and $\tau_Q \mod 4 \neq 2$ or if $P = \gamma.\mathsf{Bob}$ and $\tau_Q \mod 4 \neq 0$, then goto step 11.
9. If $\tau_0 \notin [\tau_Q - 4, \tau_Q - 1]$, then goto step 11.
10. If (id, cid) is not marked as corrupt in Γ^Q_{aux}, do:
    (a) Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^Q, P, \tau_0, z)$.
    (b) If $m = \bot$ or $\mathtt{Vfy}_{pk_P}(id, cid, \tilde{\sigma}, \mathtt{C}^Q, w^Q; s_P) \neq 1$, then goto step 11.
    (c) Output $(\text{execute–requested}, id, cid, f, z, \tau_0) \xhookrightarrow{\tau_Q} \mathcal{Z}$.
    (d) Sign $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}^Q, w^Q)$, send $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookrightarrow{\tau_Q} P$, set $\Gamma^Q := \texttt{UpdateChanSpace}(\Gamma^Q, id, cid, \tilde{\sigma}, \mathtt{C}^Q, add_A, add_R, w^Q, \{s_P, s_Q\})$, $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q + 1$.
    (e) Output $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_Q + 1} \mathcal{Z}$ and stop.
11. Mark (id, cid) as corrupt in Γ^Q. Then goto step 15.

Back to party P

12. Distinguish the following two cases

– If $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookleftarrow{\tau_2=\tau_1+2} Q$ such that $\mathtt{Vfy}_{pk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}^P, w^P; s_Q) = 1$, then set $\mathit{\Gamma}^P := \mathtt{UpdateChanSpace}(\mathit{\Gamma}^P, id, cid, \tilde{\sigma}, \mathtt{C}^P, w^P, \{s_P, s_Q\})$, $\mathit{\Gamma}^P_{aux}.\mathsf{next\text{-}version} := w^P + 1$, output $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_2} \mathcal{Z}$ and stop.

– Else mark (id, cid) as corrupt in $\mathit{\Gamma}^P_{aux}$ and execute the $\mathtt{Register}(P, id, cid)$. Once the procedure is executed (in round $\tau_3 \leq \tau_0 + 3\mathit{\Delta} + 5$) and it holds that $\mathit{\Gamma}^P(id).\mathsf{cspace}(cid).\mathsf{storage} = \tilde{\sigma}$ (i.e. Q registered the contract instance version after execution), then output $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_3} \mathcal{Z}$ and stop. Else goto the next step.

13. Send $(\text{instance–execute}, id, cid, f, z) \xhookrightarrow{\tau_3} \mathcal{F}_{scc}$.

$\boxed{\text{Back to party } Q}$

14. If $(\text{execute–requested}, id, cid, f, z, \tau) \xhookleftarrow{\tau_4 \leq \tau_0 + 4\mathit{\Delta} + 5} \mathcal{F}_{scc}$, output $(\text{execute–requested}, id, cid, f, z, \tau) \xhookrightarrow{\tau_4} \mathcal{Z}$.

$\boxed{\text{End for both parties } T = A, B}$

15. If $(\text{instance–executed}, id, cid, \widehat{\sigma}, add_L, add_R, m) \xhookleftarrow{\tau_4 \leq \tau_0 + 4\mathit{\Delta} + 5} \mathcal{F}_{scc}$, set $\mathit{\Gamma}^T := \mathtt{UpdateChanSpace}$ $(\mathit{\Gamma}^T, id, cid, \widehat{\sigma}, \mathtt{C}^T, add_L, add_R)$, output $(\text{executed}, id, cid, \widehat{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_4} \mathcal{Z}$ and stop. Else stop.

Functionality $\mathcal{F}^{\hat{\mathcal{L}}(\mathit{\Delta})}_{scc}(\mathcal{C})$: Contract instance execution

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1.

Upon $(\text{instance–execute}, id, cid, f, z, \tau) \xhookleftarrow{\tau_0} P$, within $\mathit{\Delta}$ rounds proceed as follows. Let $\gamma := \mathit{\Gamma}(id)$. If $\gamma = \perp$ or $\tau_0 - \tau > 6$, then stop. Else set $\nu := \gamma.\mathsf{cspace}(cid)$ and $\sigma := \nu.\mathsf{storage}$. If $P \neq \gamma.\mathsf{end\text{–}users}$, $\nu = \perp$ or $f \notin \nu.\mathsf{code}$, then stop. Else send $(\text{execute–requested}, id, cid, f, z, \tau) \xhookrightarrow{\tau_1 \leq \tau_0 + \mathit{\Delta}} \gamma.\mathsf{end\text{–}users}$ and compute $(\widehat{\sigma}, add_L, add_R, m) := f(\sigma, P, \tau, z)$. If $m = \perp$, then stop. Else update the channel space $\mathit{\Gamma} := \mathtt{UpdateChanSpace}(\mathit{\Gamma}, id, cid, \widehat{\sigma}, \nu.\mathsf{code}, add_L, add_R)$, send $(\text{instance–executed}, id, cid, \widehat{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_1 \leq \tau_0 + \mathit{\Delta}} \gamma.\mathsf{end\text{–}users}$ and stop.

Close a ledger state channel. In order to close a ledger state channel with identifier id by party $P \in \gamma.\mathsf{end\text{–}users}$, the environment sends the message (close, id) to the initiating party P. Before a ledger state channel can be closed, the end-users of the ledger state channel have the chance to register all the contract instances that they have constructed off-chain. Thus, the initiating party P first (in parallel) registers all the contract instances which have been updated/peacefully executed but not registered at the ledger yet. This takes up to $3\mathit{\Delta}$ rounds. Next, P asks the ideal functionality $\mathcal{F}^{\hat{\mathcal{L}}(\mathit{\Delta})}_{scc}(\mathcal{C})$ representing the state channel contract on the ledger to close the ledger state channel. Within $\mathit{\Delta}$ rounds, the ideal functionality informs both parties that the ledger state channel is being closed and gives the other end-user of the ledger state channel time $3\mathit{\Delta}$ to register contract instances that were not registered by P. If after $3\mathit{\Delta}$ rounds all registered contract instances are terminated, the ideal functionality adds $\gamma.\mathsf{cash}(\gamma.\mathsf{Alice})$ coins to $\gamma.\mathsf{Alice}$'s account on the ledger, and $\gamma.\mathsf{cash}(\gamma.\mathsf{Bob})$ coins to $\gamma.\mathsf{Bob}$'s account on the ledger, deletes the ledger state channel from its channel space and within $\mathit{\Delta}$ rounds informs both parties that the ledger state channel was successfully closed. Hence, in the pessimistic case closing can take up to $8\mathit{\Delta}$ rounds. The protocol and the contract functionality for the ledger state channel closing are presented formally below.

Protocol $\Pi(1,\mathcal{C})$: Close a ledger state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, let us denote $\mathcal{F}_{scc} := \mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$.

Party P upon $(\text{close}, id) \xhookleftarrow{\tau_0} \mathcal{Z}$

1. Let $\gamma^P := \Gamma^P(id)$. For each $cid \in \{0,1\}^*$ such that $\gamma^P.\mathsf{cspace}(cid) \neq \bot$ and (id, cid) is not marked as registered, execute $\texttt{Register}(P, id, cid)$ in round τ_0. Then send $(\text{contract–close}, id) \xhookrightarrow{\tau_1 \leq \tau_0 + 3\Delta} \mathcal{F}_{scc}$ and wait.

Party Q upon $(\text{contract–closing}, id) \xhookleftarrow{\tau_2 \leq \tau_0 + 4\Delta} \mathcal{F}_{scc}$

2. Let $\gamma^Q := \Gamma^Q(id)$. For each $cid \in \mathbb{N}$ such that (id, cid) is not marked as registered in Γ^Q and $\gamma^Q.\mathsf{cspace}(cid) \neq \bot$, call $\texttt{Register}(Q, id, cid)$ in round τ_2

Rest of the protocol for $T = P, Q$ (respectively):

3. If $(\text{contract–closed}, id) \xhookleftarrow{\tau_3 \leq \tau_0 + 8\Delta} \mathcal{F}_{scc}$, then set $\Gamma^T(id) := \bot$ and output $(\text{closed}, id) \xhookrightarrow{\tau_3} \mathcal{Z}$.

Functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$: Close a ledger state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1.

Upon $(\text{contract–close}, id) \xhookleftarrow{\tau_0} P$ let $\gamma := \Gamma(id)$ and proceed as follows:

1. Within Δ rounds send $(\text{contract–closing}, id) \xhookrightarrow{\tau_1 \leq \tau_0 + \Delta} \gamma.\mathsf{end–users}$.
2. Wait for next at most 3Δ rounds. If in round $\tau_2 \leq \tau_0 + 4\Delta$ there exists $cid \in \{0,1\}^*$ such that $\gamma.\mathsf{cspace}(cid) \neq \bot$ but the contract instance is not terminated, i.e. $\sigma_{cid}.\mathsf{locked} \neq 0$, where $\sigma_{cid} := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$, then stop.
3. Else wait for at most Δ round to add $\gamma.\mathsf{cash}(A)$ coins to A's account and $\gamma.\mathsf{cash}(B)$ coins to B's account on the ledger and set $\Gamma(id) = \bot$. Then send $(\text{contract–closed}, id) \xhookrightarrow{\tau_3 \leq \tau_0 + 5\Delta} \gamma.\mathsf{end–users}$.

7 Virtual State Channels

In this section, we first define the code of the virtual state channel contract $\texttt{VSCC}$, whose instances can be used to create virtual state channels γ. Then we describe the protocol $\Pi(i,\mathcal{C})$ that realizes the state channels ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ (see Sec. 4) for $i > 1$ and any set of contract codes $\mathcal{C}$.

7.1 Virtual State Channel Contract

The contract code $\texttt{VSCC}_i(\mathcal{C})$ is parameterized by a set $\mathcal{C}$ defining the contract codes that can be constructed in a virtual state channel γ of length $i > 0$. Consider three parties: Alice, Bob, and Ingrid, and suppose that Alice and Ingrid have opened a state channel α, and Bob and Ingrid have created a state channel β. During the creation of the virtual state channel γ between Alice and Bob, the parties Alice and Ingrid agree on updating α such that it contains the contract instance $(\sigma_A, \texttt{VSCC}_i(\mathcal{C}))$. Here, σ_A denotes the initial contract storage created by calling $\texttt{Init}_i^{\mathcal{C}}$, the constructor of $\texttt{VSCC}_i(\mathcal{C})$, on input tuple $(\text{Alice}, \tau, \gamma)$. On a very informal level, one may think of the contract storage $\sigma_A := \texttt{Init}_i^{\mathcal{C}}(\text{Alice}, \tau, \gamma)$ as being a "copy" of the virtual state channel description γ, where Ingrid plays the role of Bob. This "copy" of the virtual state channel γ will be stored in $\alpha.\mathsf{cspace}$ under the identifier $\mathit{cid}_A := \text{Alice}||\gamma.\mathsf{id}$. Symmetrically, Ingrid and Bob agree on updating their state channel β such that it contains the contract instance $(\sigma_B, \mathtt{VSCC}_i(\mathcal{C}))$, where $\sigma_B := \mathtt{Init}_i^{\mathcal{C}}(\text{Bob}, \tau, \gamma)$ is the initial state representing γ. This "copy" of the virtual state channel γ will be stored in $\beta.\mathsf{cspace}$ under the identifier $\mathit{cid}_B := \text{Bob}||\gamma.\mathsf{id}$.

Since Ingrid plays the role of Bob in the contract instance cid_A, and the role of Alice in the contract instance cid_B, in order to prevent her from losing money, she has to react to events happening in one of the contract instances and mimic them in the corresponding other contract instance. The contract functions of $\mathtt{VSCC}_i(\mathcal{C})$ are defined in such a way that they provide Ingrid with enough time to react on possible changes in cid_A or cid_B and to always keep both virtual state channel "copies" in the same state. In some sense, for the users in $\gamma.\mathsf{end{-}users}$, the contract instances referred to by cid_A and cid_B are now representing the contracts running on the ledger. They guarantee that as long as the parties $\gamma.\mathsf{end{-}users}$ behave honestly, they will never lose money.

Let us take a look at a simple example for the case when $i = 3$ (see Fig. 10). Suppose that each two consecutive parties $P_1, \dots, P_4$ have *ledger* state channel with each other. If P_1 and P_4 want to create a virtual state channel using the underlying *ledger* state channels, they can proceed recursively as follows. First, P_1 and P_3 create a virtual state channel γ' of length 2 between each other, where P_2 takes the role of Ingrid. This is done by creating a contract instance with code $\mathtt{VSCC}_2(\mathcal{C} \cup \mathtt{VSCC}_3(\mathcal{C}))$ in the ledger state channel between P_1 and P_2, resp. between P_2 and P_3. Let us take a closer look at the meaning of the contract code $\mathtt{VSCC}_2(\mathcal{C} \cup \mathtt{VSCC}_3(\mathcal{C}))$. Very informally, this contract code says that the virtual state channel is of length 2 (this is the reason for $\mathtt{VSCC}_2$), and that γ' can be used by its end-users to create contract instance with code from $\mathcal{C}$ and $\mathtt{VSCC}_3(\mathcal{C})$. The later are contracts that represent virtual state channels of length 3, which allows its end-users (of the length 3 virtual state channel) to open contract instances with code from $\mathcal{C}$. Next, parties P_1 and P_4 can open the virtual state channel of length 3, where party P_3 will take the role of Ingrid. To this end, P_1 and P_3 will use their previously created virtual state channel γ', and P_3 and P_4 will update their ledger state channel. The code of the contract instances in these two state channels is $\mathtt{VSCC}_3(\mathcal{C})$.

The contract code $\mathtt{VSCC}_i(\mathcal{C})$ will be described and formally defined together with the protocol for virtual state channels in Sec. 7.2. Here we provide only the interface of the contract $\mathtt{VSCC}_i(\mathcal{C})$.

Interface of the contract $\mathtt{VSCC}_i(\mathcal{C})$

Attributes:

- Mandatory attributes: $\mathsf{user}_L, \mathsf{user}_R, \mathsf{locked}, \mathsf{cash}_\mathsf{L}, \mathsf{cash}_\mathsf{R}$ (see Sec. 3.1)
- $\mathsf{virtual{-}channel}$: stores the initial version of the virtual state channel γ;
- cspace: stores the latest registered version of a contract instance created in γ;
- $\mathsf{preRegistered}$: stores a valid version of a contract instance of γ whose registration was not completed yet;
- $\mathsf{toExecute}$: auxiliary set storing all valid execution requests of a contract instance of γ.

Functions:

- $\mathtt{Init}_i^{\mathcal{C}}$: the constructor of the contract (see page 30);
- $\mathtt{RegisterInstance}_i^{\mathcal{C}}$: a contract function whose purpose is to register a contract instance created in the virtual state channel γ (see page 32);
- $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$: a contract function that is called during the force execution of a contract instance created in γ (see page 37);
- $\mathtt{Close}_i^{\mathcal{C}}$: a contract function called when the virtual state channel is being closed (see page 39).

7.2 Protocol for Virtual State Channels

We now describe the protocol $\Pi(i, \mathcal{C})$ that $\mathcal{E}_{res}$-realizes the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ for $i > 1$. The protocol is in the hybrid world with the hybrid ideal functionality which allows to create, update, execute and close state channels of lengths up to $i-1$ in which contract instances with code from the set $\mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C}$ can be constructed, i.e. the functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$.

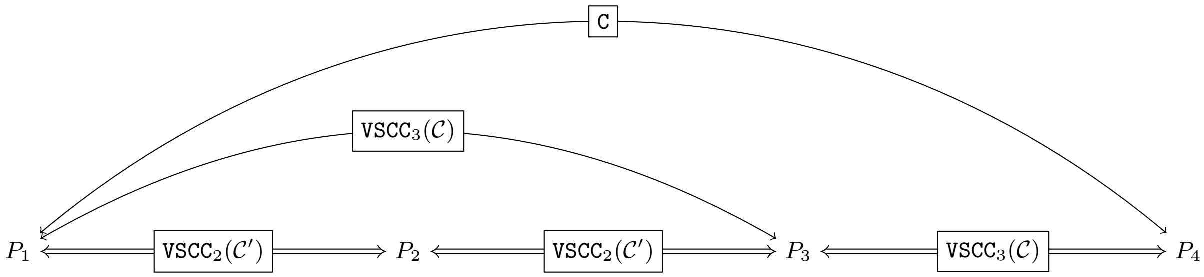

Fig. 10: The contract instance opened in state channels in order to create a virtual state channel of length 3 in which a contract instance $\mathtt{C} \in \mathcal{C}$ was opened. In the figure $\mathcal{C}' = \mathcal{C} \cup \mathtt{VSCC}_3(\mathcal{C})$.

The protocol consists of four subprotocols: Create a virtual state channel, Contract instance update, Contract instance execute and Close a virtual state channel. Similarly as for ledger state channels, we additionally define a procedure $\mathtt{Register}_i(P, id, cid)$ that registers a contract instance in a virtual state channel of length i and can be called by parties of the protocol $\Pi(i, \mathcal{C})$.

The protocol $\Pi(i, \mathcal{C})$ has to handle messages about state channels of any length j, where $1 \leq j \leq i$. If a party P of the protocol $\Pi(i, \mathcal{C})$ is instructed by the environment to create, update, execute or close a state channel of length $1 \leq j < i$, the party forwards this message (possibly with some pre-processing) to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$, and hence we focus on the protocol for virtual state channels of length exactly i.

Create a virtual state channel. To create the virtual state channel γ of length i in which contract instances with code from set $\mathcal{C}$ can be constructed, the environment sends a message (create, γ) to all three parties $\gamma.\mathsf{Alice}, \gamma.\mathsf{Bob}$ and $\gamma.\mathsf{Ingrid}$ in the same round τ_0. The creation of γ then works as follows.

As already explained in Sec. 7.1, the end-users of the virtual state channel, $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$, both need to construct a new contract instance with the code $\mathtt{VSCC}_i(\mathcal{C})$ in the subchannels they each have with $\gamma.\mathsf{Ingrid}$. Let us denote these state channels by α, β in the outline that follows below. To create these contract instances, party $\gamma.\mathsf{Alice}$ first locally computes the constructor $\mathtt{Init}_i^{\mathcal{C}}(\gamma.\mathsf{Alice}, \tau, \gamma)$ to obtain the initial admissible contract storage of $\mathtt{VSCC}_i(\mathcal{C})$. Recall that informally this contract storage can be viewed as a "copy" of the virtual state channel γ. Thereafter, she sends an update request of the state channel α to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$. At the same time, $\gamma.\mathsf{Bob}$ analogously requests the update of the state channel β. If $\gamma.\mathsf{Ingrid}$ receives update requests of both state channels α and β from $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$, she immediately confirms both of them. As already mentioned before, it is crucial for $\gamma.\mathsf{Ingrid}$ that either both both her state channels α and β are updated or none of them. Only then she is guaranteed that if she loses coins in the subchannel α, she can claim these coins back from the subchannel β.

To ensure that at the end of the protocol two honest users $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$ can conclude whether the virtual state channel γ was successfully created, there is one additional technicality in our protocol. Notice that if $\gamma.\mathsf{Ingrid}$ is honest, once $\gamma.\mathsf{Alice}$ receives a confirmation that her update request of α was successfully competed, she can conclude that the virtual state channel is created. However, we cannot assume that $\gamma.\mathsf{Ingrid}$ is honest. Hence, to guarantee that when both $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$ are honest they agree on whether γ was opened, they exchange confirmation messages at the end of the protocol. To conclude, if creation of a virtual state channel is successful, both end-users output $(\text{created}, \gamma)$ to the environment after 3 rounds.

We emphasize that creating a virtual state channel runs in constant time – independent of the ledger processing time Δ and length of the virtual state channel. This is in contrast to the *ledger* state channels with require always 2Δ time for creation. Formal description of the protocol for ledger state channel creation and the corresponding part of the contract code $\mathtt{VSCC}_i(\mathcal{C})$ are given below.

Protocol $\Pi(i, \mathcal{C})$: Create a virtual state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. We denote the hybrid functionality as $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ and the virtual state contract as $\mathtt{C} := \mathtt{VSCC}_i(\mathcal{C})$.

$\boxed{\text{Party } T \in \gamma.\mathsf{end\text{–}users} \text{ upon } (\text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{Z}\text{:}}$

1. Compute $\tilde{\sigma}_T := \mathtt{Init}_i^{\mathcal{C}}(T, \tau_0; \gamma)$ and send $(\text{update}, id_T, cid_T, \tilde{\sigma}_T, \mathtt{C}) \xhookrightarrow{\tau_0} \mathcal{F}_{ch}$, where $cid_T := T || \gamma.\mathsf{id}$ and $id_T := \gamma.\mathsf{subchan}(T)$.

$\boxed{\text{Party } I \text{ upon } (\text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{Z}\text{:}}$

2. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(\gamma.\mathsf{Alice}, \tau_0, \gamma)$ and $\tilde{\sigma}_B := \mathtt{Init}_i^{\mathcal{C}}(\gamma.\mathsf{Bob}, \tau_0, \gamma)$. Let $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$, $id_B := \gamma.\mathsf{subchan}(\gamma.\mathsf{Bob})$ and $cid_A := \gamma.\mathsf{Alice} || \gamma.\mathsf{id}$ and $cid_B := \gamma.\mathsf{Bob} || \gamma.\mathsf{id}$.
3. If both messages $(\text{update–requested}, id_A, cid_A, \tilde{\sigma}_A, \mathtt{C}) \xhookleftarrow{\tau_0+1} \mathcal{F}_{ch}$ and $(\text{update–requested}, id_B, cid_B, \tilde{\sigma}_B, \mathtt{C}) \xhookleftarrow{\tau_0+1} \mathcal{F}_{ch}$ are received, then set $\Gamma^I(\gamma.\mathsf{id}) := \gamma$ and send $(\text{update–reply}, ok, id_A, cid_A) \xhookrightarrow{\tau_0+1} \mathcal{F}_{ch}$ and $(\text{update–reply}, ok, id_B, cid_B) \xhookrightarrow{\tau_0+1} \mathcal{F}_{ch}$ and wait until time $\gamma.\mathsf{validity}$. Else stop.

$\boxed{\text{Back to } T \in \gamma.\mathsf{end\text{–}users}}$

4. If $(\text{updated}, id_T, cid_T) \xhookleftarrow{\tau_0+2} \mathcal{F}_{ch}$, then send $(\text{create–ok}, \gamma) \xhookrightarrow{\tau_0+2} \gamma.\mathsf{other\text{–}party}(T)$. If $(\text{create–ok}, \gamma) \xhookleftarrow{\tau_0+3} \gamma.\mathsf{other\text{–}party}(T)$, then set $\Gamma^T(\gamma.\mathsf{id}) := \gamma$ and output $(\text{created}, \gamma) \xhookrightarrow{\tau_0+3} \mathcal{Z}$.
5. Wait until time $\gamma.\mathsf{validity}$.

Contract $\mathtt{VSCC}_i(\mathcal{C})$: constructor $\mathtt{Init}_i^{\mathcal{C}}(P, \tau, \gamma)$

If $P \notin \gamma.\mathsf{end\text{–}users}$ or $\gamma.\mathsf{cash}(\gamma.\mathsf{Alice}) < 0$ or $\gamma.\mathsf{cash}(\gamma.\mathsf{Bob}) < 0$ or $\gamma.\mathsf{cspace}(cid) \neq \perp$ for some $cid \in \{0,1\}^*$ or $\gamma.\mathsf{validity} < \tau + 2 + 4 \cdot \text{TimeExeReq}(\lceil i/2 \rceil)$, then output $\perp$. Else output the attribute tuple σ defined as follows:

$$\begin{aligned}
(\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R) &:= \begin{cases} (\gamma.\mathsf{Alice}, \gamma.\mathsf{Ingrid}), & \text{if } P = \gamma.\mathsf{Alice}, \\ (\gamma.\mathsf{Ingrid}, \gamma.\mathsf{Bob}), & \text{if } P = \gamma.\mathsf{Bob}, \end{cases} \\
\sigma.\mathsf{locked} &:= \gamma.\mathsf{cash}(\gamma.\mathsf{Alice}) + \gamma.\mathsf{cash}(\gamma.\mathsf{Bob}), \\
(\sigma.\mathsf{cash}(\sigma.\mathsf{user}_L), \sigma.\mathsf{cash}(\sigma.\mathsf{user}_R)) &:= (\gamma.\mathsf{cash}(\gamma.\mathsf{Alice}), \gamma.\mathsf{cash}(\gamma.\mathsf{Bob})), \\
\sigma.\mathsf{virtual\text{–}channel} &:= \gamma, \\
\sigma.\mathsf{cspace}(cid) &:= \perp, \text{ for all } cid \in \{0,1\}^*, \\
\sigma.\mathsf{preRegistered} &:= \perp \\
\sigma.\mathsf{toExecute} &:= \emptyset.
\end{aligned}$$

Register a contract instance in a virtual state channel. Similarly to the procedure `Register` defined for ledger state channels, the subprotocol $\mathtt{Register}_i$ is called with parameters (P, id, cid) the first time end-users of a virtual state channel γ with identifier id disagree on a contract instance $\nu := \gamma.\mathsf{cspace}(cid)$. Intuitively, we need the intermediate party $\gamma.\mathsf{Ingrid}$ to play the role of the ledger and resolve the dispute between $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$. If the intermediary would be trusted, then both end-users could simply send their latest contract instance version to $\gamma.\mathsf{Ingrid}$, who would then decide whose contract instance version is the latest valid one. Unfortunately, the situation is more complicated since $\gamma.\mathsf{Ingrid}$ is not a trusted party. She might, for example, stop communicating or collude with one of the end-users. This is the point where the contract instances with code $\mathtt{VSCC}_i(\mathcal{C})$ created in the underlying subchannels during the virtual state channel creation play an important role. Parties instead of sending versions of ν directly to each other send them indirectly by executing the contract instances in their subchannels with $\gamma.\mathsf{Ingrid}$ on the contract function $\texttt{RegisterInstance}_i^{\mathcal{C}}$. Since this execution of the contract instance in the subchannel cannot be stopped (i.e., in the worst case it may involve the ledger which resolves the conflict), this guarantees that the end-users eventually can settle the latest state on which they both have agreed on. Let us now take a closer look at how this is achieved by $\texttt{VSCC}_i(\mathcal{C})$.

Let $cid_A := \gamma.\mathsf{Alice}||\gamma.\mathsf{id}$ be the contract instance with code $\texttt{VSCC}_i(\mathcal{C})$ stored in the state channel $\gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$ and $cid_B := \gamma.\mathsf{Bob}||\gamma.\mathsf{id}$ the contract instance stored in $\gamma.\mathsf{subchan}(\gamma.\mathsf{Bob})$. The initiating party (assume for now that it is $\gamma.\mathsf{Alice}$) first executes cid_A on the function $\texttt{RegisterInstance}_i^{\mathcal{C}}$ with input parameters (cid, ν^A), where ν^A is $\gamma.\mathsf{Alice}$'s current off-chain contract instance version. Notice that this execution is in a state channel of length strictly less than i and hence will be handled by the trusted hybrid ideal functionality $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$. The contract function $\texttt{RegisterInstance}_i^{\mathcal{C}}$ is defined in such a way that it first verifies the validity of $\gamma.\mathsf{Alice}$'s contract instance version, and if all checks pass, it stores (cid, ν^A) together with a time-stamp in the auxiliary attribute $\mathsf{preRegistered}$.

The intermediary $\gamma.\mathsf{Ingrid}$ upon receiving the information about the execution of cid_A on the function $\texttt{RegisterInstance}_i^{\mathcal{C}}$ with input parameters (cid, ν^A) can now symmetrically request execution of cid_B on $\texttt{RegisterInstance}_i^{\mathcal{C}}$ with input (cid, ν^A). We emphasize that $\gamma.\mathsf{Ingrid}$ only needs the information that cid_A is being executed and does not need to wait to start the execution until the execution of cid_A is completed.

Once $\gamma.\mathsf{Bob}$ is notified about the execution request of cid_B on $\texttt{RegisterInstance}_i^{\mathcal{C}}$ with input parameters (cid, ν^A), he immediately submits ν^B, his own off-chain contract instance version, by executing cid_B on the contract function $\texttt{RegisterInstance}_i^{\mathcal{C}}$ with input parameters (cid, ν^B). If $\gamma.\mathsf{Bob}$'s version of the contract instance with identifier cid was submitted in time and is valid, the contract function $\texttt{RegisterInstance}_i^{\mathcal{C}}$ compares the two submitted versions of cid and stores the one with higher version number in the attribute $\mathsf{cspace}(cid)$. Otherwise, ν^A will be considered as the registered one. Note that once honest $\gamma.\mathsf{Bob}$ learns about ν^A and submits ν^B, he knows whose contract instance version will be registered in cid_B. Thus, he can mark (id, cid) as registered in $\varGamma_{aux}^B$ and update his channel space accordingly without waiting for the execution of cid_B to be completed. We emphasize that there is no particular order in which parties can register state and our protocol can handle all possible variants.

Once $\gamma.\mathsf{Alice}$ receives the information about $\gamma.\mathsf{Bob}$'s version of the contract instance, she already knows whose contract instance version will be registered in cid_A. Thus, analogously to $\gamma.\mathsf{Bob}$, she can mark (id, cid) as registered $\varGamma_{aux}^A$ and update her channel space accordingly without waiting for the execution of cid_A to be completed. If $\gamma.\mathsf{Alice}$ does not receive any information about $\gamma.\mathsf{Bob}$'s version until certain round (because $\gamma.\mathsf{Bob}$ is corrupt and did not reveal his version to $\gamma.\mathsf{Ingrid}$ or because $\gamma.\mathsf{Ingrid}$ is corrupt and did not execute cid_A with $\gamma.\mathsf{Bob}$'s version in time), she can conclude that ν^A will be the registered contract instance version in cid_A and hence mark (id, cid) as registered in $\varGamma_{aux}^A$.

To conclude, the registration procedure of a virtual state channel of length i can take up to Time $\mathrm{Reg}(i) := 4 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil)$ rounds. This follows from the definition of the hybrid ideal functionality $\mathcal{F}_{ch}$ and our assumption that both subchannels have length at most $\lceil i/2 \rceil$ (see Appx. B).

Before we proceed to the formal description of the registration subprocedure, let us explain here the reason why we restrict the number of contract instances in a virtual state channel although the syntax as defined in Sec. 3 supports infinitely many contract instances as in the ledger state channel.

Assume the following scenario. Alice and Bob open a virtual state channel on top of two ledger state channels which they each have with Ingrid and thereafter they create (off-line) a large amount of contract instances in this virtual state channel. At some point Alice starts registering all the contract instances by executing the subchannel she has with Ingrid. According to the protocol, Ingrid has to symmetrically execute the subchannel she has with Bob otherwise she might lose money. If Bob is corrupt and does not react on peaceful execution requests, Ingrid has to execute all the requests forcefully on the blockchain. While in our theoretical model this is not an issue, in practice, this step would be very expensive for Ingrid due to the large amount of fees Ingrid would have to pay to the miners in common cryptocurrencies such as the Ethereum network. Thus, if Ingrid has no control over the amount of contract instances that Alice and Bob can create, the two parties can force Ingrid to pay arbitrary amount of money in fees. Therefore, we restrict the number of contract instance so that Ingrid can estimate the costs in fees that might result from the virtual state channel before she agrees to be the intermediary of that virtual state channel.

Protocol $\Pi(i,\mathcal{C})$: procedure $\texttt{Register}_i(P, id, cid)$

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, we define an auxiliary procedure `CompareVersions` whose formal definition can be found at the end of this protocol part. We denote $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\widehat{\mathcal{C}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ and $\mathrm{TER}_{sub} := \mathrm{TimeExeReq}(\lceil i/2 \rceil)$.

$\boxed{\text{Party } P\text{:}}$

1. Let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, $id_P := \gamma^P.\mathsf{subchan}(P)$, $cid_P := P||id$ and let τ_1^P be the current round. Then send $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookrightarrow{\tau_1^P} \mathcal{F}_{ch}$.

$\boxed{\text{Party } I\text{:}}$

2. Upon $(\text{execute–requested}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1^I} \mathcal{F}_{ch}$, proceed as follows. Set $\alpha := \Gamma^I(id_P)$, $P := \alpha.\mathsf{other\text{-}party}(I)$, $\sigma_P := \alpha.\mathsf{cspace}(cid_P).\mathsf{storage}$, $\gamma^I := \sigma_P.\mathsf{virtual–channel}$, $Q := \gamma^I.\mathsf{other\text{-}party}(P)$, $id_Q := \gamma^I.\mathsf{subchan}(Q)$ and $cid_Q := Q||\gamma^I.\mathsf{id}$. Then send $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookrightarrow{\tau_1^I} \mathcal{F}_{ch}$.

$\boxed{\text{Party } Q\text{:}}$

3. Upon $(\text{execute–requested}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1^Q} \mathcal{F}_{ch}$, parse $Q||id := cid_Q$, mark (id, cid) as corrupt in Γ_{aux}^Q and distinguish the following two situations:
   – If $\texttt{VerifyInstance}(id, cid, \nu^P) \neq 1$, then stop.
   – Else set $\nu^Q := \Gamma^Q(id).\mathsf{cspace}(cid)$ and send $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookrightarrow{\tau_1^Q} \mathcal{F}_{ch}$. Then execute $Q.\texttt{CompareVersions}(id, cid, \nu^P)$.

$\boxed{\text{Back to party } I\text{:}}$

4. If you receive $(\text{execute–requested}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_2^I \leq \tau_1^I + 2\cdot\mathrm{TER}_{sub}} \mathcal{F}_{ch}$, then send $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookrightarrow{\tau_2^I} \mathcal{F}_{ch}$.

$\boxed{\text{Back to party } P\text{:}}$

5. If $(\text{execute–requested}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\leq \tau_1^P + 4\cdot\mathrm{TER}_{sub}} \mathcal{F}_{ch}$ and it holds that $\texttt{VerifyInstance}(id, cid, \nu^Q) = 1$, then execute $P.\texttt{CompareVersions}(id, cid, \nu^Q)$.
6. Else mark (id, cid) as registered in Γ_{aux}^P and stop.

Auxiliary procedure: $T.\texttt{CompareVersions}(id, cid, \nu)$

Let $\nu^T := \Gamma^T(id).\mathsf{cspace}(cid)$.

1. If $\nu.\mathsf{version} \geq \nu^T.\mathsf{version}$, the set $\widehat{\nu} := \nu$. Else set $\widehat{\nu} := \nu^T$.
2. Mark (id, cid) as registered in Γ_{aux}^T and update the channel space $\Gamma^T := \texttt{UpdateChanSpace}^*(\Gamma^T, id, cid, \widehat{\nu})$.

Contract $\texttt{VSCC}_i(\mathcal{C})$

Function $\texttt{RegisterInstance}_i^{\mathcal{C}}(\sigma, P, \tau; (cid, \nu_n))$

Let $\gamma := \sigma.\mathsf{virtual\text{–}channel}, id := \gamma.\mathsf{id}, A := \gamma.\mathsf{Alice}, B := \gamma.\mathsf{Bob}, I := \gamma.\mathsf{Ingrid}$ and $\mathrm{TER}_{sub} := \mathrm{Time\ ExeReq}(\lceil i/2 \rceil)$.

1. Make the following checks:
   - $P \in \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\}$;
   - If $P \in \{A, B\}$, then $\tau \leq \gamma.\mathsf{validity}$ and if $P = I$, then $\tau \leq \gamma.\mathsf{validity} + \mathrm{TER}_{sub}$;
   - $\sigma.\mathsf{preRegistered} = \perp$ or $\sigma.\mathsf{preRegistered} = (Q, \tau^Q; cid, \nu_n^Q)$, where $P \neq Q$ and $\tau - \tau^Q \leq \mathrm{TER}_{sub}$ if $P \in \{A, B\}$ and $\tau - \tau^Q \leq 3 \cdot \mathrm{TER}_{sub}$ if $P = I$.
   - $\sigma.\mathsf{cspace}(cid') = \perp$ for every $cid' \in \{0, 1\}^*$;
   - If $\nu_n \neq \perp$, then $\texttt{VerifyInstance}(id, cid, \nu_n) = 1$; $\nu_n.\mathsf{code} \in \mathcal{C}$; and for $\sigma_n := \nu_n.\mathsf{storage}$ it holds that $\{\sigma_n.\mathsf{user}_L, \sigma_n.\mathsf{user}_R\} = \{A, B\}$; $\sigma_n \in \mathsf{C}.\Lambda$.

   If one or more checks fail, then output $(\sigma, 0, 0, \perp)$.
2. If all checks pass, then define $\tilde{\sigma} := \sigma$ and consider the following cases:
   - If $\sigma.\mathsf{preRegistered} := (P, \tau^P; cid, \nu_n^P)$, then output $(\tilde{\sigma}, 0, 0, \perp)$.
   - If $\sigma.\mathsf{preRegistered} = \perp$, then add $(P, \tau; cid, \nu_n)$ to $\tilde{\sigma}.\mathsf{preRegistered}$. In order to prevent force execution requests without parallel contract instance registration, delete from $\tilde{\sigma}.\mathsf{toExecute}$ all entries $e = (\tau', cid, P_n, \tau_n, f_n, z_n)$ such that $\tau' < \tau$ if $P \in \{A, B\}$ and $\tau' + \mathrm{TER}_{sub} < \tau$ if $P = I$. Then output $(\tilde{\sigma}, 0, 0, m)$, where $m := (\text{instance–registering}, cid, \nu_n)$.
   - If $\sigma.\mathsf{preRegistered} := (Q, \tau^Q; cid, \nu_n^Q)$ and $\nu_n = \perp$, then output $(\tilde{\sigma}, 0, 0, \perp)$.
   - If $\sigma.\mathsf{preRegistered} := (Q, \tau^Q; cid, \nu_n^Q)$ and $\nu_n \neq \perp$, then proceed as follows:
     (a) If $\nu_n^Q = \perp$ or $\nu_n.\mathsf{version} > \nu_n^Q.\mathsf{version}$, set $\widehat{\nu}_n := \nu_n$, and set $\widehat{\nu}_n := \nu_n^Q$ otherwise.
     (b) Set $\tilde{\sigma}.\mathsf{cspace}(cid) := (\widehat{\nu}_n.\mathsf{storage}, \widehat{\nu}_n.\mathsf{code})$ and modify the cash values accordingly, i.e. for $\tilde{\sigma}_n := \tilde{\sigma}.\mathsf{cspace}(cid).\mathsf{storage}$
       - If $\sigma.\mathsf{user}_L = I$, then $\tilde{\sigma}.\mathsf{cash}(I) := \sigma.\mathsf{cash}(I) - \tilde{\sigma}_n.\mathsf{cash}(A)$ and $\tilde{\sigma}.\mathsf{cash}(B) := \sigma.\mathsf{cash}(B) - \tilde{\sigma}_n.\mathsf{cash}(B)$.
       - If $\sigma.\mathsf{user}_R = I$, then $\tilde{\sigma}.\mathsf{cash}(I) := \sigma.\mathsf{cash}(I) - \tilde{\sigma}_n.\mathsf{cash}(B)$ and $\tilde{\sigma}.\mathsf{cash}(A) := \sigma.\mathsf{cash}(A) - \tilde{\sigma}_n.\mathsf{cash}(A)$.

       for $\tilde{\sigma}_n := \tilde{\sigma}.\mathsf{cspace}(cid).\mathsf{storage}$
     (c) In order to prevent double execution, if $\widehat{\nu}_n = \nu_n$ and there exists an entry $e := (\tau', cid, Q, \tau_n, f_n, z_n) \in \tilde{\sigma}.\mathsf{toExecute}$ such that for $(\tilde{\sigma}_n, add_L, add_R, m') := f_n(\nu_n^Q.\mathsf{storage}, Q, \tau_n, z_n)$ it holds that $\nu_n.\mathsf{storage} = \tilde{\sigma}_n$, then delete e from $\tilde{\sigma}.\mathsf{toExecute}$.
     (d) Finally, delete $(Q, \tau^Q; cid, \nu_n^Q)$ from $\tilde{\sigma}.\mathsf{preRegistered}$ and output $(\tilde{\sigma}, 0, 0, m)$, where $m := (\text{instance–registered}, cid, \widehat{\nu}_n)$.

Update a contract instance in a virtual state channel. As long as both end-users of a virtual state channel follow the protocol, they can update a contract instance exactly the same way as if it would be a ledger state channel. The differences between updates in a ledger state channel and in a virtual state channel appears only when end-users of the state channel run into dispute, i.e., when the parties run the contract instance registration procedure, which was defined above. The pessimistic time complexity of updating a virtual state channel of length i is equal to $\mathrm{TimeReg}(i) + 2$.

Execute a contract instance in a virtual state channel. In order to execute a contract instance in a virtual state channel γ with identifier id, the environment sends a message $(\text{execute}, id, cid, f, z)$ to one of the end-users of the virtual state channel. Let us assume for now that this party is $\gamma.\mathsf{Alice}$ and let τ_0 be the round when she received the message from the environment. The party $\gamma.\mathsf{Alice}$ first tries to execute the contract instance "peacefully", exactly as if γ would be a ledger state channel (see page 24). In case the peaceful execution fails, $\gamma.\mathsf{Alice}$ needs to register the contract instance cid by calling the sub-procedure $\texttt{Register}_i(\gamma.\mathsf{Alice}, id, cid)$ and execute the contract instance "forcefully" via the intermediary $\gamma.\mathsf{Ingrid}$. Since the intermediary is not trusted, execution must be performed by executing the contract instances with code $\texttt{VSCC}_i(\mathcal{C})$ stored in the underlying subchannels of γ (recall that the contract instance in the subchannel $\gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$ is stored under the identifier $cid_A := \gamma.\mathsf{Alice}||\gamma.\mathsf{id}$ and the contract instance in the state channel $\gamma.\mathsf{subchan}(\gamma.\mathsf{Bob})$ is stored under the identifier $cid_B := \gamma.\mathsf{Bob}||\gamma.\mathsf{id}$). Since both subchannels are state channels of length strictly less than i, the execution of their contract instances is handled by recursion via the trusted hybrid ideal functionality $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$.

The first attempt to design the force execution protocol would be to let $\gamma.\mathsf{Alice}$ execute cid_A on the function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ with parameters $param = (cid, \gamma.\mathsf{Alice}, \tau_0, f, z, s_A)$, where s_A is $\gamma.\mathsf{Alice}$'s signature on the tuple $(cid, \gamma.\mathsf{Alice}, \tau_0, f, z)$. We call the value τ_0 the *time-stamp* of the execution request. The contract function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ would be defined such that it verifies the execution request (for example, checks that $\gamma.\mathsf{Alice}$'s signature is valid, etc.) and then executes the contract instance with identifier cid. After successful execution of cid_A, $\gamma.\mathsf{Ingrid}$ symmetrically executes cid_B on the same contract function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ with the same input parameters $param$. It is important to emphasize that $\gamma.\mathsf{Ingrid}$ is not able to modify the tuple $param$ (for example change the time-stamp) in a way that would be accepted by $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ since she would need to forge $\gamma.\mathsf{Alice}$'s signature.

Unfortunately, this straightforward solution does not work since we allow parties to interact fully concurrently. To illustrate the problem consider an example where while the execution between $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Ingrid}$ is running, $\gamma.\mathsf{Bob}$ also wants to forcefully execute the contract instance with identifier cid on different inputs. This means that before $\gamma.\mathsf{Ingrid}$ has time to execute cid_B on $\gamma.\mathsf{Alice}$'s request, $\gamma.\mathsf{Bob}$ executes cid_B on the function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ with his own parameters $param' = (cid, \gamma.\mathsf{Bob}, \tau_0', f', z', s_B)$. Consequently, the order of executions of the contract instance cid is different in cid_A and cid_B. Depending on the contract code of cid, this asymmetry may lead to $\gamma.\mathsf{Ingrid}$ losing money.

Therefore, the contract function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}$ is defined in such a way that it verifies the validity of the submitted contract instance execution request as before and if all checks pass, then it only *stores* the execution request in an auxiliary attribute $\mathsf{toExecute}$. In other words, during the lifetime of the virtual state channel γ, the contract instances cid_A and cid_B in the subchannels of γ only collect information about the force executions of cid but they do not perform any of them. Thus, if $\gamma.\mathsf{Ingrid}$ mimics all requests from cid_A to cid_B and vice versa, then, after the last accepted force execution, the (unordered) set $\mathsf{toExecute}$ stored in cid_A is equal to the (unordered) set $\mathsf{toExecute}$ stored in cid_B. This is illustrated in Fig. 11 on a concrete example of three force execution requests.

All the internal executions are postponed until the virtual state channel is being closed and the contract instance cid_A and cid_B are being terminated. Looking ahead, the contract function $\mathtt{Close}_i^{\mathcal{C}}$ first sorts the elements of the set $\mathsf{toExecute}$ by their time-stamp and only then preforms all the internal executions. This guarantees the same order of internal executions in both cid_A and cid_B which implies that both of these contract instances terminate with the same money distribution. Detailed description of the closing procedure is given later in this section (see page 38).

To complete the description of our force execution protocol, it remains to discuss how do end-users of the virtual state channel learn the result of the force execution and when do they output it to the environment. Since internal executions of cid are postponed until the virtual state channel closure, end-users of the virtual state channel cannot wait until they learn the results from the hybrid ideal functionality $\mathcal{F}_{ch}$ as they did in the straightforward solution. Instead, they have to derive the results themselves. They proceed as follows. Party $\gamma.\mathsf{Alice}$, after initiating the force execution, waits for $2 \cdot \mathrm{TimeExe}(\lceil i/2 \rceil) + 5$ rounds to be sure that $\gamma.\mathsf{Bob}$ did not initiate force execution of cid that should be performed before her own force execution request. After the waiting is over, she performs her execution of cid locally and outputs the result to the environment. The other party acts similarly. Once $\gamma.\mathsf{Bob}$ learns about $\gamma.\mathsf{Alice}$'s force execution, he checks if he has some pending execution requests that should take place before the one requested by $\gamma.\mathsf{Alice}$. If this is the case then he locally executes them first. Thereafter, he locally executes the newly requested by $\gamma.\mathsf{Alice}$ and outputs the result to the environment.

Execution of a contract instance in a virtual state channel of length i as described above would take in the pessimistic case up to $5 + \mathrm{TimeReg}(i) + 2 \cdot \mathrm{TimeExe}(\lceil i/2 \rceil)$ rounds. Unfortunately, it turns out that the above time complexity is polynomial in the length of the virtual state channel. In order to achieve *linear* pessimistic time complexity, we make two important observations which optimize our protocol.

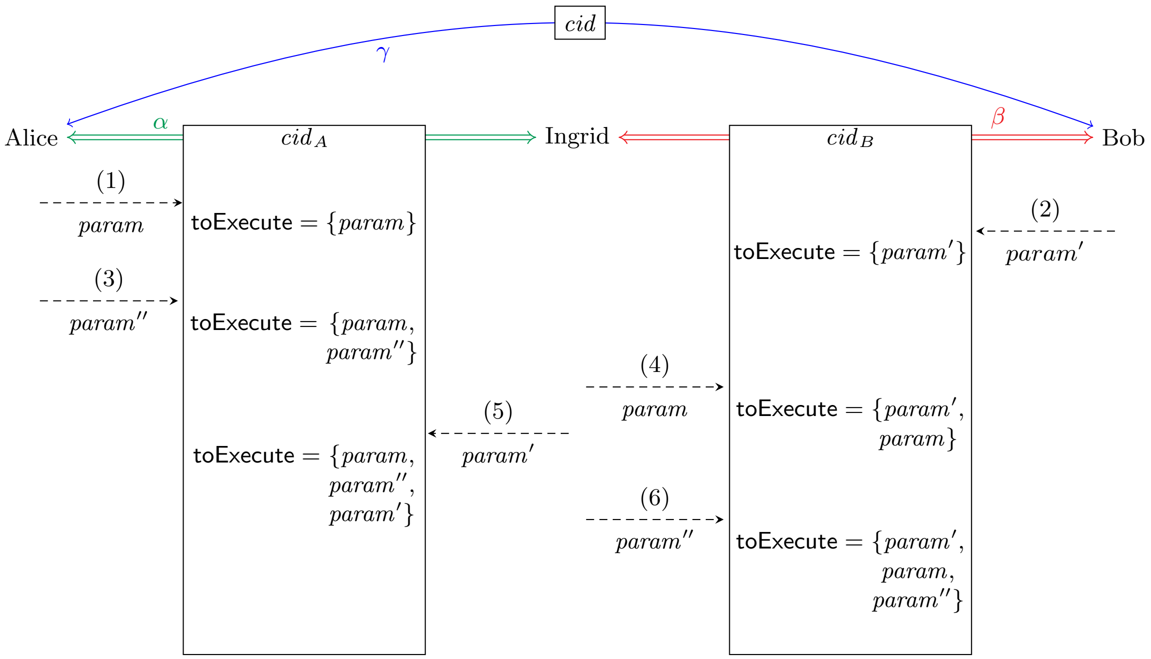

cid
γ
Alice
α
cid_A
Ingrid
cid_B
β
Bob
(1)
$param$
$\mathsf{toExecute} = \{param\}$
(2)
$param'$
$\mathsf{toExecute} = \{param'\}$
(3)
$param''$
$\mathsf{toExecute} = \{param, param''\}$
(4)
$param$
$\mathsf{toExecute} = \{param', param\}$
(5)
$param'$
$\mathsf{toExecute} = \{param, param'', param'\}$
(6)
$param''$
$\mathsf{toExecute} = \{param', param, param''\}$

Fig. 11: In the depicted example, Alice and Bob make three force execution requests of the contract instance cid created in their virtual state channel γ. In Step (1), Alice executes cid_A on the function $\texttt{ExecuteInstance}_i^{\mathcal{C}}$ with parameters $param = (cid, \gamma.\mathsf{Alice}, \tau_0, f, z, s_A)$. In Step (2) Bob executes cid_B on the function $\texttt{ExecuteInstance}_i^{\mathcal{C}}$ with parameters $param' = (cid, \gamma.\mathsf{Bob}, \tau_0, f', z', s_B)$ and in Step (3) Alice executes cid_A on the function $\texttt{ExecuteInstance}_i^{\mathcal{C}}$ with parameters $param'' = (cid, \gamma.\mathsf{Alice}, \tau_0'', f'', z'', s_A'')$, where $\tau_0 < \tau_0''$. Both Alice's requests are stored in the set $\mathsf{toExecute}$ in cid_A before Ingrid mimics the execution request made by Bob (see Step (5)). The example assumes that Ingrid mimics both Alice's requests (see Steps (4) and (6)). Although the force execution requests were received by cid_A and cid_B in a different order, the (unordered) sets $\mathsf{toExecute}$ are at the end identical.[19]

First note that party $\gamma.\mathsf{Ingrid}$ does not need to wait until execution of cid_A is completed in order to initiate the execution of cid_B. Similarly, $\gamma.\mathsf{Bob}$ does not need to wait until execution of cid_B is completed to locally execute the contract instance cid and output the result to the environment. This reduces the time complexity to $5 + \mathrm{TimeReg}(i) + 2 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil)$ rounds. Secondly, we observe that the registration subprocedure can be run in parallel with the force execution phase which reduces the pessimistic time complexity to:

$$\begin{aligned} \mathrm{TimeExeReq}(i) &:= 5 + 2 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil), \\ \mathrm{TimeExe}(i) &:= 5 + 4 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil). \end{aligned} \tag{1}$$

The previous description omits some technicalities and we refer the reader for further details to the full specification of the protocol and the corresponding part of the contract code $\texttt{VSCC}_i(\mathcal{C})$ which can be found below.

[19] Let us emphasize that the figure does not depict the only possible order in which the three force execution requests can be added to the sets $\mathsf{toExecute}$. It might, for instance, happen that the request $param$ is added to $\mathsf{toExecute}$ after the request $param''$. This is possible since the executions of cid_A and cid_B are performed in a black-box way via the hybrid ideal functionality $\mathcal{F}_{ch}$ and thus for each request we know only a time interval in which it must be added to $\mathsf{toExecute}$. The exact round is determined by the adversary (see Sec. 4). Hence, if the execution requests $param$ and $param''$ are made shortly after each other, the adversary has the power to swap the order in which they are stored.

Protocol $\Pi(i,\mathcal{C})$: Contract instance execution

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, we define an auxiliary procedure `VerifyInstance` whose formal description can be found at the end of this protocol part. We denote $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ and $\mathrm{TER}_{sub} := \mathrm{TimeExeReq}(\lceil i/2 \rceil)$.

Party P upon $(\text{execute}, id, cid, f, z) \xhookleftarrow{\tau_0} \mathcal{Z}$

1. Let $\gamma^P := \Gamma^P(id), \nu^P := \gamma^P.\mathsf{cspace}(cid), \sigma^P := \nu^P.\mathsf{storage}, \mathtt{C}^P := \nu^P.\mathsf{code}, w^P := \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}$ and $Q := \gamma^P.\mathsf{other\text{-}party}(P)$.
2. Set $\tau_1 := \tau_0 + x$, where x is the smallest offset such that $\tau_1 = 1 \mod 4$ if $P = \gamma^P.\mathsf{Alice}$ and $\tau_1 = 3 \mod 4$ if $P = \gamma^P.\mathsf{Bob}$.
3. For round $\tau \in [\tau_0, \tau_1]$ proceed as follows: If (id, cid) is marked as corrupt in Γ^P_{aux}, goto step step 11.
4. In round τ_1, compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^P, P, \tau_0, z)$. If $m = \bot$, then stop. Otherwise compute $s_P := \mathtt{Sign}_{sk_P}(id, cid, \tilde{\sigma}, \mathtt{C}^P, w^P)$, send $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookrightarrow{\tau_1} Q$ and goto step 10.

Party Q upon $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookleftarrow{\tau_1^Q} P$

5. Let $\gamma^Q := \Gamma^Q(id), \nu^Q := \gamma^Q.\mathsf{cspace}(cid), \sigma^Q := \nu^Q.\mathsf{storage}, \mathtt{C}^Q := \nu^Q.\mathsf{code}, w^Q := \Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version}$. If $\gamma^Q = \bot$, $P, Q \notin \gamma^Q.\mathsf{end\text{–}users}$, $\nu^Q = \bot$ or $f \notin \mathtt{C}^Q$, then goto step 9.
6. If $P = \gamma^Q.\mathsf{Alice}$ and $\tau_1^Q \mod 4 \neq 2$ or if $P = \gamma.\mathsf{Bob}$ and $\tau_1^Q \mod 4 \neq 0$, then goto step 9.
7. If $\tau_0 \notin [\tau_1^Q - 4, \tau_1^Q - 1]$ or $\tau_0 \geq \gamma.\mathsf{validity}$, then goto step 9.
8. If (id, cid) is not marked as corrupt in Γ^Q_{aux}, do:
   (a) Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^Q, P, \tau_0, z)$. If $m = \bot$ or $\mathtt{Vfy}_{pk_P}(id, cid, \tilde{\sigma}, \mathtt{C}^Q, w^Q; s_P) \neq 1$, then goto step 9.
   (b) Output $(\text{execute–requested}, id, cid, f, z) \xhookrightarrow{\tau_1^Q} \mathcal{Z}$.
   (c) Sign $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}^Q, w^Q)$, send $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookrightarrow{\tau_1^Q} P$, set $\Gamma^Q := \mathtt{UpdateChanSpace}(\Gamma^Q, id, cid, \tilde{\sigma}, \mathtt{C}^Q, add_L, add_R, w^Q, \{s_P, s_Q\})$ and $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q + 1$.
   (d) Output $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_1^Q + 1} \mathcal{Z}$ and stop.
9. Mark (id, cid) as corrupt in Γ^Q_{aux}. Then goto step 14.

Back to party P

10. Distinguish the following two cases
    - If $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookleftarrow{\tau_2 = \tau_1 + 2} Q$ such that $\mathtt{Vfy}_{pk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}^P, w^P; s_Q) = 1$, then set $\Gamma^P := \mathtt{UpdateChanSpace}(\Gamma^P, id, cid, \tilde{\sigma}, \mathtt{C}^P, add_L, add_R, w^P, \{s_P, s_Q\}), \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version} := w^P + 1$, output $(\text{executed}, id, cid, \tilde{\sigma}, add_L, add_R, m) \xhookrightarrow{\tau_2} \mathcal{Z}$ and stop.
    - Else mark (id, cid) as corrupt in Γ^P_{aux} and goto step 11.
11. Let τ_3 be the current round $(\tau_3 \leq \tau_0 + 5)$, let $id_P := \gamma^P.\mathsf{subchan}(P)$, $cid_P := P||id$, $s_n := \mathtt{Sign}_{sk_P}(cid, P, \tau_0, f, z)$ and $p_n := (P, \tau_0, f, z, s_n)$. Then send $(\text{execute}, id_P, cid_P, \mathtt{ExecuteInstance}_i^{\mathcal{C}}, (cid, p_n)) \xhookrightarrow{\tau_3} \mathcal{F}_{ch}$ and add (f, P, z, τ_0) to $\Gamma^P_{aux}(id, cid).\mathsf{toExecute}$. If (id, cid) is not marked as registered in Γ^P_{aux}, then run in parallel $\mathtt{Register}_i(P, id, cid)$.
12. Wait until round $\tau_4 := \tau_0 + 4 \cdot \mathrm{TER}_{sub} + 5$. In order to prevent double execution, first check if $\Gamma^P(id).\mathsf{cspace}(cid).\mathsf{storage} = \tilde{\sigma}$. If this is the case (i.e. Q registered the contract instance version after

execution), then delete (f, P, z, τ_0) from $\Gamma_{aux}(id, cid).\mathsf{toExecute}$, output (executed, id, cid, $\tilde{\sigma}$, add_L, $add_R, m) \xhookrightarrow{\tau_4} \mathcal{Z}$ and stop. Otherwise execute $P.\mathtt{LocalExe}(\tau_4, id, cid, \tau_0)$.

Party I:

13. Upon receiving (execute–requested, $id_P, cid_P, \mathtt{ExecuteInstance}_i^{\mathcal{C}}, (cid, p_n)) \xhookleftarrow{\tau_1^I} \mathcal{F}_{ch}$, proceed as follows. Set $\alpha := \Gamma^I(id_P)$, $P := \alpha.\mathsf{other\text{-}party}(I)$, $\sigma_P := \alpha.\mathsf{cspace}(cid_P).\mathsf{storage}$, $\gamma^I := \sigma_P.\mathsf{virtual\text{–}channel}$, $Q := \gamma^I.\mathsf{other\text{-}party}(P)$, $id_Q := \gamma^I.\mathsf{subchan}(Q)$ and $cid_Q := Q||\gamma^I.\mathsf{id}$. Then send (execute, id_Q, cid_Q, $\mathtt{ExecuteInstance}_i^{\mathcal{C}}, (cid, p_n)) \xhookrightarrow{\tau_1^I} \mathcal{F}_{ch}$.

Party Q:

14. Upon receiving (execute–requested, $id_Q, cid_Q, \mathtt{ExecuteInstance}_i^{\mathcal{C}}, (cid, p_n)) \xhookleftarrow{\tau_2^Q} \mathcal{F}_{ch}$, then parse $Q||id := cid_Q$ and $(P, \tau_0, f, z, s_n) := p_n$. Let $\gamma^Q := \Gamma^Q(id)$ and $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$. If $\gamma^Q = \bot$, $\nu^Q = \bot$, then stop. Else mark (id, cid) as corrupt in Γ_{aux}^Q. If $\mathtt{Vfy}_{pk_P}(cid, P, \tau_0, f, z; s_n) \neq 1$ or $f \notin \nu^Q.\mathsf{code}$, then stop. Else add (f, P, z, τ_0) to $\Gamma_{aux}^Q(id, cid).\mathsf{toExecute}$.
15. If (id, cid) is marked as registered in Γ_{aux}^Q, then output (execute–requested, id, cid, f, $z) \xhookrightarrow{\tau_2^Q} \mathcal{Z}$ and execute the subprocedure $Q.\mathtt{LocalExe}(\tau_2^Q, id, cid, \tau_0)$.
16. Else wait until round $\tau_3^Q := \tau_0 + 2 \cdot \mathrm{TER}_{sub} + 5$. If (id, cid) is not marked as registered after the waiting is over (i.e. P requested force execution without starting the contract instance registration in parallel), then delete (f, P, z, τ_0) from $\Gamma_{aux}(id, cid).\mathsf{toExecute}$ and stop. Otherwise output (execute–requested, id, cid, f, $z) \xhookrightarrow{\tau_3^Q} \mathcal{Z}$ and execute the subprocedure $Q.\mathtt{LocalExe}(\tau_3^Q, id, cid, \tau_0)$.

Auxiliary procedure: $T.\mathtt{LocalExe}(\tau, id, cid, \tau_0)$

1. Let $\gamma := \Gamma^T(id)$ and let $\sigma^{(0)} := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$.
2. Let $E \subseteq \Gamma_{aux}^T(id, cid).\mathsf{toExecute}$ consist of all tuples (f', T', z', τ_0'), where $\tau_0' \leq \tau_0$.
3. Let $|E| = \ell$ and $(e^{(1)}, \dots, e^{(\ell)})$ be such that $e^{(k)} = (\tau^{(k)}, T_n^{(k)}, \tau_n^{(k)}, f_n^{(k)}, z_n^{(k)}) \in E$ for every $k \in [1, \ell]$, $\tau_n^{(1)} \leq \dots \leq \tau_n^{(\ell)}$ and if $\tau_n^{(i)} = \tau_n^{(j)}$ for some $i < j$, then the following holds:
   - If $T^{(i)} \neq T^{(j)}$, then $T^{(i)} = A$ and $T^{(j)} = B$.
   - If $T^{(i)} = T^{(j)}$, then either $f_i <_{\mathtt{C}} f_j$, where $<_{\mathtt{C}}$ is total ordering of the contract functions defined by the contract code $\mathtt{C}$, or $f_i = f_j$ and $z_n^{(i)} \leq_{\mathrm{lex}} z_n^{(j)}$, where $\leq_{\mathrm{lex}}$ is the lexicographic ordering of binary strings.
4. For $k = 1$ to ℓ
   (a) Compute $(\sigma^{(k)}, add_L^{(k)}, add_R^{(k)}, m^{(k)}) := f(\sigma^{(k-1)}, T^{(k)}, \tau_0^{(k)}, z^{(k)})$.
   (b) Output (executed, id, cid, $\sigma^{(k)}$, $add_L^{(k)}$, $add_R^{(k)}, m^{(k)}) \xhookrightarrow{\tau} \mathcal{Z}$
   (c) Set $\Gamma^T := \mathtt{UpdateChanSpace}(\Gamma^T, id, cid, \sigma^{(k)}, \mathtt{C}, add_L^{(k)}, add_R^{(k)})$, where $\mathtt{C} := \gamma.\mathsf{cspace}(cid).\mathsf{code}$.
   (d) Delete $e^{(k)}$ from $\Gamma_{aux}^T(id, cid).\mathsf{toExecute}$.

Contract $\mathtt{VSCC}_i(\mathcal{C})$

Function $\mathtt{ExecuteInstance}_i^{\mathcal{C}}(\sigma, P, \tau, (cid, P_n, \tau_n, f_n, z_n, s_n))$

Let $\gamma := \sigma.\mathsf{virtual\text{–}channel}, A := \gamma.\mathsf{Alice}, B := \gamma.\mathsf{Bob}, I := \gamma.\mathsf{Ingrid}$. In addition, let $\mathrm{TER}_{sub} :=$ Time ExeReq($\lceil i/2 \rceil$). First, make the following checks:

- $P \in \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\}$ and $P_n \in \{A, B\}$;
- $\mathtt{Vfy}_{pk_{P_n}}(cid, P_n, \tau_n, f_n, z_n; s_n) = 1$;
- If $P \in \{A, B\}$, then $\tau - \tau_n \leq 5$ and $P = P_n$;
- If $P = I$, then $\tau - \tau_n \leq 5 + \mathrm{TER}_{sub}$ and $P_n \in \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\} \cap \{A, B\}$;
- $\tau_n \leq \gamma.\mathsf{validity}$.

If one of the checks fails, then output $(\sigma, 0, 0, \perp)$. Else let $\tilde{\sigma} := \sigma$, add $(\tau, cid, P_n, \tau_n, f_n, z_n)$ to $\tilde{\sigma}.\mathsf{toExecute}$ and output $(\tilde{\sigma}, 0, 0, m)$ for $m :=$ (instance–executing, $cid, P_n, \tau_n, f_n, z_n, s_n$).

Close a virtual state channel. Recall that in case of ledger state channels, the environment instructs one party to close the ledger state channel. The parties of the ledger state channel have some time to register all contract instances that were opened in the ledger state channel off-chain. If thereafter there is a contract instance in the ledger state channel which is not terminated, then the ledger state channel is not closed.

For virtual state channels the situation is different. We require that the closing procedure of a virtual state channel γ always starts in round $\gamma.\mathsf{validity}$ and always results in γ being closed. In other words, both contract instances with code $\mathtt{VSCC}_i(\mathcal{C})$ that were opened in the subchannels of γ must be terminated. This ensures that virtual state channels can never infinitely block closure of ledger state channels. Let us now explain how the protocol "Close a virtual state channel" works.

In round $\gamma.\mathsf{validity}$ both end-users of the virtual state channel start registering the contract instance if it has been created in the virtual state channel γ but has never been registered before. Thereafter, $\gamma.\mathsf{Alice}$ requests execution of the contract instance $cid_A := \gamma.\mathsf{Alice}||\gamma.\mathsf{id}$ stored in the subchannel $\gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$, on the contact function $\mathtt{Close}_i^{\mathcal{C}}$. In case $\gamma.\mathsf{Alice}$ is corrupt and does not request execution of cid_A on the function $\mathtt{Close}_i^{\mathcal{C}}$, $\gamma.\mathsf{Ingrid}$ can request it herself after certain time has passed. We proceed similarly for $\gamma.\mathsf{Bob}$.

The contract function $\mathtt{Close}_i^{\mathcal{C}}$ is defined in such a way that it first (if necessary) finalizes registration of a contract instance cid.[20] This step is needed in case one party initiated registration of a contract instance with identifier cid but the other party did not react. As a next step the function $\mathtt{Close}_i^{\mathcal{C}}$ internally executes all the force execution requests stored in $\mathsf{toExecute}$. As already discussed before (see page 34), the order in which the force execution requests of the contract instance cid are processed is very important. To this end, the contract function $\mathtt{Close}_i^{\mathcal{C}}$ sorts the elements of the set $\mathsf{toExecute}$, which are tuples of the form (cid, P, τ_0, f, z), according to the following rules:

1. **Time-stamp** τ_0: requests with lower time-stamp are processed first;
2. **Party** P: requests made by $\gamma.\mathsf{Alice}$'s are processed first;
3. **Function call** f: we assume that every contract code contains a total ordering of its functions;
4. **Input parameter** z: the lexicographic ordering of binary string is applied.

Consider again the example from Fig. 11. Recall that the set $\mathsf{toExecute}$ has at the end of execution phase three element: $param = (cid, \gamma.\mathsf{Alice}, \tau_0, f, z, s_A)$; $param' = (cid, \gamma.\mathsf{Bob}, \tau_0, f', z', s_B)$ and $param'' = (cid, \gamma.\mathsf{Alice}, \tau_0'', f'', z'', s_A'')$, where $\tau_0 < \tau_0''$. By applying the rules from above, the execution request defined by $param''$ will be processed last since it has the highest time-stamp. Since both $param$ and $param'$ have the same time-stamp, the Rule 2. is applied. To conclude, the function $\mathtt{Close}_i^{\mathcal{C}}$ will first process $param$, then $param'$ and lastly $param''$.

Let us now discuss what happens if there exists a registered contract instance cid which is however not terminated (the amount of locked coins is not equal to zero). The first idea would be to let $\mathtt{Close}_i^{\mathcal{C}}$ ignore the contract instance. However, this would lead to the problem that the intermediary of the virtual state channel, $\gamma.\mathsf{Ingrid}$, loses money (because some money may still be locked in the contract) without ever having the chance to react to virtual state channel closing. Instead, the contract function $\mathtt{Close}_i^{\mathcal{C}}$ gives all the locked coins in the contract instance to the intermediary. This implies that end-users of a virtual channels are responsible to open a contract instance only if they are certain that they can terminate it before the channel validity expires since otherwise they will lose money.

[20] Recall that we assume that there can be at most one contract instance in a virtual state channel.

Finally, the contract function verifies that the current value of the attribute cash is non-negative for both users and that the amount of coins that were originally invested into the virtual state channel is equal to the current amount of coins in the virtual state channel. If this is the case, $\texttt{Close}_i^{\mathcal{C}}$ unlocks for each user the current amount of coins it holds in the channel contract. If one of the users have negative balance in the virtual state channel or the amount of invested coins is not equal to the current amount of coins, then any trading that happened between the end-users of γ is reverted by $\texttt{Close}_i^{\mathcal{C}}$. This again guarantees that $\gamma.\mathsf{Ingrid}$ cannot lose money when $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$ are malicious.

The time complexity of closing a virtual state channel of length i can be computed as $2 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil) + 2 \cdot \mathrm{TimeExe}(\lceil i/2 \rceil)$. This follows from the simple observation that in case parties need to register a contract instance before closing the channel, both end-users should initiate the registration procedure in the same round (i.e. $\texttt{Register}_i(\gamma.\mathsf{Alice}, id, cid)$ and $\texttt{Register}_i(\gamma.\mathsf{Bob}, id, cid)$ are run in parallel) which reduces the time complexity of the registration phase.

Before we provide the full specification of the protocol and the corresponding part of $\texttt{VSCC}_i(\mathcal{C})$, let us briefly explain one additional technicality. Recall that in case $\gamma.\mathsf{Ingrid}$ is corrupt, it can happen that the contract instances with code $\texttt{VSCC}_i(\mathcal{C})$ are opened in the subchannels of γ although the virtual state channel γ was is not successfully created. This in particular means that the coins needed to create γ are locked in the subchannels and can be unlocked only after round $\gamma.\mathsf{validity}$ by executing the contact function $\texttt{Close}_i^{\mathcal{C}}$.

Protocol $\Pi(i, \mathcal{C})$: Close a virtual state channel

Let γ the the virtual state channel requested to be created in round τ_0 and let $cid_T := T||\gamma.\mathsf{id}$ and $id_T := \gamma.\mathsf{subchan}(T)$. We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, we denote $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$, $\mathrm{TV} := \gamma.\mathsf{validity}$,$\mathrm{TE}_{sub} := \mathrm{TimeExe}(\lceil i/2 \rceil)$ and $\mathrm{TER}_{sub} := \mathrm{TimeExeReq}(\lceil i/2 \rceil)$.

$\boxed{\text{Party } T \in \gamma.\mathsf{end\text{–}users} \text{ in round TV}}$

1. If the virtual channel was not created, i.e. $\mathit{\Gamma}^T(\gamma.\mathsf{id}) = \perp$ but $(\mathrm{updated}, id_T, cid_T) \xhookleftarrow{\leq \tau_0 + 2 + 4 \cdot \mathrm{TER}_{sub}} \mathcal{F}_{ch}$ was received, then goto step 3.
2. If $\gamma^T := \mathit{\Gamma}^T(\gamma.\mathsf{id}) \neq \perp$, then for $cid \in \{0,1\}^*$ such that $\gamma^T.\mathsf{cspace}(cid) \neq \perp$ and (id, cid) is not marked as registered in $\mathit{\Gamma}_{aux}^T$, call $\texttt{Register}_i(T, id, cid)$.
3. Send $(\mathrm{execute}, id_T, cid_T, \texttt{Close}_i^{\mathcal{C}}, \emptyset) \xhookrightarrow{\mathrm{TV} + \mathrm{TER}_{sub} + \mathrm{TE}_{sub}} \mathcal{F}_{ch}$.

$\boxed{\text{Party } I}$

For both $T \in \{A, B\}$ behave as follows:

4. If you did not receive $(\mathrm{execute\text{–}requested}, id_T, cid_T, \texttt{Close}_i^{\mathcal{C}}, \emptyset) \xhookleftarrow{\leq \mathrm{TV} + 2 \cdot \mathrm{TER}_{sub} + \mathrm{TE}_{sub}} \mathcal{F}_{ch}$, then send $(\mathrm{execute}, id_T, cid_T, \texttt{Close}_i^{\mathcal{C}}, \emptyset) \xhookrightarrow{\mathrm{TV} + 2 \cdot \mathrm{TER}_{sub} + \mathrm{TE}_{sub}} \mathcal{F}_{ch}$.

$\boxed{\text{Party } T = A, B}$

5. Upon $(\mathrm{executed}, id_T, cid_T, \sigma_T, L_T, R_T, m_T) \xhookleftarrow{\tau \leq \mathrm{TV} + 2 \cdot \mathrm{TER}_{sub} + 2 \cdot \mathrm{TE}_{sub}} \mathcal{F}_{ch}$, where $m_T = (\mathrm{contract\text{–}closed})$, set $\gamma^T(id) := \perp$ and output $(\mathrm{closed}, id) \xhookrightarrow{\tau} \mathcal{Z}$.

Contract $\texttt{VSCC}_i(\mathcal{C})$: function $\texttt{Close}_i^{\mathcal{C}}(\sigma, P, \tau)$

Let $L := \sigma.\mathsf{user}_L$, $R := \sigma.\mathsf{user}_R$, $\gamma := \sigma.\mathsf{virtual\text{–}channel}$, $A := \gamma.\mathsf{Alice}$, $B := \gamma.\mathsf{Bob}$, $I := \gamma.\mathsf{Ingrid}$.

1. Make the following checks: $\gamma \neq \perp$; $P \in \{L, R\}$; if $P \in \{A, B\}$, then $\tau < \gamma.\mathsf{validity} + \mathrm{TER}_{sub} + \mathrm{TE}_{sub}$; if $P = I$, then $\tau < \gamma.\mathsf{validity} + 2 \cdot \mathrm{TER}_{sub} + \mathrm{TE}_{sub}$. If one of the checks fails, the output $(\sigma, 0, 0, \perp)$.
2. Let $\sigma^{(0)} := \sigma$.

3. If $\sigma.\mathsf{cspace}(\mathit{cid}) = \bot$ for every $\mathit{cid} \in \{0,1\}^*$ and $\sigma.\mathsf{preRegistered} \neq \bot$, then parse $(T, \tau^T; \mathit{cid}^T, \nu^T) := \sigma.\mathsf{preRegistered}$ and define $\sigma^{(0)}.\mathsf{cspace}(\mathit{cid}^T) := \nu^T$.
4. If $\sigma^{(0)}.\mathsf{cspace}(\mathit{cid}) \neq \bot$ for some $\mathit{cid} \in \{0,1\}^*$, let us denote this identifier cid^*, then let $E \subseteq \sigma^{(0)}.\mathsf{toExecute}$ consist of all tuples $(\tau', \mathit{cid}^*, T_n, \tau_n, f_n, z_n)$, where f_n is a contract function with respect to $\mathtt{C} := \sigma^{(0)}.\mathsf{cspace}(\mathit{cid}^*).\mathsf{code}$.
5. Let $|E| = \ell$ and $(e^{(1)}, \ldots, e^{(\ell)})$ be such that $e^{(k)} = (\tau^{(k)}, T_n^{(k)}, \tau_n^{(k)}, f_n^{(k)}, z_n^{(k)}) \in E$ for every $k \in [1, \ell]$, $\tau_n^{(1)} \leq \cdots \leq \tau_n^{(\ell)}$ and if $\tau_n^{(i)} = \tau_n^{(j)}$ for some $i < j$, then the following holds:
   - If $T^{(i)} \neq T^{(j)}$, then $T^{(i)} = A$ and $T^{(j)} = B$.
   - If $T^{(i)} = T^{(j)}$, then either $f_i <_{\mathtt{C}} f_j$, where $<_{\mathtt{C}}$ is total ordering of the contract functions defined by the contract code $\mathtt{C}$, or $f_i = f_j$ and $z_n^{(i)} \leq_{\text{lex}} z_n^{(j)}$, where $\leq_{\text{lex}}$ is the lexicographic ordering of binary strings.
6. For $k = 1$ to ℓ do the following: Compute $\sigma^{(k)} := \mathtt{Evaluate}(\sigma^{(k-1)}, \mathit{cid}^*, T_n^{(k)}, f_n^{(k)}, z_n^{(k)})$ and delete $e^{(k)}$ from $\sigma^{(k)}.\mathsf{toExecute}$.
7. Set $\tilde{\sigma} := \sigma^{(\ell)}$. Let $\mathit{invest}_L := \gamma.\mathsf{cash}(A)$, $\mathit{invest}_R := \gamma.\mathsf{cash}(B)$ denote the balance when the contract was opened and let $\mathit{final}_L := \tilde{\sigma}.\mathsf{cash}(L)$ and $\mathit{final}_R := \tilde{\sigma}.\mathsf{cash}(R)$ denote the current balance. Distinguish the following two situations:
   - If $X := (\mathit{invest}_L - \mathit{final}_L) + (\mathit{invest}_R - \mathit{final}_R) \geq 0$, then set $\tilde{\sigma}.\mathsf{cash}(L) := (\mathit{invest}_L - \mathit{final}_L)$ and $\mathit{add}_L := \mathit{final}_L$. Analogously for $\tilde{\sigma}.\mathsf{cash}(R)$ and add_R. If $X < 0$, then set $\tilde{\sigma}.\mathsf{cash}(I) := \tilde{\sigma}.\mathsf{cash}(I) + X$. In addition, add X coins to add_L if $I = L$ and to add_R if $I = R$.
   - Otherwise set $\tilde{\sigma}.\mathsf{cash}(L) := 0$, $\tilde{\sigma}.\mathsf{cash}(R) := 0$ and $(\mathit{add}_L, \mathit{add}_R) := (\mathit{invest}_L, \mathit{invest}_R)$.
8. Set $\tilde{\sigma}.\mathsf{locked} := 0$, $\tilde{\sigma}.\mathsf{virtual\text{–}channel} := \bot$ and output $(\tilde{\sigma}, \mathit{add}_L, \mathit{add}_R, m)$, where $m =$ (contract–closed).

Auxiliary procedure: $\mathtt{Evaluate}(\sigma, \mathit{cid}, P_n, \tau_n, f_n, z_n)$

Let $\gamma := \sigma.\mathsf{virtual\text{–}channel}, I := \gamma.\mathsf{Ingrid}$, $\nu := \sigma.\mathsf{cspace}(\mathit{cid})$, $\sigma_n := \nu.\mathsf{storage}$ and $P := \gamma.\mathsf{end\text{–}users} \cap \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\}$

1. Compute $(\tilde{\sigma}_n, \mathit{add}_L, \mathit{add}_R, m_n) = f_n(\sigma_n, P_n, \tau_n, z_n)$.
2. If $m_n = \bot$, then output σ. Otherwise let $\tilde{\sigma} := \sigma$ and make the following changes:
   (a) Set $\tilde{\sigma}.\mathsf{cspace}(\mathit{cid}) := (\tilde{\sigma}_n, \nu.\mathsf{code})$
   (b) If $P = \sigma_n.\mathsf{user}_L$, then $\tilde{\sigma}.\mathsf{cash}(P) := \sigma.\mathsf{cash}(P) + \mathit{add}_L$ and $\tilde{\sigma}.\mathsf{cash}(I) := \sigma.\mathsf{cash}(I) + \mathit{add}_R$.
   (c) If $P = \sigma_n.\mathsf{user}_R$, then $\tilde{\sigma}.\mathsf{cash}(P) := \sigma.\mathsf{cash}(P) + \mathit{add}_R$ and $\tilde{\sigma}.\mathsf{cash}(I) := \sigma.\mathsf{cash}(I) + \mathit{add}_L$.
   Then output $\tilde{\sigma}$.

7.3 Time complexity

Let us summarize the time complexity of our virtual state channel construction and formally define the timing functions $\text{TimeReg}(i, \Delta)$, $\text{TimeExeReq}(i, \Delta)$ and $\text{TimeExe}(i, \Delta)$ for $i > 1$.

The optimistic time complexity of updating and executing a contract instance in a virtual state channel is the same as in case of ledger state channels, i.e. 2 reps. 5 rounds. However, let us emphasize that in case of virtual state channels also the optimistic time complexity of channel creation is *independent of the channel length* since it take 3 rounds for any virtual state channel length $i > 1$.

The pessimistic time complexities of the protocol $\Pi(i, \mathcal{C})$ for a virtual state channel of length $i > 1$ can be expressed in terms of the time complexities to execute its subchannels (which are state channels of length $\lceil i/2 \rceil$), using recursively Eq.(1). Since we know that $\text{TimeExeReq}(1, \Delta) = 5 + 4\Delta$ we can solve the recurrence and obtain

$$\begin{aligned} \text{TimeExeReq}(i, \Delta) &:= 5 + 2 \cdot \text{TimeExeReq}(\lceil i/2 \rceil, \Delta) \\ &\leq 2i \cdot (10 + 4\Delta) - 5 = O(\Delta \cdot i) \,. \end{aligned} \tag{2}$$

Registering a contact instance in a virtual state channel of length i takes at most $\mathrm{TimeReg}(i, \mathit{\Delta}) := 4 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil, \mathit{\Delta})$ rounds. Updating a contract instance in a virtual state channel of length i is upper bounded by $2 + 4 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil, \mathit{\Delta})$. Contract instance execution takes at most $\mathrm{TimeExe}(i, \mathit{\Delta}) := 5 + 4 \cdot \mathrm{TimeExeReq}(\lceil i/2 \rceil, \mathit{\Delta})$. In addition, we have the guarantee that the virtual state channel will be closed before round $\gamma.\mathsf{validity} + 2\mathrm{TimeExeReq}(\lceil i/2 \rceil, \mathit{\Delta}) + 2\mathrm{TimeExe}(\lceil i/2 \rceil, \mathit{\Delta})$.

8 Conclusion

We showed how to build general state channel networks, i.e., state channels of arbitrary length in which arbitrary contracts can be opened and executed off-chain. Our modular approach allows for a recursive construction of state channels (i.e. a virtual channel of length i is build on top of *two* state channels of length $\lceil i/2 \rceil$) which significantly simplifies the description of our construction. All protocols were proven to be secure in the global UC model and their optimistic time complexity is independent of the channel length. In the pessimistic case when malicious parties try to delay the protocol execution as much as possible, the time complexity of our construction is linear in channel length. We did not aim to optimize the pessimistic time complexity of our protocols since this would make their description even more complex. More fine grained timing analysis, which would reduce the constants in the pessimistic time complexity, and corresponding optimization of our state channel protocol would be highly recommended before the implementation. Another question is whether virtual state channels with time complexity independent of the channel length could be designed (for example using techniques from [22]).

Incentivizing intermediaries An important practical question is why would a party want to become an intermediary of a virtual state channel. Although our construction does guarantee that an honest intermediary will never lose coins, the fact that an intermediary has to lock coins for the entire lifetime of the virtual channel makes this role unattractive. This problem can be solved by adding the concept of *service fees* to our construction. Let us sketch how this could be done: both Alice and Bob would lock some additional coins in the $\mathtt{VSCC}_i$ contract instance each of them opens in their channel with Ingrid during the virtual state channel creation. More precisely, in order to create a virtual state channel γ, Alice would lock $\gamma.\mathsf{cash}(A) + \mathsf{serviceFee}$ coins in the channel α she has with Ingrid and Bob would lock $\gamma.\mathsf{cash}(B) + \mathsf{serviceFee}$ coins in the channel β he has with Ingrid. During the closure of γ (assuming that it was successfully created), the service fee would be unlocked from the $\mathtt{VSCC}$ contract instances in favor of Ingrid in both channels α and β.

Suitable contract codes – a cautionary note We would like to point out one subtle issue, that users of future real-life implementations need to be aware of. As discussed in Sec. 3.2, the security guarantees provided to the end-users of a state channel are strongly dependent on the code of the contract instance that is opened in the state channel (in other words: our system is only as secure as the contract that the user run in the channel). In principle, this is the same as in case of the standard contracts on the ledger, however there are several additional aspects that have to be taken into account when designing contract codes for state channels. Recall that all coins that are locked in a contract instance when the underlying virtual state channel is closed are assigned to the intermediary of the channel. Therefore, it is important that a contract instance is terminated by any end-user before the validity of the underlying virtual state channel expires. Another important point to keep in mind is that although our construction guarantees that end-user of a state channel can execute a contract instance in any round and on any contract function, it might take (in the pessimistic case) up to $\mathrm{TimeExeReq}(i)$ rounds before the other party is notified about the execution and $\mathrm{TimeExe}(i)$ rounds before the execution takes place (where i denotes the length of the state channel). Thus, compared to the contract deployment directly on the blockchain, the notification and execution delay might be longer.

Acknowledgments

This work has been supported by the Polish National Science Centre grant 2014/13/B/ST6/03540 and by the German Research Foundation (DFG) as part of project S7 within the CRC 1119 CROSSING and by the Emmy Noether Program FA 1320/1-1.

We thank Jeff Coleman for several useful comments and in particular for pointing out a weakness of an earlier version of our protocol when taking fees into account.

References

[1] Ian Allison. *Ethereum's Vitalik Buterin explains how state channels address privacy and scalability.* 2016.

[2] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. "Secure Multiparty Computations on Bitcoin". In: *2014 IEEE Symposium on Security and Privacy.* Berkeley, CA, USA: IEEE Computer Society Press, 2014, pp. 443–458. DOI: `10.1109/SP.2014.35`.

[3] Iddo Bentov and Ranjit Kumaresan. "How to Use Bitcoin to Design Fair Protocols". In: *Advances in Cryptology – CRYPTO 2014, Part II.* Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8617. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2014, pp. 421–439. DOI: `10.1007/978-3-662-44381-1_24`.

[4] Iddo Bentov, Ranjit Kumaresan, and Andrew Miller. "Instantaneous Decentralized Poker". In: *Advances in Cryptology – ASIACRYPT 2017.* Ed. by Tsuyoshi Takagi and Thomas Peyrin. Cham: Springer International Publishing, 2017, pp. 410–440. ISBN: 978-3-319-70697-9.

[5] *Bitcoin Wiki: Payment Channels.* https://en.bitcoin.it/wiki/Payment_channels. 2018.

[6] *Bitcoin Wiki: Scalability.* `https://en.bitcoin.it/wiki/Nanopayments`. 2018.

[7] Ran Canetti. "Universally Composable Security: A New Paradigm for Cryptographic Protocols". In: *42nd Annual Symposium on Foundations of Computer Science.* Las Vegas, NV, USA: IEEE Computer Society Press, 2001, pp. 136–145.

[8] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. "Universally Composable Security with Global Setup". In: *TCC 2007: 4th Theory of Cryptography Conference.* Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, 2007, pp. 61–85.

[9] *Counterfactual.* `https://counterfactual.com/`. 2018.

[10] Christian Decker and Roger Wattenhofer. "A Fast and Scalable Payment Network with Bitcoin Duplex Micropayment Channels". In: *Stabilization, Safety, and Security of Distributed Systems: 17th International Symposium, SSS 2015, Edmonton, AB, Canada, August 18-21, 2015, Proceedings.* Ed. by Andrzej Pelc and Alexander A. Schwarzmann. Cham: Springer International Publishing, 2015, pp. 3–18. ISBN: 978-3-319-21741-3. DOI: `10.1007/978-3-319-21741-3_1`. URL: `http://dx.doi.org/10.1007/978-3-319-21741-3_1`.

[11] Stefan Dziembowski, Lisa Eckey, Sebastian Faust, and Daniel Malinowski. *Perun: Virtual Payment Hubs over Cryptographic Currencies.* conference version accepted to the 40th IEEE Symposium on Security and Privacy (IEEE S&P) 2019. 2017. URL: `http://eprint.iacr.org/2017/635`.

[12] Oded Goldreich. *Foundations of Cryptography: Volume 1.* New York, NY, USA: Cambridge University Press, 2006. ISBN: 0521035368.

[13] Dennis Hofheinz and Joern Mueller-Quade. *A Synchronous Model for Multi-Party Computation and the Incompleteness of Oblivious Transfer.* Cryptology ePrint Archive, Report 2004/016. `http://eprint.iacr.org/2004/016`. 2004.

[14] Yael Tauman Kalai, Yehuda Lindell, and Manoj Prabhakaran. "Concurrent Composition of Secure Protocols in the Timing Model". In: *Journal of Cryptology* 20.4 (Oct. 2007), pp. 431–492.

[15] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography (Chapman & Hall/Crc Cryptography and Network Security Series).* Chapman & Hall/CRC, 2007. ISBN: 1584885513.

[16] Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. “Universally Composable Synchronous Computation”. In: *TCC 2013: 10th Theory of Cryptography Conference*. Ed. by Amit Sahai. Vol. 7785. Lecture Notes in Computer Science. Tokyo, Japan: Springer, Heidelberg, Germany, 2013, pp. 477–498. DOI: `10.1007/978-3-642-36594-2_27`.

[17] Rami Khalil and Arthur Gervais. “Revive: Rebalancing Off-Blockchain Payment Networks”. In: *ACM CCS 17: 24th Conference on Computer and Communications Security*. Ed. by Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu. Dallas, TX, USA: ACM Press, 2017, pp. 439–453.

[18] Joshua Lind, Ittay Eyal, Florian Kelbert, Oded Naor, Peter R. Pietzuch, and Emin Gün Sirer. “Teechain: Scalable Blockchain Payments using Trusted Execution Environments”. In: *CoRR* abs/1707.05454 (2017). arXiv: `1707.05454`. URL: `http://arxiv.org/abs/1707.05454`.

[19] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, Matteo Maffei, and Srivatsan Ravi. “Concurrency and Privacy with Payment-Channel Networks”. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*. 2017, pp. 455–471.

[20] Patrick McCorry, Surya Bakshi, Iddo Bentov, Andrew Miller, and Sarah Meiklejohn. “Pisa: Arbitration Outsourcing for State Channels”. In: *IACR Cryptology ePrint Archive* 2018 (2018), p. 582. URL: `https://eprint.iacr.org/2018/582`.

[21] Silvio Micali and Ronald L. Rivest. “Micropayments Revisited”. In: *Topics in Cryptology – CT-RSA 2002*. Ed. by Bart Preneel. Vol. 2271. Lecture Notes in Computer Science. San Jose, CA, USA: Springer, Heidelberg, Germany, 2002, pp. 149–163.

[22] Andrew Miller, Iddo Bentov, Ranjit Kumaresan, and Patrick McCorry. “Sprites: Payment Channels that Go Faster than Lightning”. In: *CoRR* abs/1702.05812 (2017). URL: `http://arxiv.org/abs/1702.05812`.

[23] Satoshi Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. `http://bitcoin.org/bitcoin.pdf`. 2009.

[24] Jesper Buus Nielsen. “On Protocol Security in the Cryptographic Model”. PhD thesis. Aarhus University, 2003.

[25] Olaoluwa Osuntokun. *Hardening Lightning*. BPASE. 2018. URL: `https://cyber.stanford.edu/sites/default/files/hardening_lightning_updated.pdf`.

[26] Rafael Pass and Abhi Shelat. “Micropayments for Decentralized Currencies”. In: *ACM CCS 15: 22nd Conference on Computer and Communications Security*. Ed. by Indrajit Ray, Ninghui Li, and Christopher Kruegel: Denver, CO, USA: ACM Press, 2015, pp. 207–218.

[27] Joseph Poon and Thaddeus Dryja. *The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments*. Draft version 0.5.9.2, available at `https://lightning.network/lightning-network-paper.pdf`. Jan. 2016.

[28] Ronald L. Rivest. “Electronic Lottery Tickets as Micropayments”. In: *FC’97: 1st International Conference on Financial Cryptography*. Ed. by Rafael Hirschfeld. Vol. 1318. Lecture Notes in Computer Science. Anguilla, British West Indies: Springer, Heidelberg, Germany, 1997, pp. 307–314.

[29] Stefanie Roos, Pedro Moreno-Sanchez, Aniket Kate, and Ian Goldberg. “Settling Payments Fast and Private: Efficient Decentralized Routing for Path-Based Transactions”. In: *CoRR* abs/1709.05748 (2017). arXiv: `1709.05748`. URL: `http://arxiv.org/abs/1709.05748`.

[30] David Siegel. *Understanding The DAO Attack*. CoinDesk, `http://www.coindesk.com/understanding-dao-hack-journalists/`. 2016.

[31] *Update from the Raiden team on development progress, announcement of raidEX*. `https://tinyurl.com/z2snp9e`. Feb. 2017.

[32] David Wheeler. “Transactions Using Bets”. In: *Proceedings of the International Workshop on Security Protocols*. London, UK, UK: Springer-Verlag, 1997, pp. 89–92. ISBN: 3-540-62494-5. URL: `http://dl.acm.org/citation.cfm?id=647214.720381`.

[33] Gavin Wood. *ETHEREUM: A SECURE DECENTRALISED GENERALISED TRANSACTION LEDGER*. `http://gavwood.com/paper.pdf`. 2014.

A Routing payments using hash-locked transactions

Consider the situation when Alice has a payment channel with Ingrid and Ingrid has a payment channel with Bob. Assume that Alice wants to send one coin to Bob and route the payment via Ingrid. The first idea would be to let Alice update the channel with Ingrid such that Alice pays one coin to Ingrid and then let Ingrid symmetrically update the channel with Bob such that Ingrid pays one coin to Bob. However, this naive solution allows a malicious Ingrid to abort after receiving the coin from Alice and never pay anything to Bob.

Let us briefly explain how to solve the above problem using *hash-locked transactions*. Let H be some fixed hash function. Bob first picks a random value $x \in \{0,1\}^*$ and sends the hash value $h = H(x)$ to Alice who creates a hash-locked transaction HLT_A. Informally, this transaction promises to update the channel between Alice and Ingrid such that Ingrid earns one coin if she publishes a preimage of h before a timeout t_A. Ingrid, upon receiving the hash-locked transaction HLT_A from Alice, creates a hash-locked transaction HLT_B which promises to update the channel between Ingrid and Bob such that Bob earns one coin if he publishes a preimage of h before the timeout $t_B < t_A$. Hence, if Bob reveals x before time t_B, he gets one coin from Ingrid. Since $t_B < t_A$, Ingrid has time to use the value x to get one coin from Alice and thus finalize the payment. In case Bob does not reveal x to Ingrid before the timeout t_B, Ingrid can refund her coin locked in HLT_B. Analogously, in case Ingrid does not reveal x, Alice can refund her coin locked in HLT_A after round t_A.

B Restrictions on the Environment

In order to simplify the description of the state channel ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{C}}(\Delta)}(i,\mathcal{C})$ and the protocol $\Pi(i,\mathcal{C})$ realizing it, we define a set of restricted environments $\mathcal{E}_{res}$ by the restrictions give below (some of the restrictions were already informally introduced in Sec. 4). Every environment $\mathcal{Z} \in \mathcal{E}_{res}$ has to satisfy the following:

- $\mathcal{Z}$ never sends the same message to the same party twice.
- $\mathcal{Z}$ sends a message (create, γ), where γ is a ledger state channel, to all honest parties in the set $\gamma.\mathsf{end\text{–}users}$ in the same round τ_0 (and it never send this message to any other honest party). In addition, we assume the following: there does not exist a state channel γ' with $\gamma.\mathsf{id} = \gamma'.\mathsf{id}$ (and no such state channel is currently being created); parties of the ledger state channel are from the set $\mathcal{P}$; $\gamma.\mathsf{cash}(A) \geq 0, \gamma.\mathsf{cash}(B) \geq 0$; both parties of the ledger state channel have enough funds on the ledger for the channel creation;[21] the set of contract instances is empty; and $\gamma.\mathsf{length} = 1$. In addition, we assume that if A is honest and the environment does not receive the message $(\text{created}, \gamma)$ from A within 2Δ rounds, it sends the message (refund, γ) to party A.
- $\mathcal{Z}$ sends the message (create, γ), where γ is a virtual state channel , to all honest parties in the set $\gamma.\mathsf{end\text{–}users} \cup \{I\}$ in the same round τ_0 (and it never send this message to any other honest party). In addition, we assume the following: there does not exists a state channel γ' with $\gamma.\mathsf{id} = \gamma'.\mathsf{id}$ (and no such state channel is currently being created); parties of the virtual state channel are from the set $\mathcal{P}$; $\gamma.\mathsf{cash}(A) \geq 0, \gamma.\mathsf{cash}(B) \geq 0$; the set of contract instances is empty; $\gamma.\mathsf{validity} < \tau_0 + 2 + 4 \cdot \mathrm{TimeExe\ Req}(\lceil i/2 \rceil)$; $j := \gamma.\mathsf{length} \leq i$. Additionally, we assume the following about the subchannels of γ:
  - If honest $P \in \gamma.\mathsf{end\text{–}users}$ receives the message (create, γ), then the following must be satisfied: the subchannel $\alpha := \gamma.\mathsf{subchan}(P)$ must exist; $\alpha.\mathsf{end\text{–}users} = \{P, I\}$; $\alpha.\mathsf{length} \leq \lceil j/2 \rceil$; $\gamma.\mathsf{validity} > \alpha.\mathsf{validity} + 2\mathrm{TimeExeReq}(\lceil j/2 \rceil) + 2\mathrm{TimeExe}(\lceil j/2 \rceil)$; $\alpha.\mathsf{cspace}(cid) = \perp$ for every $cid \in \{0,1\}^*$ if α is a virtual state channel; both P and I have enough funds in α.
  - If honest I receives the message (create, γ), then both subchannels $\alpha := \gamma.\mathsf{subchan}(A)$ and $\beta := \gamma.\mathsf{subchan}(B)$ exist; $\alpha.\mathsf{end\text{–}users} = \{A, I\}$ and $\beta.\mathsf{end\text{–}users} = \{B, I\}$; $j = \alpha.\mathsf{length} + \beta.\mathsf{length}$, $\alpha.\mathsf{length} \leq \lceil j/2 \rceil$ and $\beta.\mathsf{length} \leq \lceil j/2 \rceil$; $\gamma.\mathsf{validity} > \max\{\alpha.\mathsf{validity}, \beta.\mathsf{validity}\} + 2\mathrm{TimeExeReq}(\lceil j/2 \rceil) + 2\mathrm{Time}$

[21] In case the environment requests opening more ledger state channels at the same time, we require that all parties have enough funds for all ledger state channels that are being created.

Exe($\lceil j/2 \rceil$); $\alpha.\mathsf{cspace}(cid) = \perp$ for every $cid \in \{0,1\}^*$ if α is a virtual state channel; similarly $\beta.\mathsf{cspace}(cid) = \perp$ for every $cid \in \{0,1\}^*$ if β is a virtual state channel; A and I have enough funds in α and B and I have enough funds in β.

- If $\mathcal{Z}$ sends the message (update, $id, cid, \tilde{\sigma}, \mathsf{C}$) or (update–reply, ok, id, cid) to an honest party P, then a state channel γ with identifier id exists in Γ; $P \in \gamma.\mathsf{end\text{–}users}$; the state channel supports the contract code ; the new contract instance $\tilde{\sigma}$ is admissible with respect to C, i.e. $\tilde{\sigma} \in \mathsf{C}.\Lambda$; it holds that $\tilde{\sigma}.\mathsf{locked} = \tilde{\sigma}.\mathsf{cash}(\tilde{\sigma}.\mathsf{user}_L) + \tilde{\sigma}.\mathsf{cash}(\tilde{\sigma}.\mathsf{user}_R)$ and both parties have enough cash in the state channel for the contract instance update.[22] If the contract instance has been updated before, i.e., if $\nu := \gamma.\mathsf{cspace}(cid) \neq \perp$, then the following must hold: the contract instance code remains the same, i.e., $\nu.\mathsf{code} = \mathsf{C}$; the users of the contract instance remain the same, i.e., for $\sigma := \nu.\mathsf{storage}$ we have $\sigma.\mathsf{user}_L = \tilde{\sigma}.\mathsf{user}_L$ and $\sigma.\mathsf{user}_R = \tilde{\sigma}.\mathsf{user}_R$; and $\sigma.\mathsf{locked} = \sigma.\mathsf{cash}(\sigma.\mathsf{user}_L) + \sigma.\mathsf{cash}(\sigma.\mathsf{user}_R)$. $\mathcal{Z}$ never asks to update a contract instance that is currently being updated or executed. In addition, if $\Gamma(id)$ is a virtual state channel, then we assume that there is no other contract instance in the virtual state channel (and no other instance is being created) and the message was send before round $\gamma.\mathsf{validity}$.
- If $\mathcal{Z}$ sends the message (execute, id, cid, f, z) to an honest party P, then a state channel γ with identifier id exists in Γ, $P \in \gamma.\mathsf{end\text{–}users}$, the contract instance cid has already been defined in γ, i.e. $\gamma.\mathsf{cspace}(cid) \neq \perp$, and f is a contract function with respect to $\gamma.\mathsf{cspace}(cid).\mathsf{code}$, i.e. $f \in \gamma.\mathsf{cspace}(cid).\mathsf{code}$. If γ is a virtual state channel, then we assume that the message is sent before round $\gamma.\mathsf{validity}$.
- If $\mathcal{Z}$ sends the message (close, id) to honest party P, then state channel γ with identifier id exists in Γ, γ is a ledger state channel and $P \in \gamma.\mathsf{end\text{–}users}$.

C Balance security

Let us prove that the state channels ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ satisfies the balance security property. The remaining security and efficiency goals that were defined in Sec. 3.2 follow directly from the formal definition of the ideal functionality given in Figure 6 and were already discussed in Sec. 4.

Let γ be a virtual state channel of length j requested to be created in round τ_0. First note that if creation of a virtual state channel γ fails (at least one of the parties does not confirm the creation within three rounds), then all coins locked in the subchannels in Step 1 (see virtual state channel creation Figure 6) are unlocked back to the subchannels latest in round $\gamma.\mathsf{validity} + 2\mathrm{TimeExeReq}(\lceil j/2 \rceil) + 2\mathrm{TimeExe}(\lceil j/2 \rceil)$ (see Step 3 of virtual state channel creation in Figure 6).

In order to argue about the balance security in case a virtual state channel γ is successfully created, let us state an auxiliary lemma. Intuitively, it says that the amount of coins in γ does not change during the lifetime of the channel. Before we state the lemma, let us recall that we do not allow a virtual state channel γ to contain multiple contract instances, i.e. there exists *at most one* $cid \in \{0,1\}^*$ such that $\gamma.\mathsf{cspace}(cid) \neq \perp$.

Lemma 1. *Let γ be a virtual state channel of length $j > 1$ created in round τ_1 and let X be the amount of coins initially locked in the virtual state channel, i.e. $X := \gamma.\mathsf{cash}(A) + \gamma.\mathsf{cash}(B)$. Let $\hat{\gamma}$ be the version of the virtual state channel in round $\tau \in [\tau_0, \gamma.\mathsf{validity} + 2\mathrm{TimeExeReq}(\lceil j/2 \rceil) + 2\mathrm{TimeExe}(\lceil j/2 \rceil)]$. If there exists $cid \in \{0,1\}^*$ such that $\sigma_{cid}.\mathsf{locked} \neq 0$ for $\sigma_{cid} := \hat{\gamma}.\mathsf{cspace}(cid).\mathsf{storage}$, then set $\hat{c} := \sigma_{cid}.\mathsf{locked}$; else set $\hat{c} := 0$. It holds that $X = \hat{\gamma}.\mathsf{cash}(A) + \hat{\gamma}.\mathsf{cash}(B) + \hat{c}$.*

Proof. Let γ_1 be a successfully created virtual state channel in round τ_1 and let τ_2 be the first round in which the state channel is successfully updated. First note that for $\tau \in (\tau_1, \tau_2)$ the lemma holds trivially since the channel did not change. Let γ_2 be the channel after the update and let us denote $\sigma_2 := \gamma_2.\mathsf{cspace}(cid).\mathsf{storage}$. By definition of the ideal functionality and the auxiliary procedure `UpdateChanSpace` it holds that $\gamma_2.\mathsf{cash}(A) = \gamma_1.\mathsf{cash}(A) - \sigma_2.\mathsf{cash}(A)$ and $\gamma_2.\mathsf{cash}(B) = \gamma_1.\mathsf{cash}(B) - \sigma_2.\mathsf{cash}(B)$. By the

[22] In case the environment requests constructing more contract instances at the same time, we require that both parties have enough funds in the state channel for all of them.

restrictions on the environment (see Appx. B), we know that $\sigma_2.\mathsf{locked} = \sigma_2.\mathsf{cash}(A) + \sigma_2.\mathsf{cash}(B)$. Thus

$$\begin{aligned}
&\gamma_1.\mathsf{cash}(A) + \gamma_1.\mathsf{cash}(B) = \\
&\qquad = \gamma_2.\mathsf{cash}(A) + \sigma_2.\mathsf{cash}(A) + \gamma_2.\mathsf{cash}(B) + \sigma_2.\mathsf{cash}(B) \\
&\qquad = \gamma_2.\mathsf{cash}(A) + \gamma_2.\mathsf{cash}(B) + \sigma_2.\mathsf{locked}
\end{aligned}$$

which is exactly what we needed to prove.

By the restrictions on the environment (see Appx. B), the only way how the channel can change in round $\tau \in (\tau_2, \gamma_1.\mathsf{validity} + 2\mathrm{TimeExeReq}(\lceil j/2 \rceil) + 2\mathrm{TimeExe}(\lceil j/2 \rceil)]$ is via another successful update or a successful execution the contract instance cid. Since the argument in case of another successful update of the contract instance cid is very similar to the one above, let us discuss only the case of successful contract instance execution execution in round τ.

Let γ_3 be the channel before the execution of a contract function and let γ_4 be the channel after the execution. In addition, let $\sigma_3 := \gamma_3.\mathsf{cspace}(cid).\mathsf{storage}$ and let $\sigma_4 := \gamma_4.\mathsf{cspace}(cid).\mathsf{storage}$. By definition of a contract function, we know that the output of the function contains two value add_A and add_B which are such that $\sigma_3.\mathsf{locked} - \sigma_4.\mathsf{locked} = add_A + add_B$. By definition of the auxiliary procedure `UpdateChanSpace` called by the ideal functionality on inputs add_A and add_B, we know that $\gamma_4.\mathsf{cash}(A) = \gamma_3.\mathsf{cash}(A) + add_A$ and $\gamma_4.\mathsf{cash}(B) = \gamma_3.\mathsf{cash}(B) + add_B$. Thus,

$$\begin{aligned}
&\gamma_3.\mathsf{cash}(A) + \gamma_3.\mathsf{cash}(B) + \sigma_3.\mathsf{locked} = \\
&\qquad = \gamma_4.\mathsf{cash}(A) - add_A + \gamma_4.\mathsf{cash}(B) - add_B + \sigma_3.\mathsf{locked} \\
&\qquad = \gamma_4.\mathsf{cash}(A) + \gamma_4.\mathsf{cash}(B) + \sigma_4.\mathsf{locked}
\end{aligned}$$

which is exactly what we needed to prove.

The balance security now easily follows. Let $\hat{\gamma}$ be the virtual state channel when the channel is being closed. If there is no contract instance with locked coins in $\hat{\gamma}$, then the intermediary gets $\hat{\gamma}.\mathsf{cash}(A) + \hat{\gamma}.\mathsf{cash}(B)$ coins which by Lemma 1 is equal to the amount of coins the intermediary had to lock in Step 1. If there is a contract instance with $\hat{c}$ locked coins in $\hat{\gamma}$, then the intermediary gets $\hat{\gamma}.\mathsf{cash}(A) + \hat{\gamma}.\mathsf{cash}(B) + 2\hat{c}$ coins which is by Lemma 1 more than what the intermediary initially locked in Step 1 (concretely, the intermediary gains $\hat{c}$ coins in this case).

D The Lottery contract

The contract code $\mathtt{C_{lot}}(i)$ allows parties to play a lottery in a state channel of length up to i. The contract $\mathtt{C_{lot}}(i)$ was informally described in Sec. 4.3, here we present it formally (see Fig. 12).

Lottery contract $\texttt{C}_{\texttt{lot}}(i)$

Constructor $\texttt{Init}_{\texttt{lot}}(P, \tau, (A, B))$

Output a contract storage σ defined as follows:
$\sigma.\mathsf{user}_L := A$, $\sigma.\mathsf{user}_R := B$, $\sigma.\mathsf{cash}(\sigma.\mathsf{user}_L) := 1$, $\sigma.\mathsf{cash}(\sigma.\mathsf{user}_R) := 1$, $\sigma.\mathsf{locked} := 2$, $\sigma.\mathsf{com} = \perp$ and $\sigma.\mathsf{bit} = \perp$.

Function $\texttt{Com}(\sigma, P, \tau; c)$

If $P \neq \sigma.\mathsf{user}_L$ or if $\sigma.\mathsf{com}(P) \neq \perp$, then output $(\sigma, 0, 0, \perp)$. Else set $\tilde{\sigma} := \sigma$, set $\tilde{\sigma}.\mathsf{com} := (c, \tau)$ and output $(\tilde{\sigma}, 0, 0, (\text{committed}, P, c, \tau))$.

Function $\texttt{Reveal}(\sigma, P, \tau; r_B)$

If $P \neq \sigma.\mathsf{user}_R$ or if $\sigma.\mathsf{com}(P) = \perp$, then output $(\sigma, 0, 0, \perp)$. Else set $\tilde{\sigma} := \sigma$, set $\tilde{\sigma}.\mathsf{bit} := (r_B, \tau)$ and output $(\tilde{\sigma}, 0, 0, (\text{revealed}, P, r_B, \tau))$.

Function $\texttt{Open}(\sigma, P, \tau; (r_A, s))$

If $P \neq \sigma.\mathsf{user}_L$, $\sigma.\mathsf{com} = \perp$, $\sigma.\mathsf{bit}(P) \neq \perp$ or (r_A, s) is not a valid opening of $\sigma.\mathsf{com}$, then output $(\sigma, 0, 0, \perp)$. Else parse $(r_B, \tau') := \sigma.\mathsf{bit}$ and compute $x := r_A \oplus r_B$. If $x = 0$, then set $(add_A, add_B) := (2, 0)$ else set $(add_A, add_B) := (0, 2)$. Define $\tilde{\sigma} := \sigma$, set $\tilde{\sigma}.\mathsf{locked} = 0$, $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_L) := 1 - add_L$ and $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_R) := 1 - add_R$. Then output $(\tilde{\sigma}, add_L, add_R, (\text{opened}, P, r_A, \tau))$.

Function $\texttt{Punish}(\sigma, P, \tau; z)$

If $P \in \{\sigma.\mathsf{user}_L, \sigma.\mathsf{user}_R\}$ and $\sigma.\mathsf{locked} \neq 0$, then consider the following three situations:

1. If $\sigma.\mathsf{com} = \perp$, $P = \sigma.\mathsf{user}_R$ and $\tau > \sigma.\mathsf{start} + \text{TimeExe}(i)$, then set $\tilde{\sigma} := \sigma$ and define $\tilde{\sigma}.\mathsf{locked} = 0$, $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_L) := 1$, $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_R) := -1$ and output $(\tilde{\sigma}, 0, 2, (\text{punished}))$.
2. If $\sigma.\mathsf{com} \neq \perp$, $P = \sigma.\mathsf{user}_L$, $\sigma.\mathsf{bit} = \perp$ and $\tau > \sigma.\mathsf{start} + 2 \cdot \text{TimeExe}(i)$, then set $\tilde{\sigma} := \sigma$, define $\tilde{\sigma}.\mathsf{locked} = 0$, $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_L) := -1$ and $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_R) := 1$ and output $(\tilde{\sigma}, 2, 0, (\text{punished}))$.
3. If $\sigma.\mathsf{com} \neq \perp$, $P = \sigma.\mathsf{user}_R$, $\sigma.\mathsf{bit} \neq \perp$ and $\tau > \sigma.\mathsf{start} + 3 \cdot \text{TimeExe}(i)$, then proceed as follows. Set $\tilde{\sigma} := \sigma$, define $\tilde{\sigma}.\mathsf{locked} = 0$, $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_L) := 1$ and $\tilde{\sigma}.\mathsf{cash}(\mathsf{user}_R) := -1$ and output $(\tilde{\sigma}, 0, 2, (\text{punished}))$.

Else output $(\sigma, 0, 0, \perp)$

Fig. 12: The lottery contract.

E Security analysis for ledger state channels

In this section, we will show that for any set of contract codes $\mathcal{C}$, the $\Pi(1, \mathcal{C})$ protocol $\mathcal{E}_{res}$-emulates the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ in $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$-hybrid world. In other words, for any PPT adversary Adv we construct a simulator Sim_1 that operates in the $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ world and simulates the $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$-hybrid world to any environment $\mathcal{Z} \in \mathcal{E}_{res}$.

The two main challenges of our analysis are the following: (i) ensure the consistency of timings (if an honest party P outputs a message m in round τ in the hybrid world, then P must output the same message m in the same round τ in the ideal world as well) and (ii) ensure the consistency of balances of parties on the ledger (i.e. if the state of accounts on the ledger in round τ is equal $(x_1, \ldots, x_n)$ in the hybrid world, then the state of user's accounts in round τ must be $(x_1, \ldots, x_n)$ in the ideal world as well). Recall that the ledger $\widehat{\mathcal{L}}$ is a global ideal functionality thus the environment can read its state at any point in time. Inconsistencies on the ledger could therefore reveal to the environment whether it is communicating with the real or ideal world.

The simulator Sim_1 constructed in this section will internally run a copy the hybrid world. It will maintain a channel space Γ^T and the auxiliary channel space Γ^T_{aux} for every honest party T and the channel space Γ for the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$. In addition, the simulator will generate a key pair $(pk_T, sk_T) \leftarrow_{\$} \mathsf{KGen}(1^\lambda)$ for every honest party T during the setup phase. Recall that since there are no private inputs or messages being sent, we implicitly assume that the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ on receiving a message m from party P immediately sends a message (P, m) to the simulator Sim_1 (this convention was introduced in Sec. 4). The simulator Sim_1 thus receives all the input messages of the honest parties from the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$. Since Sim_1 receives messages addressed to the adversary Adv (which it internally runs) from the environment $\mathcal{Z}$, it knows the behavior of corrupt parties in the protocol as well as the instruction given by the adversary to the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$. This, in particular, means that our simulator Sim_1 can instruct the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ to make changes on the ledger $\widehat{\mathcal{L}}$ in the same round as the adversary Adv would instruct the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ to make the changes on the ledger. To simplify the pseudocode description of the simulator Sim_1, we do not write these instructions explicitly.

We will discuss each part of the protocol separately and for each of them distinguish all possible corruption combinations: both parties are honest, only one party is honest and both parties are corrupt. We present a full description of the simulator Sim_1 for all of these cases and provide a detailed proof sketch of the ideal and hybrid world indistinguishability for the ledger state channel creation when A is honest and B is corrupt. The argumentation in the remaining cases is very similar and thus omitted from this version of the paper.

Create a ledger state channel. Let us begin with the description of the simulator Sim_1 for the ledger state channel creation. We will first discuss in detail the case when A is honest and B is corrupt (the corresponding pseudocode description of the simulator can be found below).

According to the protocol $\Pi(1, \mathcal{C})$, honest party A upon receiving the message (create, γ) from the environment $\mathcal{Z}$ sends the message $(\text{construct}, \gamma)$ to the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$. Since we assume that $\mathcal{Z} \in \mathcal{E}_{res}$, all checks made by the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ will pass. This can be verified by careful inspection of $\mathcal{E}_{res}$ definition (see page 44) and the description of the ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ for ledger state channel creation (see page 22). The hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ within Δ rounds removes coins from A's account on the ledger $\widehat{\mathcal{L}}$. The exact round is determined by the adversary Adv. The simulator Sim_1 is receiving messages from $\mathcal{Z}$ addressed to the adversary Adv; thus, it can instruct the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ to remove coins from A's account in the same round (recall our convention that these messages from the simulator to the ideal functionality are implicit in our descriptions to for better readability). After removing the coins from A's account, the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ sends the message $(\text{initializing}, \gamma)$ to party B which is exactly what the simulator Sim_1 does as well.

If B is instructed by the environment $\mathcal{Z}$ to immediately reply to the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ with the message $(\text{confirm}, \gamma)$, the ledger state channel γ will be created in the hybrid world. Therefore, the simulator Sim_1 sends the message (create, γ) to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ on behalf of B which ensures the channel creation in the ideal world as well. The simulator again instructs the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1,\mathcal{C})$ to remove coins from B's account in the same round the adversary Adv would instruct the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$. After removing coins from B's account, the hybrid ideal functionality immediately sends the message $(\text{initialized},\gamma)$ to both $\mathsf{end\text{–}users}$ which makes honest A output the message $(\text{created},\gamma)$ to the environment. Therefore, the simulator Sim_1 sends the message $(\text{initialized},\gamma)$ to B right after the coins are removed. In addition, the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ after removing coins from B's account sens the message $(\text{created},\gamma)$ to both end-users which results in the honest party A forwarding it to the environment. Thus, the content and timing of the honest party's output message to the environment is the same in both worlds. Finally, the simulator Sim_i stores the new ledger state channel γ in the channel space of the hybrid ideal functionality Γ and the channel space of the honest party Γ^A and stops.

If B is not instructed by the environment $\mathcal{Z}$ to confirm the channel creation by sending the message $(\text{confirm},\gamma)$ to the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$, the ledger state channel γ will not be created in the hybrid world. Thus, simulator Sim_1 does not send any message to $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1,\mathcal{C})$ on B's behalf in this case. By our assumption that $\mathcal{Z}\in\mathcal{E}_{res}$, the honest party A receives the message (refund,γ). In the hybrid world, A forwards this message to the hybrid ideal functionality $\mathcal{F}_{scc}^{\widehat{\mathcal{L}}(\Delta)}(\mathcal{C})$ who adds coins back to A's account on the ledger within Δ rounds. In the ideal world, A is a dummy party and thus forwards the message to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1,\mathcal{C})$. Hence, the only thing that the simulator Sim_1 has to do is to instruct the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1,\mathcal{C})$ to add the coins back to A's account in the correct round and then stop.

The pseudocode description of the simulator Sim_1 that we just defined as well as the description of Sim_1 for the remaining corruption combinations with at last one corrupted party can be found below.

Simulator Sim_1: Create a ledger state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(1,\mathcal{C})$.

Case A is honest and B is corrupt:

Upon $(A,\text{create},\gamma)\overset{\tau_0}{\hookleftarrow}\mathcal{F}_{ch}$, proceed as follows:

1. Wait until round $\tau_1\leq\tau_0+\Delta$ to send $(\text{initializing},\gamma)\overset{\tau_1}{\hookrightarrow}B$.
2. If $(\text{confirm},\gamma)\overset{\tau_1}{\hookleftarrow}B$, then send $(\text{create},\gamma)\overset{\tau_1}{\hookrightarrow}\mathcal{F}_{ch}$ on behalf of B. Send $(\text{initialized},\gamma)\xhookrightarrow{\tau_2\leq\tau_1+\Delta}B$ and set $\Gamma^A(\gamma.\mathsf{id}):=\gamma,\Gamma(\gamma.\mathsf{id}):=\gamma$ and stop.

Case A is corrupt and B is honest:

Upon $(\text{construct},\gamma)\overset{\tau_0}{\hookleftarrow}A$ proceed as follows:

1. If A does not have enough funds on the ledger, there already exists a state channel γ' such that $\gamma.\mathsf{id}=\gamma'.\mathsf{id}$ in Γ, $\gamma.\mathsf{cspace}\neq\emptyset$, or $\gamma.\mathsf{cash}(A)<0$ or $\gamma.\mathsf{cash}(B)<0$, then stop.
2. Else send $(\text{create},\gamma)\overset{\tau_0}{\hookrightarrow}\mathcal{F}_{ch}$ on behalf of A and in round $\tau_1\leq\tau_0+\Delta$ send $(\text{initializing},\gamma)\overset{\tau_1}{\hookrightarrow}A$.
3. Distinguish the following two situations:
   - If $(B,\text{create},\gamma)\overset{\tau_0}{\hookleftarrow}\mathcal{F}_{ch}$, then send $(\text{initialized},\gamma)\xhookrightarrow{\tau_0+2\Delta}A$ and set $\Gamma^B(\gamma.\mathsf{id}):=\gamma$, $\Gamma(\gamma.\mathsf{id}):=\gamma$ and stop.
   - Else wait. If $(\text{refund},\gamma)\xhookleftarrow{\tau_3>\tau_0+2\Delta}A$, then send $(\text{refund},\gamma)\overset{\tau_3}{\hookrightarrow}\mathcal{F}_{ch}$.

Case A and B are corrupt:

Upon $(\text{construct},\gamma)\overset{\tau_0}{\hookleftarrow}A$ proceed as follows:

1. If A does not have enough funds on the ledger, there already exists a state channel γ' such that $\gamma.\mathsf{id}=\gamma'.\mathsf{id}$, $\gamma.\mathsf{cspace}\neq\emptyset$ or $\gamma.\mathsf{cash}(A)<0$ or $\gamma.\mathsf{cash}(B)<0$, then stop.

2. Else send $(\text{create}, \gamma) \overset{\tau_0}{\hookrightarrow} \mathcal{F}_{ch}$ on behalf of A and in round $\tau_1 \leq \tau_0 + \Delta$ send $(\text{initializing}, \gamma) \overset{\tau_1}{\hookrightarrow}$ $\gamma.\mathsf{end\text{–}users}$.
3. Distinguish the following two situations:
   – If $(\text{confirm}, \gamma) \overset{\tau_1}{\hookleftarrow} B$ and B has sufficient funds on the ledger, then $(\text{create}, \gamma) \overset{\tau_1}{\hookrightarrow} \mathcal{F}_{ch}$ and on behalf of B and wait until round $\tau_2 \leq \tau_0 + 2\Delta$ to send $(\text{initialized}, \gamma) \overset{\tau_2}{\hookrightarrow} \gamma.\mathsf{end\text{–}users}$. Then set $\Gamma(\gamma.\mathsf{id}) := \gamma$ and stop.
   – Else wait if $(\text{refund}, \gamma) \xleftarrow{\tau_3 > \tau_0 + 2\Delta} A$. In such a case send $(\text{refund}, \gamma) \overset{\tau_3}{\hookrightarrow} \mathcal{F}_{ch}$ and stop.

What remains to discuss the case when both parties of the ledger state channel are honest. The only thing the simulator has to do is to instruct the ideal functionality to remove coins from ledger accounts in the correct round which can be done since it received the message addressed to the adversary Adv. After removing the coins from both user's accounts, the simulator updates the channel space sets, i.e. defines $\Gamma^A(\gamma.\mathsf{id}) = \Gamma^B(\gamma.\mathsf{id}) = \Gamma(\gamma.\mathsf{id}) = \gamma$.

Registration of a contract instance in a ledger state channel. Since registration of a contract instance is defined as a separate procedure that can be called by parties of the protocol $\Pi(1, \mathcal{C})$, we define a "subsimulator" $\texttt{SimRegister}(P, id, cid)$ which can be called as a procedure by the simulator Sim_1. We define the subsimulator formally below. Let us here discuss one technicality.

As already mentioned, one of the main challenges of the simulation is to ensure the consistency of the ledger accounts in the ideal and hybrid world. In particular, if two parties created a ledger state channel between them (i.e. their coins were subtracted from their ledger accounts), the simulator has to ensure that once this ledger state channel is closed, the amount of coins returned to each party's account on the ledger is the same in the real and hybrid world. In case at least one party of the ledger state channel is honest, every time the channel is updated or executed, the ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ receives the corresponding message from the honest party and thus has the same view on the channel's state as the honest party in the hybrid world. The situation is more tricky in case both parties are corrupt.

If two corrupt parties have a ledger state channel between them, they can update its state arbitrarily (even to an invalid state). As long as these updates are done off-chain (parties exchange messages with each other and do not send any message to the hybrid ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$), no changes in the channel space Γ of ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$ are needed. Only when parties successfully register a contract instance with the hybrid ideal functionality $\mathcal{F}_{scc}^{\hat{\mathcal{L}}(\Delta)}(\mathcal{C})$, the update of the ledger state channel resulting from the new contract instance becomes "official". Thus, the simulator has to ensure that these changes to the ledger state channel are also made in the ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$. This is the reason, why the simulator has to send update message to the ideal functionality on behalf of the corrupt parties, in case they successfully register a contract instance in the hybrid world.

Sub-simulator : $\texttt{SimRegister}(P, id, cid)$

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$.

Case P and Q are honest:

1. Let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$ and $\gamma^Q := \Gamma^Q(id)$, $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$.
2. Wait up to 2Δ rounds and then proceed as follows. If $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$, then set $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$. Else set $\tilde{\nu} := (\nu^Q.\mathsf{storage}, \nu^Q.\mathsf{code})$.
3. Mark (id, cid) as registered in $\Gamma^P_{aux}, \Gamma^Q_{aux}$ and update all three sets $\Gamma, \Gamma^P, \Gamma^Q$, i.e. set $\Gamma :=$ $\texttt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$, $\Gamma^P := \texttt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$, $\Gamma^Q := \texttt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$.

Case P is honest and Q is corrupt:

1. Let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, $\sigma^P := \nu^P.\mathsf{storage}$.
2. Set τ_0 be the current round. Send $(\text{instance–registering}, id, cid, \nu^P) \xhookrightarrow{\tau_1 \leq \tau_0 + \Delta} Q$.
3. If $(\text{instance–register}, id, cid, \nu^Q) \xleftrightarrow{\tau_1} Q$ where ν^Q is a valid contract instance (both signatures $\nu^Q.\mathsf{sign}(A)$ and $\nu^Q.\mathsf{sign}(B)$ are valid, the amount of locked coins in ν^Q is non-negative, users of the contract instance are A and B, the contract instance storage is admissible and the contract code is from the set $\mathcal{C}$), then proceed as follows. If $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$, then $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$ and otherwise set $\tilde{\nu} := (\nu^Q.\mathsf{storage}, \nu^Q.\mathsf{code})$. Thereafter $(\text{instance–registered}, id, cid, \tilde{\nu}) \xhookrightarrow{\tau_2 \leq \tau_1 + \Delta} Q$ and goto step 5.
4. Else define $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$, send $(\text{instance–registered}, id, cid, \tilde{\nu}) \xhookrightarrow{\tau_2 \leq \tau_0 + 3\Delta} Q$ and goto step 5.
5. Mark (id, cid) as registered in Γ^P_{aux} , set $\Gamma := \texttt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$ and $\Gamma^P := \texttt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$.

Case P is corrupt and Q is honest:

Upon $(\text{instance–register}, id, cid, \nu^P) \xleftrightarrow{\tau_0} P$, s.t. $\Gamma(id) \neq \bot$, $\Gamma(id).\mathsf{cspace}(cid) = \bot$, ν^P is a valid contract instance (both signatures $\nu^P.\mathsf{sign}(A)$ and $\nu^P.\mathsf{sign}(B)$ are valid, the amount of locked coins in ν^P is non-negative, users of the contract instance are A and B, the contract instance storage is admissible and the contract code is from the set $\mathcal{C}$), then do:

1. Within Δ rounds, send $(\text{instance–registering}, id, cid, \nu^P) \xhookrightarrow{\tau_1 \leq \tau_0 + \Delta} P$.
2. Let $\gamma^Q := \Gamma^Q(id)$, $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$. If $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$, then $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$ and otherwise set $\tilde{\nu} := (\nu^Q.\mathsf{storage}, \nu^Q.\mathsf{code})$.
3. Send $(\text{instance–registered}, id, cid, \tilde{\nu}) \xhookrightarrow{\tau_2 \leq \tau_1 + \Delta} P$, mark (id, cid) as registered in Γ^Q_{aux} and then set $\Gamma := \texttt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$ and $\Gamma^Q := \texttt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$.

Case P and Q are corrupt :

Upon $(\text{instance–register}, id, cid, \nu^P) \xleftrightarrow{\tau_0} P$, s.t. $\Gamma(id) \neq \bot$, $\Gamma(id).\mathsf{cspace}(cid) = \bot$, ν^P is a valid contract instance (both signatures $\nu^P.\mathsf{sign}(A), \nu^P.\mathsf{sign}(B)$ are valid, the amount of locked money in ν^P is non-negative, users of the contract instance are A and B, the contract instance storage is admissible and the contract code is from the set $\mathcal{C}$), then do:

1. Within Δ rounds, send $(\text{instance–registering}, id, cid, \nu^P) \xhookrightarrow{\tau_1 \leq \tau_0 + \Delta} \Gamma(id).\mathsf{end\text{–}users}$.
2. If $(\text{instance–register}, id, cid, \nu^Q) \xleftrightarrow{\tau_1} Q$ s.t. ν^Q is a valid contract instance (both $\nu^Q.\mathsf{sign}(A)$ and $\nu^Q.\mathsf{sign}(B)$ are valid signatures, the amount of locked money in ν^Q is non-negative, users of the contract instance are A and B, the contract instance storage is admissible and the contract code is from the set $\mathcal{C}$), then proceed as follows. If $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$, then $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$ and otherwise set $\tilde{\nu} := (\nu^Q.\mathsf{storage}, \nu^Q.\mathsf{code})$. Thereafter send $(\text{instance–registered}, id, cid, \tilde{\nu}) \xhookrightarrow{\tau_2 \leq \tau_1 + \Delta} \Gamma(id).\mathsf{end\text{–}users}$ and goto step 4.
3. Else proceed as follows. If $(\text{finalize–register}, id, cid) \xhookleftarrow{\tau_0 + 2\Delta} P$, then define $\tilde{\nu} := (\nu^P.\mathsf{storage}, \nu^P.\mathsf{code})$, send $(\text{instance–registered}, id, cid, \tilde{\nu}) \xhookrightarrow{\tau_2 \leq \tau_0 + 3\Delta} \Gamma(id).\mathsf{end\text{–}users}$ and goto step 4.
4. Update the channel space $\Gamma := \texttt{UpdateChanSpace}^*(\Gamma, id, cid, \tilde{\nu})$. Send $(\text{update}, id, cid, \tilde{\nu}.\mathsf{storage}, \tilde{\nu}.\mathsf{code}) \hookrightarrow \mathcal{F}_{ch}$ on behalf of P and $(\text{update–reply}, ok, id, cid) \hookrightarrow \mathcal{F}_{ch}$ on behalf of Q.

Update a contract instance in a ledger state channel If both parties are honest, the simulator does not need to give any instructions to the ideal functionality and only updates the sets Γ^P, Γ^Q,Γ^P_{aux}, Γ^Q_{aux}, when the messages $(P, \text{update}, id, cid, \tilde{\sigma}, \mathtt{C})$ and $(Q, \text{update–reply}, ok, id, cid)$ are received from the ideal functionality. In case both parties are corrupt, the simulator can internally simulate the communication of the two corrupt parties and in case the registration procedure is started by one of them, it executes the subsimulator `SimRegister` for the case when both parties are corrupt. Note that if the registration procedure is successful (a contract instance gets registered), the subsimulator `SimRegister` instructs the ideal functionality to update the contract instance accordingly. We define the simulator Sim_1 for the remaining two case, i.e. when only the initiating party is corrupt and if only the reacting party is corrupt, below.

Simulator Sim_1: Contract instance update

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$.

Case P is honest and Q is corrupt:

Upon $(P, \text{update}, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}$ do:

1. If (id, cid) is marked as corrupt in Γ^P_{aux}, then stop. Else let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, $\sigma^P := \nu^P.\mathsf{storage}$. If $\nu^P = \perp$, then set $w^P := 1$, else set $w^P := \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}$.
2. Sign $s_P := \mathtt{Sign}_{sk_P}(id, cid, \tilde{\sigma}, \mathtt{C}, w^P)$ and send $(\text{update}, s_P, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookrightarrow{\tau_0+1} Q$ of behalf of P.
3. Distinguish the following cases:
   - If $(\text{update–ok}, s_Q) \xhookleftarrow{\tau_1 \leq \tau_0+1} Q$ such that $\mathtt{Vfy}_{pk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}, w^P; s_B) = 1$, then send $(\text{update–reply}, ok, id, cid) \xhookrightarrow{\tau_1} \mathcal{F}_{ch}$ on behalf of Q, set $\Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version} := w + 1$ and $\Gamma^P := \mathtt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\sigma}, \mathtt{C}, w^P, \{s_P, s_Q\})$.
   - If $(\text{update–not–ok}, s_Q) \xhookleftarrow{\tau_1 \leq \tau_0+1} Q$ such that $\mathtt{Vfy}_{pk_B}(id, cid, \sigma^P, \mathtt{C}, w^P + 1; s_Q) = 1$, then compute $s_P := \mathtt{Sign}_{sk_P}(id, cid, \sigma^P, \mathtt{C}, w^P + 1)$ and set $\Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}^P := w + 2$ and $\Gamma^P := \mathtt{UpdateChanSpace}^*(\Gamma^P, id, cid, \sigma^P, \mathtt{C}, w^P + 1, \{s_P, s_Q\})$.
   - Else mark (id, cid) as corrupt in Γ^P_{aux} and execute $\mathtt{SimRegister}(P, id, cid)$. If after the subsimulator is executed (in round $\tau_2 \leq \tau_0 + 3\Delta + 1$) it holds that $\Gamma^P(id).\mathsf{cspace}(cid) = (\tilde{\sigma}, \mathtt{C})$, then $(\text{update–reply}, ok, id, cid) \xhookrightarrow{\tau_2} \mathcal{F}_{ch}$ on behalf of Q.

Case P is corrupt and Q is honest:

Upon $(\text{update}, s_P, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookleftarrow{\tau_0} P$ do:

1. If (id, cid) is marked as corrupt in Γ^Q_{aux}, then stop. Let $\gamma^Q := \Gamma^Q(id)$. If $\gamma^Q = \perp$ or there exists $cid' \neq cid$ such that $\gamma.\mathsf{cspace}(cid) \neq \perp$, then stop; else let $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$. If $\nu^Q = \perp$, then set $w^Q := 1$, else set $w^Q := \Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version}$.
2. If $\mathtt{Vfy}_{pk_P}(id, cid, \tilde{\sigma}, \mathtt{C}, w^Q; s_P) \neq 1$, then mark (id, cid) as corrupt in Γ^Q_{aux} and stop. Else send $(\text{update}, id, cid, \tilde{\sigma}, \mathtt{C}) \xhookrightarrow{\tau_0} \mathcal{F}_{ch}$ on behalf of P.
3. Distinguish the following cases:
   - If $(Q, \text{update–reply}, ok, id, cid) \xhookleftarrow{\tau_1 \leq \tau_0+1} \mathcal{F}_{ch}$, then compute $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \tilde{\sigma}, \mathtt{C}, w^Q)$, set $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q{+}1$ and $\Gamma^Q := \mathtt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\sigma}, \mathtt{C}, w^Q, \{s_P, s_Q\})$ and send $(\text{update–ok}, s_Q) \xhookrightarrow{\tau_0+2} P$ on behalf of Q and stop.
   - Else set $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q{+}2$ compute $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \nu^Q.\mathsf{storage}, \nu^Q.\mathsf{code}, w^Q + 1)$ and on behalf of Q send $(\text{update–not–ok}, s_Q) \xhookrightarrow{\tau_0+2} P$.

Execute a contract instance in a ledger state channel In case both parties are honest, the simulator only has to instruct the ideal functionality to output the result in the correct round. Let τ_0 be the round in which the environment instructed the initiating party P to execute. Then the simulator sets $\tau_1 := \tau_0 + x$, where x is the smallest offset such that $\tau_1 = 1 \mod 4$ if $P = \gamma.\mathsf{Alice}$ and $\tau_1 = 3 \mod 4$ if $P = \gamma.\mathsf{Bob}$ and waits until round τ_1 to instruct the ideal functionality to output the result. Then it updates both channel spaces

Γ^P, Γ^Q, Γ^P_{aux} and Γ^Q_{aux} accordingly. We formally describe the situation when one or two parties are corrupt below.

Simulator Sim_1: Contract instance execution.

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$.

Case P is honest and Q is corrupt:

Upon $(P, \text{execute}, id, cid, f, z) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}$, let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, $\sigma^P := \nu^P.\mathsf{storage}$ and $w^P := \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}$. In addition, set $\tau_1 := \tau_0 + x$, where x is the smallest offset such that $\tau_1 = 1 \mod 4$ if $P = \gamma^P.\mathsf{Alice}$ and $\tau_1 = 3 \mod 4$ if $P = \gamma^P.\mathsf{Bob}$. Wait until round τ_1 and then proceed as follows:

1. If (id, cid) is not marked as corrupt in Γ^P_{aux}, do:
   (a) Set $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^P, P, \tau_0, z)$. If $m = \bot$, then stop.
   (b) Else compute $s_P := \mathtt{Sign}_{sk_P}(id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, w^P)$, send $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookrightarrow{\tau_1+1} Q$ and instruct the ideal functionality to output the execute requested message.
   (c) If $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookleftarrow{\tau_1+1} Q$ such that $\mathtt{Vfy}_{pk_Q}(id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, w^P; s_Q) = 1$, then set $\Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version} := w^P + 1$ and $\Gamma^P := \mathtt{UpdateChanSpace}(\Gamma^P, id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, add_L, add_R, w^P, \{s_P, s_Q\})$. Then instruct the ideal functionality to output the result. Else mark (id, cid) as corrupt in Γ^P_{aux} and execute the subsimulator $\mathtt{SimRegister}(P, id, cid)$ in round $\tau_1 + 2$. If after the execution of the sub-simulator (in round $\tau_1 \leq \tau_0 + 3\Delta + 5$) it holds that $\sigma^P = \tilde{\sigma}$, then set $\Gamma^P := \mathtt{UpdateChanSpace}(\Gamma^P, id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, add_L, add_R)$, instruct the ideal functionality to output the result and stop. Else goto step 2b.
2. If (id, cid) is marked as corrupt in Γ^P_{aux}
   (a) If (id, cid) is not marked as registered in Γ^P_{aux}, then execute the sub-simulator $\mathtt{SimRegister}(P, id, cid)$.
   (b) Let τ_3 be the current round. If $(\text{executed}, id, cid, \sigma, add_L, add_R, m) \xhookleftarrow{\tau_4 \leq \tau_3 + \Delta} \mathcal{F}_{ch}$, then update the channel space Γ^P and Γ, i.e. set $\Gamma := \mathtt{UpdateChanSpace}(\Gamma, id, cid, \sigma, \nu^P.\mathsf{code}, add_L, add_R)$ and $\Gamma^P := \mathtt{UpdateChanSpace}(\Gamma^P, id, cid, \sigma, \nu^P.\mathsf{code}, add_L, add_R)$. Thereafter send $(\text{instance–executed}, id, cid, \sigma, add_L, add_R, m) \xhookrightarrow{\tau_4} Q$ and stop. Else stop.

Case P is corrupt and Q is honest:

Upon $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookleftarrow{\tau_1} P$

1. Let $\gamma^Q := \Gamma^Q(id)$, $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$, $\sigma^Q := \nu^Q.\mathsf{storage}$, $w^Q := \Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version}$. If $\gamma^Q = \bot$, $P \notin \gamma^Q.\mathsf{end\text{–}users}$, $\nu^Q = \bot$ or $f \notin \nu^Q.\mathsf{code}$, then goto step 4.
2. If $P = \gamma^Q.\mathsf{Alice}$ and $\tau_1 \mod 4 \neq 1$ or if $P = \gamma.\mathsf{Bob}$ and $\tau_1 \mod 4 \neq 3$, then goto step 4.
3. If (id, cid) is not marked as corrupt in Γ^Q_{aux}, do:
   (a) Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^Q, P, \tau_0, z)$.
   (b) If $m = \bot$ or $\mathtt{Vfy}_{pk_P}(id, cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, w^Q; s_P) \neq 1$, then goto step 4.
   (c) Send $(\text{execute}, id, cid, f, z) \xhookrightarrow{\tau_1} \mathcal{F}_{ch}$ on behalf of P and instruct the ideal functionality to set $\tau := \tau_0$.
   (d) Compute the signature $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, w^Q)$, send $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookrightarrow{\tau_1+1} P$, set $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q + 1$, $\Gamma^Q := \mathtt{UpdateChanSpace}(\Gamma^Q, id, cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, add_L, add_R, w^Q, \{s_P, s_Q\})$, instruct the functionality to deliver the result and stop.

4. Mark (id, cid) as corrupt in Γ^Q_{aux} and stop.

Upon P starting the registration procedure for id, cid, then execute the sub-simulator $\texttt{SimRegister}(P, id, cid)$.
Upon $(\text{instance–execute}, id, cid, f, z, \tau_0) \xleftarrow{\tau_2} P$, then

1. If $\tau_2 - \tau_0 > 5$, then stop. Let $\gamma := \Gamma(id)$. If $\gamma = \perp$ or $P \notin \gamma.\textsf{end–users}$, then stop. Else let $\nu := \gamma.\textsf{cspace}(cid)$, $\sigma := \nu.\textsf{storage}$. If $\nu = \perp$ or $f \notin \nu.\textsf{code}$, stop.
2. Else send $(\text{execute}, id, cid, f, z) \xrightarrow{\tau_2} \mathcal{F}_{ch}$ on behalf of P, instruct the functionality $\mathcal{F}_{ch}$ to set $\tau := \tau_0$ and within Δ round instruct the functionality to output the result.
3. When $(\text{executed}, id, cid, \sigma, add_L, add_R, m) \xleftarrow{\tau_3 \leq \tau_2 + \Delta} \mathcal{F}_{ch}$, then update the sets Γ^Q and Γ, i.e. set $\Gamma := \texttt{UpdateChanSpace}(\Gamma, id, cid, \sigma, \nu^Q.\textsf{code}, add_L, add_R)$ and $\Gamma^Q := \texttt{UpdateChanSpace}(\Gamma^Q, id, cid, \sigma, \nu^Q.\textsf{code}, add_L, add_R)$. Then send $(\text{instance–executed}, id, cid, \sigma, add_L, add_R, m) \xrightarrow{\tau_3} P$ and stop.

Case P and Q are corrupt:

Internally simulate the communication of the corrupt parties. If P starting the registration procedure for id, cid, then execute the sub-simulator $\texttt{SimRegister}(P, id, cid)$ for the case when both parties are corrupt. Note that if the registration procedure is successful (a contract instance gets registered), the subsimulator $\texttt{SimRegister}$ instructs the ideal functionality to update the contract instance accordingly.
If $(\text{instance–execute}, id, cid, f, z, \tau_0) \xleftarrow{\tau_2} P$, then

1. If $\tau_2 - \tau_0 > 5$, then stop. Let $\gamma := \Gamma(id)$. If $\gamma = \perp$ or $P \notin \gamma.\textsf{end–users}$, then stop. Else let $\nu := \gamma.\textsf{cspace}(cid)$, $\sigma := \nu.\textsf{storage}$. If $\nu = \perp$ or $f \notin \nu.\textsf{code}$, stop.
2. Else send $(\text{execute}, id, cid, f, z) \xrightarrow{\tau_2} \mathcal{F}_{ch}$ on behalf of P, instruct the ideal functionality $\mathcal{F}_{ch}$ to set $\tau := \tau_0$ and within Δ round instruct the $\mathcal{F}_{ch}$ to output the result.
3. When $(\text{executed}, id, cid, \sigma, add_L, add_R, m) \xleftarrow{\tau_3 \leq \tau_2 + \Delta} \mathcal{F}_{ch}$, then set $\Gamma := \texttt{UpdateChanSpace}(\Gamma, id, cid, \sigma, \nu.\textsf{code}, add_L, add_R)$, send $(\text{instance–executed}, id, cid, \sigma, add_L, add_R, m) \xrightarrow{\tau_3} P$ and stop.

Close a ledger state channel. The simulator Sim_1 is formally defined for all four possible situations below .

Simulator Sim_1: Close a ledger state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\mathcal{F}_{ch} := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(1, \mathcal{C})$.

Case P, Q are honest

Upon $(P, \text{close}, id) \xleftarrow{\tau_0} \mathcal{F}_{ch}$, proceed as follows. Let $\gamma^P := \Gamma^P(id)$. If there exists $cid \in \{0,1\}^*$ such that $\gamma^P.\textsf{cspace}(cid) \neq \perp$, the execute $\texttt{SimRegister}(P, id, cid)$ for the case when both parties are honest. In round $\tau_1 \leq \tau_0 + 8\Delta$ instruct the ideal functionality to output the result. If $(\text{closed}, id) \xleftarrow{\tau_1 \leq \tau_0 + 8\Delta} \mathcal{F}_{ch}$, set $\Gamma(id) := \perp$, $\Gamma^P(id) := \perp$, $\Gamma^Q(id) := \perp$ and stop.

Case P is honest and Q is corrupt:

Upon $(P, \text{close}, id) \xleftarrow{\tau_0} \mathcal{F}_{ch}$, do:
1. Let $\gamma^P := \Gamma^P(id)$. If there exists $cid \in \{0,1\}^*$ such that $\gamma^P.\textsf{cspace}(cid) \neq \perp$ but the contract instance has never been registered, execute $\texttt{SimRegister}(P, id, cid)$.

2. After the execution of the subsimulator, wait for at most Δ rounds to send the message (contract–closing, id) $\xhookrightarrow{\tau_2 \leq \tau_0 + 4\Delta}$ Q.
3. Execute the sub-simulator $\texttt{SimRegister}(Q, id, cid)$ if registration started by Q for some cid.
4. In round $\tau_3 \leq \tau_0 + 8\Delta$ instruct the ideal functionality to output the result. If (closed, id) $\xhookleftarrow{\tau_3}$ $\mathcal{F}_{ch}$, set $\Gamma(id) := \perp$ $\Gamma^P(id) := \perp$ and send (contract–closed, id) $\xhookrightarrow{\tau_3}$ Q. Then stop.

$\boxed{\textbf{Case } P \textbf{ is corrupt and } Q \textbf{ is honest:}}$

1. Execute the sub-simulator $\texttt{SimRegister}(P, id, cid)$ if registration started by P for some cid in round τ_0.
2. After the execution (in round $\tau_1 \leq \tau_0 + 2\Delta$), if (contract–close, id) $\xhookleftarrow{\tau_1}$ P, where $\Gamma(id) \neq \perp$, then send (close, id) $\xhookrightarrow{\tau_1}$ $\mathcal{F}_{ch}$ on behalf of P.
3. Wait at most Δ rounds to (contract–closing, id) $\xhookrightarrow{\tau_2 \leq \tau_0 + 3\Delta}$ P.
4. Let $\gamma^Q := \Gamma^Q(id)$. If there exists cid such that $\gamma^Q.\mathsf{cspace}(cid) \neq \perp$ but the contract instance has never been registered, execute the sub-simulator $\texttt{SimRegister}(Q, id, cid)$.
5. Upon (closed, id) $\xhookleftarrow{\tau_5 \leq \tau_0 + 8\Delta}$ $\mathcal{F}_{ch}$, $\Gamma^Q(id) := \perp$ and (contract–closed, id) $\xhookrightarrow{\tau_5}$ P and stop.

$\boxed{\textbf{Case } P \textbf{ and } Q \textbf{ are corrupt:}}$

1. Execute the sub-simulator $\texttt{SimRegister}(P, id, cid)$ if registration started by P for some cid. Note that if the registration procedure is successful (a contract instance gets registered), the subsimulator $\texttt{SimRegister}$ instructs the ideal functionality to update the contract instance accordingly.
2. After the execution (in round $\tau_1 \leq \tau_0 + 2\Delta$), if (contract–close, id) $\xhookleftarrow{\tau_1}$ P, where $\Gamma(id) \neq \perp$, then send (close, id) $\xhookrightarrow{\tau_1}$ $\mathcal{F}_{ch}$ on behalf of P.
3. Wait at most Δ rounds to (contract–closing, id) $\xhookrightarrow{\tau_2 \leq \tau_0 + 4\Delta}$ $\Gamma(id).\mathsf{end\text{–}users}$.
4. Execute the sub-simulator $\texttt{SimRegister}(Q, id, cid)$ if registration started by Q for some cid. Again, if the registration procedure is successful, the subsimulator $\texttt{SimRegister}$ instructs the ideal functionality to update the contract instance accordingly.
5. In round $\tau_3 \leq \tau_0 + 8\Delta$ instruct the ideal functionality to output the result. If (closed, id) $\xhookleftarrow{\tau_1 \leq \tau_0 + 8\Delta}$ $\mathcal{F}_{ch}$, set $\Gamma(id) := \perp$ and stop.

F Security analysis for virtual state channels

The purpose of this section is to show that for any $i > 1$ and any set $\mathcal{C}$ of contract codes, the protocol $\Pi(i, \mathcal{C})$ emulates the ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ in $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$-hybrid world against environments from the set $\mathcal{E}_{res}$.

The proof consists of two parts. First, we need to prove an auxiliary lemma stating that an instance of the protocol $\Pi(i, \mathcal{C})$ called by an environment $\mathcal{Z} \in \mathcal{E}_{res}$ is $\mathcal{E}_{res}$-respecting. This is because the hybrid ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ is emulated by the protocol $\Pi(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ only against environments from the set $\mathcal{E}_{res}$. This proves that the hybrid world is well defined and the composition of state channel protocols is possible. Thereafter we can construct the simulator Sim_i in order to prove that the protocol $\Pi(i, \mathcal{C})$ in the hybrid world of $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ emulates the ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ against environments from the set $\mathcal{E}_{res}$.

Lemma 2. *For any $i > 1$, set of contract codes $\mathcal{C}$, PPT adversary* Adv *and environment $\mathcal{Z} \in \mathcal{E}_{res}$, the protocol $\Pi(i, \mathcal{C})$ is $\mathcal{E}_{res}$-respecting.*

Proof. We need to prove that for any PPT adversary Adv and any environment $\mathcal{Z} \in \mathcal{E}_{res}$, honest parties of the protocol $\Pi(i, \mathcal{C})$ make calls to the hybrid ideal functionality $\mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ according to the restrictions defining the set $\mathcal{E}_{res}$. In other words, honest parties of the protocol jointly represent an environment from the set $\mathcal{E}_{res}$.

If the environment $\mathcal{Z}$ sends a message to an honest party in the protocol regarding a state channel of length $j < i$, then the party simply forwards the message to the hybrid ideal functionality. Since $\mathcal{Z} \in \mathcal{E}_{res}$, no invalid calls can be made to the hybrid functionality in this way. It remains to show that the protocol is $\mathcal{E}_{res}$-respecting even if the environments sends a message regarding a virtual state channel of length i.

First note that honest parties in the protocol $\Pi(i, \mathcal{C})$ upon receiving a message about a virtual state channel of length i only ask the hybrid ideal functionality to update or execute a contract instance in a state channel but never to create or close a state channel. Thus, none of the restrictions regarding creating or closing a state channel can be violated.

Parties of the protocol send messages regarding update of a contract instance to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ only during the protocol "Create a virtual state channel". Since we assume that parties of the protocol receive messages from an environment $\mathcal{Z} \in \mathcal{E}_{res}$, we have the guarantee that they all receive the message (create, γ) in the same round τ_0. According to the protocol, party $\gamma.\mathsf{Alice}$ sends in round τ_0 the message $(\text{update}, id_A, cid_A, \tilde{\sigma}_A, \mathtt{VSCC}_i(\mathcal{C}))$, where $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(\gamma.\mathsf{Alice}, \tau_0, \gamma)$, $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{Alice})$ and $cid_A := \gamma.\mathsf{Alice}||\gamma.\mathsf{id}$. Hence clearly $\tilde{\sigma}_A$ is admissible with respect to $\mathtt{VSCC}_i(\mathcal{C})$ and $\tilde{\sigma}_A.\mathsf{locked} = \tilde{\sigma}_A.\mathsf{cash}(A) + \tilde{\sigma}_A.\mathsf{cash}(I)$. We can argue similarly with the update of the subchannel between $\gamma.\mathsf{Ingrid}$ and $\gamma.\mathsf{Bob}$. Since $\mathcal{Z} \in \mathcal{E}_{res}$, we know that both subchannels of the virtual state channel γ exist, that they contain no contract instances and that they have enough funds. In addition, the subchannels do support contracts with code $\mathtt{VSCC}_i(\mathcal{C})$ since they were created via the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$.

Parties of the protocol send messages regarding execution of a contract instance to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ (i) during the protocol "Update a contract instance in a virtual state channel" (more specifically in the procedure $\mathtt{Register}_i$), (ii) during the protocol "Execute a contract instance in a virtual state channel" and (iii) during the protocol "Close virtual state channel". Since $\mathcal{Z} \in \mathcal{E}_{res}$, we know that none these protocols is ever called for a state channel that does not exists. This in particular implies that the contract instance that is being executed by parties of the protocol in the underlying subchannels must have been constructed and could not have been closed yet. In other words, we know that $\alpha.\mathsf{cspace}(cid_A) \neq \perp$ and $\beta.\mathsf{cspace}(cid_B) \neq \perp$, where $cid_A := \gamma.\mathsf{Alice}||\gamma.\mathsf{id}$, $\alpha := \Gamma^A(\gamma.\mathsf{subchan}(\gamma.\mathsf{Alice}))$ and $cid_B := \gamma.\mathsf{Bob}||\gamma.\mathsf{id}$, $\beta := \Gamma^B(\gamma.\mathsf{subchan}(\gamma.\mathsf{Bob}))$, where Γ^A and Γ^B are the channel spaces of $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Bob}$, respectively.

In order to complete the proof that $\Pi(i, \mathcal{C})$ protocol $\mathcal{E}_{res}$-emulates the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ in $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$-hybrid world for any set of contract codes $\mathcal{C}$, we need for every adversary Adv to construct a simulator Sim_i that simulates the hybrid world for any environment $\mathcal{Z} \in \mathcal{E}_{res}$.

The simulator Sim_i constructed in this section maintains a channel space Γ^T and auxiliary channel space Γ_{aux}^T for every honest party $T \in \mathcal{P}$ and Γ for the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$. In addition, the simulator generates a key pair $(pk_T, sk_T) \leftarrow\!\!\$\, \mathsf{KGen}(1^\lambda)$ for every honest party T during the setup phase which allows Sim_i to internally run a copy of the hybrid world. Recall that there are no private inputs or messages being sent, thus we assume that the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ upon receiving a message m from party P immediately sends the message (P, m) to the simulator Sim_i.

We discuss in detail the most interesting case, when the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ sends message about a virtual state channel of length exactly i or when a corrupt party P is instructed by the environment to update or execute a subchannel of a virtual state channel of length exactly i, where the other user of the subchannel is not corrupt. The simulation in the remaining cases is straightforward. Let us describe it here only briefly.

If the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ sends a message about a state channel of length j, where $1 \leq j < i$, the simulator internally executes the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ on the received message and sends the result to the adversary Adv (recall that honest parties in the protocol $\Pi(i, \mathcal{C})$ act like dummy parties and only forward messages to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$). If the corrupt parties are instructed to send valid replies to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$, the simulator Sim_i sends the messages to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ on their behalf and further instructs the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ as the simulator Sim_j would do. Thus specially, if all parties of a state channel are honest, then the simulator Sim_i is defined exactly as the simulator Sim_j. Let us give one example on how the simulator is defined in case there are corrupt parties.

Let us consider the situation when $\gamma.\mathsf{Alice}$ and $\gamma.\mathsf{Ingrid}$ are honest, $\gamma.\mathsf{Bob}$ is corrupt and the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ sends the messages $(\gamma.\mathsf{Alice}, \text{create}, \gamma)$ and $(\gamma.\mathsf{Ingrid}, \text{create}, \gamma)$, where $1 < \gamma.\mathsf{length} < i$, in round τ_0. Then the simulator waits until round $\tau_0 + 3$ if the corrupt party $\gamma.\mathsf{Bob}$ is instructed to send (create, γ) to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$. In that case, Sim_i forwards the message to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ on behalf of $\gamma.\mathsf{Bob}$, adds the new virtual state channel γ to the channel spaces Γ^A and Γ. The simulator then waits until round $\gamma.\mathsf{validity}$.

The simulator Sim_i is defined similarly in the remaining case when it does not receive any message from the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ but a corrupt party is instructed to send a message to the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$ about a state channel of length $1 \leq j < i$. This happens if a corrupt party is the initiator of execute or update procedure or when all parties of the state channel are corrupt. In this situation, the simulator Sim_i internally executes the hybrid ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$. In case the message satisfies the restrictions on the environment, Sim_i forwards it to the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ on behalf of the corrupt party and further instructs the ideal functionality $\mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ as the simulator Sim_j would do.

From now on, we will focus on the simulator Sim_i for the most challenging case when at least one party of a virtual state channel of length exactly i is honest. We begin with the definition of the simulator for virtual state channel creation. Then, similarly as for the ledger state channels, we separately define a subsimulator $\mathtt{SimRegister}_i$ which can be called as a procedure by the simulator Sim_i. The description of the simulator Sim_i for the contract instance update in a virtual state channel of length i will be very similar to the simulator Sim_1. Therefore, we refer the reader to the described in Appx. E and discuss here only the main differences. Firstly, the simulator Sim_i internally calls the subsimulator $\mathtt{SimRegister}_i$ instead of the subsimulator $\mathtt{SimRegister}$ and secondly, in case the initiating party P is corrupt the simulator Sim_i also checks if there is no other contract instance cid' already created in the virtual state channel (recall that we allow only one contract instance to be opened in each virtual state channel). The simulator Sim_i for the execution of a contract instance in case both end-users of the virtual state channel are honest is defined exactly as the simulator Sim_1, see Appx. E. The remaining cases are formally described below. We finalize the definition of the simulator Sim_i by defining its behavior in time $\gamma.\mathsf{validity}$, where γ is a virtual state channel of length i whose creation environment requested earlier.

Simulator Sim_i: Create a virtual state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. We denote the ideal functionality $\mathcal{F}_{ch}(i) := \mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i,\mathcal{C})$ and $\mathcal{F}_{ch}(i-1) := \mathcal{F}_{ch}^{\widehat{\mathcal{L}}(\Delta)}(i-1,\mathcal{C})$.

Case A, I, B are honest

Upon receiving $(A, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$, $(B, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ and $(I, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ proceed as follows:

1. Set $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{id})$, $cid_A := A||\gamma.\mathsf{id}$ and $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B := B||\gamma.\mathsf{id}$. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(A, \tau_0, \gamma)$ and $\tilde{\sigma}_B = \mathtt{Init}_i^{\mathcal{C}}(B, \tau_0, \gamma)$.
2. For both $T \in \{A, B\}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update}, id_T, cid_T, \tilde{\sigma}_T, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} T$.
3. For both $T \in \{A, B\}$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update–reply}, ok, id_T, cid_T) \xhookleftarrow{\tau_0+1} I$.
4. Set $\Gamma^A(\gamma.\mathsf{id}) := \gamma$, $\Gamma^B(\gamma.\mathsf{id}) := \gamma$ and wait until round $\gamma.\mathsf{validity}$.

Case A, B are honest and I is corrupt:

Upon receiving $(A, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ and $(B, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ proceed as follows:

1. Set $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{id})$, $cid_A := A||\gamma.\mathsf{id}$ and $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B := B||\gamma.\mathsf{id}$. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(A, \tau_0, \gamma)$ and $\tilde{\sigma}_B = \mathtt{Init}_i^{\mathcal{C}}(B, \tau_0, \gamma)$.
2. For both $T \in \{A, B\}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update}, id_T, cid_T, \tilde{\sigma}_T, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} T$ and forward the result to I.
3. If $(\text{update–reply}, ok, id_T, cid_T) \xhookleftarrow{\tau_0+1} I$ for $T \in \{A, B\}$, then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving this message and forward the result to I.
4. If in round $\tau_0 + 1$, party I confirms both updates, then send $(\text{create}, \gamma) \xhookrightarrow{\tau_0+1} \mathcal{F}_{ch}(i)$ on behalf of I, set $\Gamma^A(\gamma.\mathsf{id}) := \gamma$, $\Gamma^B(\gamma.\mathsf{id}) := \gamma$
5. Wait until round $\gamma.\mathsf{validity}$.

Simulator Sim_i: Create a virtual state channel

Case A, I are honest and B is corrupt:

Upon $(A, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ and $(I, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$, proceed as follows:

1. Set $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{id})$, $cid_A := A||\gamma.\mathsf{id}$ and $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B := B||\gamma.\mathsf{id}$. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(A, \tau_0, \gamma)$ and $\tilde{\sigma}_B = \mathtt{Init}_i^{\mathcal{C}}(B, \tau_0, \gamma)$.
2. In round τ_0, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{update}, id_A, cid_A, \tilde{\sigma}_A, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} A$.
3. If $(\text{update}, id_B, cid_B, \tilde{\sigma}_B, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} B$, then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving this message. Else stop.
4. For both $T \in \{A, B\}$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update–reply}, ok, id_T, cid_T) \xhookleftarrow{\tau_0+1} I$ and forward the result of updating id_B to B.
5. Send $(\text{create–ok}, \gamma) \xhookrightarrow{\tau_0+3} B$ on behalf of A.
6. If $(\text{create–ok}, \gamma) \xhookleftarrow{\tau_0+2} B$, then send $(\text{create}, \gamma) \xhookrightarrow{\tau_0+3} \mathcal{F}_{ch}(i)$ on behalf of B, add γ to Γ^A.
7. Wait until round $\gamma.\mathsf{validity}$.

Case I, B are honest and A is corrupt:

Analogous to the case when only B is corrupt.

Case I, B are corrupt and A is honest:

Upon receiving $(A, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ proceed as follows:

1. Set $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{id})$, $cid_A := A||\gamma.\mathsf{id}$ and $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B \neq B||\gamma.\mathsf{id}$. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(A, \tau_0, \gamma)$.
2. In round τ_0, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{update}, id_A, cid_A, \tilde{\sigma}_A, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} A$ and forward the result to I.
3. If $(\text{update–reply}, ok, id_A, cid_A) \xhookleftarrow{\tau_0+1} I$, then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving this message, send $(\text{create}, \gamma) \xhookrightarrow{\tau_0+1} \mathcal{F}_{ch}(i)$ on behalf of I and send $(\text{create–ok}, \gamma) \xhookrightarrow{\tau_0+3} B$ on behalf of A.
4. If $(\text{create–ok}, \gamma) \xhookleftarrow{\tau_0+2} B$, then send $(\text{create}, \gamma) \xhookrightarrow{\tau_0+3} \mathcal{F}_{ch}(i)$ on behalf of B and add γ to Γ^A.
5. Wait until round $\gamma.\mathsf{validity}$.

Case A, I are corrupt and B is honest:

Analogous to the case when only A is honest.

Case A, B are corrupt and I is honest:

Upon receiving $(I, \text{create}, \gamma) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}(i)$ proceed as follows:

1. Set $id_A := \gamma.\mathsf{subchan}(\gamma.\mathsf{id}), cid_A := A||\gamma.\mathsf{id}$ and $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B \neq B||\gamma.\mathsf{id}$. Compute $\tilde{\sigma}_A := \mathtt{Init}_i^{\mathcal{C}}(A, \tau_0, \gamma)$ and $\tilde{\sigma}_B := \mathtt{Init}_i^{\mathcal{C}}(B, \tau_0, \gamma)$.
2. If $(\text{update}, id_A, cid_A, \tilde{\sigma}_A, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} A$ and in the same round $(\text{update}, id_B, cid_B, \tilde{\sigma}_B, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} B$, then proceed. Else stop.
3. For both $T \in \{A, B\}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update}, id_T, cid_T, \tilde{\sigma}_T, \mathtt{VSCC}_i(\mathcal{C})) \xhookleftarrow{\tau_0} T$.
4. For both $T \in \{A, B\}$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{update–reply}, ok, id_T, cid_T) \xhookleftarrow{\tau_0+1} I$ and forward the result to T.
5. Wait until round $\gamma.\mathsf{validity}$.

Subsimulator: $\mathtt{SimRegister}_i(P, id, cid)$

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, let $\text{TER}_{sub} := \text{TimeExeReq}(\lceil i/2 \rceil)$, $\text{TE}_{sub} := \text{TimeExe}(\lceil i/2 \rceil)$, $\mathcal{F}_{ch}(i) := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ and $\mathcal{F}_{ch}(i-1) := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \mathcal{C})$.

All parties are honest:

1. Let $\gamma := \Gamma^P(id)$, $id_P := \gamma.\mathsf{subchan}(P)$, $cid_P := P||\gamma.\mathsf{id}$, $id_Q := \gamma.\mathsf{subchan}(Q)$, $cid_Q := Q||\gamma^P.\mathsf{id}$, $\nu^P := \gamma.\mathsf{cspace}(cid)$, $\nu^Q := \Gamma^Q(id).\mathsf{cspace}(cid)$ and let τ_0 be the current round. Denote $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise.
2. In round τ_0, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$.
3. In round $\tau_1 \leq \tau_0 + 4$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1} I$.
4. In round $\tau_2 \leq \tau_1 + 4$, mark (id, cid) as registered in Γ_{aux}^Q, update the channels space $\Gamma^Q := \mathtt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_2} Q$.
5. In round $\tau_3 \leq \tau_2 + 4$, mark (id, cid) as registered in Γ_{aux}^P, update the channels space $\Gamma^P := \mathtt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_3} I$.

Case P, I are honest and Q is corrupt:

1. Let $\gamma := \Gamma^P(id)$, $id_P := \gamma.\mathsf{subchan}(P)$, $cid_P := P||\gamma.\mathsf{id}$, $id_Q := \gamma.\mathsf{subchan}(Q)$, $cid_Q := Q||\gamma.\mathsf{id}$, $\nu^P := \gamma.\mathsf{cspace}(cid)$.
2. In the current round τ_0, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$.
3. In round $\tau_1 \leq \tau_0 + 4$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \mathtt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1} I$, in round $\tau_2 \leq \tau_1 + \text{TER}_{sub}$ send the message $(\text{execute–requested}, id_Q,$

$cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P))$ to Q and once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to Q.

4. In τ_2 distinguish the following two cases:
   - If (execute, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \overset{\tau_2}{\hookleftarrow} Q$, where $\texttt{VerifyInstance}(id, cid, \nu^Q) = 1$, then
     (a) Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \overset{\tau_2}{\hookleftarrow} Q$ and once the internal execution is completed (before round $\tau_2 + \mathrm{TE}_{sub}$), forward the result to Q.
     (b) In round $\tau_3 \leq \tau_2 + \mathrm{TER}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \overset{\tau_3}{\hookleftarrow} I$.
     (c) Let $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise. In round $\tau_4 \leq \tau_3 + 4$, mark (id, cid) as registered in Γ^P_{aux} and update the channel space $\Gamma^P := \texttt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$.
   - Otherwise in round $\tau_5 := \tau_0 + 8 + 2 \cdot \mathrm{TER}_{sub}$ mark (id, cid) as registered in Γ^P_{aux}.

Case P, Q are honest and I is corrupt:

1. Let $\gamma := \Gamma^P(id)$, $id_P := \gamma.\mathsf{subchan}(P)$, $cid_P := P||\gamma.\mathsf{id}$, $id_Q := \gamma.\mathsf{subchan}(Q)$, $cid_Q := Q||\gamma.\mathsf{id}$. Let $\nu^P := \gamma.\mathsf{cspace}(cid)$ and $\nu^Q := \Gamma^Q(id).\mathsf{cspace}(cid)$ and $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise. Let τ_0 be the current round.
2. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message (execute, $id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \overset{\tau_0}{\hookleftarrow} P$, in round $\tau_1 \leq \tau_0 + \mathrm{TER}_{sub}$ send the message (execute–requested, $id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P))$ to I and once the internal execution is completed (before round $\tau_0 + \mathrm{TE}_{sub}$), forward the result to I.
3. If (execute, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \overset{\tau_1}{\hookleftarrow} I$, then
   (a) Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \overset{\tau_1}{\hookleftarrow} I$ and once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to I.
   (b) In round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ mark (id, cid) as registered in Γ^Q_{aux} and update the channel space $\Gamma^Q := \texttt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$.
   (c) In round τ_2 internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \overset{\tau_2}{\hookleftarrow} Q$, in round $\tau_3 \leq \tau_2 + \mathrm{TER}_{sub}$ send the message (execute–requested, $id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q))$ to I and once the internal execution is completed (before round $\tau_2 + \mathrm{TE}_{sub}$), forward the result to I.
   (d) If (execute, $id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \overset{\tau_3}{\hookleftarrow} I$, then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving this message and once the internal execution is completed (before round $\tau_3 + \mathrm{TE}_{sub}$), forward the result to I. In round $\tau_4 \leq \tau_0 + 3 \cdot \mathrm{TE}_{sub}$ mark (id, cid) as registered in Γ^P_{aux}, update the channel space $\Gamma^P := \texttt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$ and stop.
4. In round $\tau_0 + 4 \cdot \mathrm{TER}_{sub}$ mark (id, cid) as registered in Γ^P_{aux} and stop.

Case I, Q are honest and P is corrupt:

Upon (execute, $id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \overset{\tau_0}{\hookleftarrow} P$, such that $\alpha \neq \bot$ for $\alpha := \Gamma(id_P)$, $\nu \neq \bot$ for $\nu := \alpha.\mathsf{cspace}(cid_P)$, $\nu.\mathsf{code} = \texttt{VSCC}_i(\mathcal{C})$, $\nu.\mathsf{storage}.\mathsf{cspace}(cid') = \bot$ for every $cid' \in \{0,1\}^*$ and $\texttt{VerifyInstance}(id, cid, \nu^P) = 1$, proceed as follows

1. Set $\gamma := \nu.\mathsf{storage}.\mathsf{virtual\text{–}channel}$, $Q := \gamma.\mathsf{other\text{-}party}(P)$, $id_Q := \gamma.\mathsf{subchan}(Q)$, $cid_Q := Q||\gamma.\mathsf{id}$ and $\nu^Q := \Gamma^Q(\gamma.\mathsf{id}).\mathsf{cspace}(cid)$. If $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$, then set $\tilde{\nu} := \nu^P$, otherwise let $\tilde{\nu} := \nu^Q$.

2. In round τ_0, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$ and once the internal execution is completed (before round $\tau_0 + \mathrm{TE}_{sub}$), forward the result to P.
3. In round $\tau_1 \leq \tau_0 + \mathrm{TER}_{sub}$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1} I$.
4. In round $\tau_2 \leq \tau_1 + 4$, mark (id, cid) as registered in Γ^Q_{aux}, update the channel space $\Gamma^Q := \texttt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$ and then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_2} Q$.
5. In round $\tau_3 \leq \tau_2 + 4$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_3} I$, in round $\tau_4 \leq \tau_3 + \mathrm{TER}_{sub}$ send the message $(\text{execute–requested}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookrightarrow{\tau_4} P$ and once the internal execution is completed (before round $\tau_3 + \mathrm{TE}_{sub}$), forward the result to P.

Case I is honest and P, Q are corrupt:

Upon $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$, such that $\alpha \neq \perp$ for $\alpha := \Gamma(id_P)$, $\nu \neq \perp$ for $\nu := \alpha.\mathsf{cspace}(cid_P)$, $\nu.\mathsf{code} = \texttt{VSCC}_i(\mathcal{C})$, $\nu.\mathsf{storage}.\mathsf{cspace}(cid') = \perp$ for every $cid' \in \{0,1\}^*$ and $\texttt{VerifyInstance}(id, cid, \nu^P) = 1$, proceed as follows

1. Set $\gamma := \nu.\mathsf{storage}.\mathsf{virtual–channel}$, $Q := \gamma.\mathsf{other\text{-}party}(P)$, $id_Q := \gamma.\mathsf{subchan}(Q)$, $cid_Q := Q || \gamma.\mathsf{id}$.
2. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$ and once the internal execution is completed (before round $\tau_0 + \mathrm{TE}_{sub}$), forward the result to P.
3. In round $\tau_1 \leq \tau_0 + \mathrm{TE}_{sub}$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookleftarrow{\tau_1} I$, in round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ send the message $(\text{execute–requested}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^P)) \xhookrightarrow{\tau_2} Q$ and once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to Q.
4. If not $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_2} Q$, where $\texttt{VerifyInstance}(id, cid, \nu^Q) = 1$, then set $\tilde{\nu} := \nu^P$ and goto step 8.
5. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_2} Q$ and once the internal execution is completed (before round $\tau_2 + \mathrm{TE}_{sub}$), forward the result to Q.
6. In round $\tau_3 := \tau_2 + \mathrm{TER}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookleftarrow{\tau_3} I$, in round $\tau_4 \leq \tau_3 + \mathrm{TER}_{sub}$ send the message $(\text{execute–requested}, id_P, cid_P, \texttt{RegisterInstance}_i^{\mathcal{C}}, (cid, \nu^Q)) \xhookrightarrow{\tau_4} P$ and once the internal execution is completed (before round $\tau_3 + \mathrm{TE}_{sub}$), forward the result to P.
7. Let $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise.
8. Let τ_5 be the current round. Send $(\text{update}, id, cid, \tilde{\nu}.\mathsf{storage}, \tilde{\nu}.\mathsf{code}) \xhookrightarrow{\tau_5} \mathcal{F}_{ch}(i)$ on behalf of P and $(\text{update–reply}, ok, id, cid) \xhookrightarrow{\tau_5+1} \mathcal{F}_{ch}(i)$ on behalf of Q.

Subsimulator: $\texttt{SimRegister}_i(P, id, cid)$

Case Q is honest and P, I are corrupt:

1. If $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^P)) \xhookleftarrow{\tau_0} I$, such that $\beta \neq \bot$ for $\beta := \Gamma(id_Q)$, $\nu \neq \bot$ for $\nu := \beta.\mathsf{cspace}(cid_Q)$, $\nu.\mathsf{code} = \texttt{VSCC}_i(\mathcal{C})$, $\nu.\mathsf{storage.cspace}(cid') = \bot$ for every $cid' \in \{0,1\}^*$ and $\texttt{VerifyInstance}(id, cid, \nu^P) = 1$, then proceed. Otherwise stop.
2. Set $\gamma := \nu.\mathsf{storage.virtual{-}channel}$, $Q := \gamma.\mathsf{other\text{-}party}(P)$. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^P)) \xhookleftarrow{\tau_0} I$ and once the internal execution is completed (before round $\tau_0 + \mathrm{TE}_{sub}$), forward the result to I.
3. In round $\tau_1 \leq \tau_0 + \mathrm{TER}_{sub}$ let $\nu^Q := \Gamma^Q(\gamma.\mathsf{id}).\mathsf{cspace}(cid)$ and set $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise. Mark (id, cid) as registered in Γ^Q_{aux} and update the channel space $\Gamma^Q := \texttt{UpdateChanSpace}^*(\Gamma^Q, id, cid, \tilde{\nu})$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^Q)) \xhookleftarrow{\tau_1} I$, in round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ send the message $(\text{execute–requested}, id_Q, cid_Q, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^Q)) \xhookrightarrow{\tau_2} I$ and once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to I.

$\boxed{\textbf{Case } P \textbf{ is honest and } Q, I \textbf{ are corrupt:}}$

1. Let $\gamma := \Gamma^P(id)$, $id_P := \gamma.\mathsf{subchan}(P)$, $cid_P := P||\gamma.\mathsf{id}$ and $\nu^P := \gamma.\mathsf{cspace}(cid)$.
2. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^P)) \xhookleftarrow{\tau_0} P$, in round $\tau_1 \leq \tau_0 + \mathrm{TER}_{sub}$ send the message $(\text{execute–requested}, id_P, cid_P, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^P)) \xhookrightarrow{\tau_1} I$ and once the internal execution is completed (before round $\tau_0 + \mathrm{TE}_{sub}$), forward the result to I.
3. Then distinguish the following two situations
   - If $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^Q)) \xhookleftarrow{\tau_1 \leq \tau_0 + 3\cdot\mathrm{TER}_{sub}} I$ and $\texttt{VerifyInstance}(id, cid, \nu^Q) = 1$, then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \texttt{RegisterInstance}^{\mathcal{C}}_{i}, (cid, \nu^Q)) \xhookleftarrow{\tau_1} I$ and once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to I. Set $\tilde{\nu} := \nu^P$ if $\nu^P.\mathsf{version} \geq \nu^Q.\mathsf{version}$ and $\tilde{\nu} := \nu^Q$ otherwise. In round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ mark (id, cid) as registered in Γ^P_{aux} and update the channel space $\Gamma^P := \texttt{UpdateChanSpace}^*(\Gamma^P, id, cid, \tilde{\nu})$
   - Else, in round $\tau_2 := \tau_0 + 4 \cdot \mathrm{TE}_{sub}$, mark (id, cid) as registered in Γ^P_{aux} and stop.

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. In addition, we define an auxiliary procedure $\texttt{SimLocalExe}$ whose formal description can be found at the end of this simulator. Let $\mathrm{TER}_{sub} := \mathrm{TimeExeReq}(\lceil i/2 \rceil)$, $\mathrm{TE}_{sub} := \mathrm{TimeExe}(\lceil i/2 \rceil)$, $\mathcal{F}_{ch}(i) := \mathcal{F}^{\hat{\mathcal{L}}(\Delta)}_{ch}(i, \mathcal{C})$ and $\mathcal{F}_{ch}(i-1) := \mathcal{F}^{\hat{\mathcal{L}}(\Delta)}_{ch}(i-1, \texttt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$.

$\boxed{\textbf{Case } P \textbf{ and } I \textbf{ are honest and } Q \textbf{ is corrupt:}}$

Upon $(P, \text{execute}, id, cid, f, z) \xhookleftarrow{\tau_0} \mathcal{F}_{ch}$, let $\gamma^P := \Gamma^P(id)$, $\nu^P := \gamma^P.\mathsf{cspace}(cid)$, $\sigma^P := \nu^P.\mathsf{storage}$ and $w^P := \Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version}$. In addition, set $\tau_1 := \tau_0 + x$, where x is the smallest offset such that $\tau_1 = 1 \mod 4$ if $P = \gamma^P.\mathsf{Alice}$ and $\tau_1 = 3 \mod 4$ if $P = \gamma^P.\mathsf{Bob}$. Wait until round τ_1 and then proceed as follows:

1. If (id, cid) is not marked as corrupt in Γ^P_{aux}, do:
   (a) Set $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^P, P, \tau_0, z)$. If $m = \bot$, then stop. Else compute $s_P := \texttt{Sign}_{sk_P}(id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, w^P)$ and send $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \xhookrightarrow{\tau_1+1} Q$ and instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the execute requested message.
   (b) If $(\text{peaceful–confirm}, id, cid, f, z, s_Q) \xhookleftarrow{\tau_1+1} Q$ such that $\texttt{Vfy}_{pk_Q}(id, cid, \tilde{\sigma}, \nu^P.\mathsf{code}, w^P; s_Q) = 1$, then set $\Gamma^P_{aux}(id, cid).\mathsf{next\text{-}version} := w^P + 1$ and $\Gamma^P := \texttt{UpdateChanSpace}(\Gamma^P, id, cid, \tilde{\sigma},$

$\nu^P.\mathsf{code}, add_L, add_R, w^P, \{s_P, s_Q\})$ and instruct the ideal functionality to output the result and stop. Else mark (id, cid) as corrupt in Γ^P_{aux} and goto step 2.

2. Let τ_3 be the current round ($\tau_3 \leq \tau_0 + 5$), let $id_P := \gamma^P.\mathsf{subchan}(P)$, $cid_P := P||id$, $s_n := \mathtt{Sign}_{sk_P}(cid, P, \tau_0, f, z)$ and $p_n := (P, \tau_0, f, z, s_n)$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_P, cid_P, \mathtt{ExecuteInstance}^{\mathcal{C}}_i, (cid, p_n)) \overset{\tau_3}{\hookleftarrow} P$ and add (f, P, z, τ_0) to the set $\Gamma^P_{aux}(id, cid).\mathsf{toExecute}$. If (id, cid) is not marked as registered in Γ^P_{aux}, then run in parallel the subprocedure $\mathtt{SimRegister}_i(P, id, cid)$.
3. In round $\tau_4 \leq \tau_3 + 4$ internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving the message (execute, id_Q, $cid_Q, \mathtt{ExecuteInstance}^{\mathcal{C}}_i, (cid, p_n)) \overset{\tau_4}{\hookleftarrow} I$, in round $\tau_5 \leq \tau_4 + \mathrm{TER}_{sub}$ send the message (execute–requested, $id_Q, cid_Q, \mathtt{ExecuteInstance}^{\mathcal{C}}_i, (cid, p_n)) \overset{\tau_5}{\hookrightarrow} Q$ and instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the execution requested message. Once the internal execution is completed (before round $\tau_4 + \mathrm{TE}_{sub}$) forward the result to Q.
4. Wait until round $\tau_5 := \tau_0 + 4 \cdot \mathrm{TER}_{sub} + 5$. In order to prevent double execution, first check if $\Gamma^P(id).\mathsf{cspace}(cid).\mathsf{storage} = \tilde{\sigma}$. If this is the case (i.e. Q registered the contract instance version after execution), then delete (f, P, z, τ_0) from $\Gamma_{aux}(id, cid).\mathsf{toExecute}$ and instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the result. Otherwise execute $P.\mathtt{SimLocalExe}(\tau_5, id, cid, \tau_0)$.

Case P honest and I and Q are corrupt:

Upon $(P, \text{execute}, id, cid, f, z) \overset{\tau_0}{\hookleftarrow} \mathcal{F}_{ch}$, make the same initialization as in the case when P and I are honest.

1. Same as in the case when P and I are honest.
2. Let τ_3 be the current round ($\tau_3 \leq \tau_0 + 5$), let $id_P := \gamma^P.\mathsf{subchan}(P)$, $cid_P := P||id$, $s_n := \mathtt{Sign}_{sk_P}(cid, P, \tau_0, f, z)$ and $p_n := (P, \tau_0, f, z, s_n)$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, $id_P, cid_P, \mathtt{ExecuteInstance}^{\mathcal{C}}_i, (cid, p_n)) \overset{\tau_3}{\hookleftarrow} P$ and add (f, P, z, τ_0) to the set $\Gamma^P_{aux}(id, cid).\mathsf{toExecute}$. If (id, cid) is not marked as registered in Γ^P_{aux}, then run in parallel the subprocedure $\mathtt{SimRegister}_i(P, id, cid)$.
3. In round $\tau_4 \leq \tau_3 + \mathrm{TER}_{sub}$ send the message (execute–requested, $id_P, cid_P, \mathtt{ExecuteInstance}^{\mathcal{C}}_i, (cid, p_n)) \overset{\tau_4}{\hookrightarrow} I$ and once the internal execution is completed (before round $\tau_3 + \mathrm{TE}_{sub}$) forward the result to I.
4. Same as in the case when P and I are honest.

Simulator Sim_i: Contract instance execution

Case Q and I are honest and P is corrupt:

Upon (peaceful–request, $id, cid, f, z, s_P, \tau_0) \overset{\tau_1}{\hookleftarrow} P$

1. Let $\gamma^Q := \Gamma^Q(id)$, $\nu^Q := \gamma^Q.\mathsf{cspace}(cid)$, $\sigma^Q := \nu^Q.\mathsf{storage}$, $w^Q := \Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version}$. If $\gamma^Q = \perp$ or $P \notin \gamma^Q.\mathsf{end\text{–}users}$ or $\nu^Q = \perp$ or $f \notin \nu^Q.\mathsf{code}$, then goto step 4.
2. If $P = \gamma^Q.\mathsf{Alice}$ and $\tau_1 \mod 4 \neq 1$ or if $P = \gamma.\mathsf{Bob}$ and $\tau_1 \mod 4 \neq 3$, then goto step 4.
3. If (id, cid) is not marked as corrupt in Γ^Q_{aux}, do:
   (a) Compute $(\tilde{\sigma}, add_L, add_R, m) := f(\sigma^Q, P, \tau_0, z)$.
   (b) If $m = \perp$ or $\mathtt{Vfy}_{pk_P}(id, cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, w^Q; s_P) \neq 1$, then goto step 4.
   (c) Send $(\text{execute}, id, cid, f, z) \overset{\tau_1}{\hookrightarrow} \mathcal{F}_{ch}(i)$ on behalf of P, instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to set $\tau := \tau_0$ and instruct the functionality to output the execution requested message.
   (d) Compute the signature $s_Q := \mathtt{Sign}_{sk_Q}(id, cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, w^Q)$, send (peaceful–confirm, id, cid, $f, z, s_Q) \xhookrightarrow{\tau_1+1} P$, set $\Gamma^Q_{aux}(id, cid).\mathsf{next\text{-}version} := w^Q + 1$ and $\Gamma^Q := \mathtt{UpdateChanSpace}(\Gamma^Q, id,$

$cid, \tilde{\sigma}, \nu^Q.\mathsf{code}, add_L, add_R)$. Then instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the result and stop.

4. Mark (id, cid) as corrupt in $\mathit{\Gamma}^Q_{aux}$ and stop.

Upon $(\text{execute}, id_P, cid_P, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_2}{\hookleftarrow} P$, let $\alpha := \mathit{\Gamma}(id_P)$, $\nu_P := \alpha.\mathsf{cspace}(cid_P)$ and $\gamma := \nu_P.\mathsf{storage.virtual\text{–}channel}$. Then set $Q := \gamma.\mathsf{other\text{–}party}(P)$, $id_Q := \gamma.\mathsf{subchan}(Q)$ and $cid_Q := Q||\gamma.\mathsf{id}$. In addition, parse $p_n := (P, \tau_0, f, z, s_n)$. If $\texttt{Vfy}_{pk_P}(cid, P, \tau_0, f, z; s_n) \neq 1$, then stop. Otherwise proceed as follows:

1. If (id, cid) is marked as registered, then goto step 2. Else if P request execution of cid_P on function $\texttt{RegisterInstance}^{\mathcal{C}}_i$ with parameters (cid, ν^P) such that $\texttt{VerifyInstance}(id, cid, \nu^P) = 1$, then execute the subprocedure $\texttt{SimRegister}_i(P, id, cid)$ and in parallel goto step 2. Otherwise stop.
2. Send $(\text{execute}, \gamma.\mathsf{id}, cid, f, z) \overset{\tau_3}{\hookrightarrow} \mathcal{F}_{ch}(i)$ on behalf of P, instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to set $\tau := \tau_0$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_2}{\hookleftarrow} P$ and once the internal execution is completed (before round $\tau_2 + \mathrm{TE}_{sub}$) forward the result to P.
3. In round $\tau_3 \leq \tau_2 + \mathrm{TER}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_3}{\hookleftarrow} I$.
4. In round $\tau_4 \leq \tau_3 + 4$, instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the execute requested message, mark (id, cid) as corrupt in $\mathit{\Gamma}^Q_{aux}$ and add (f, P, z, τ_0) to $\mathit{\Gamma}^Q_{aux}(id, cid).\mathsf{toExecute}$.
5. Wait until (id, cid) is marked as registered in $\mathit{\Gamma}^Q_{aux}$ and then execute $Q.\texttt{SimLocalExe}(\tau_4, id, cid, \tau_0)$.

Case Q honest and P and I are corrupt:

Upon $(\text{peaceful–request}, id, cid, f, z, s_P, \tau_0) \overset{\tau_1}{\hookleftarrow} P$ behave exactly as in the case when Q and I are honest.
Upon $(\text{execute}, id_Q, cid_Q, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_2}{\hookleftarrow} I$, let $\beta := \mathit{\Gamma}(id_Q)$, $\nu_Q := \alpha.\mathsf{cspace}(cid_Q)$ and $\gamma := \nu_Q.\mathsf{storage.virtual\text{–}channel}$. In addition, parse $p_n := (P, \tau_0, f, z, s_n)$. If it holds that $\texttt{Vfy}_{pk_P}(cid, P, \tau_0, f, z; s_n) \neq 1$, then stop. Otherwise proceed as follows:

1. If (id, cid) is marked as registered in $\mathit{\Gamma}^Q_{aux}$, then goto step 2. Else wait until round $\tau_2 + \mathrm{TER}_{sub}$ if I requests execution of cid_Q on function $\texttt{RegisterInstance}^{\mathcal{C}}_i$ with parameters (cid, ν^P), where $\texttt{VerifyInstance}(id, cid, \nu^P) = 1$. If so, goto step 2. Else stop.
2. Let τ_3 be the current round. Send $(\text{execute}, \gamma.\mathsf{id}, cid, f, z) \overset{\tau_3}{\hookrightarrow} \mathcal{F}_{ch}(i)$ on behalf of P and instruct $\mathcal{F}_{ch}(i)$ to set $\tau := \tau_0$.
3. Internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_3}{\hookleftarrow} I$ and once the internal execution is completed (before round $\tau_3 + \mathrm{TE}_{sub}$) forward the result to I.
4. In round $\tau_4 \leq \tau_3 + \mathrm{TER}_{sub}$, instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the execute requested message, mark (id, cid) as corrupt in $\mathit{\Gamma}^Q_{aux}$ and add (f, P, z, τ_0) to $\mathit{\Gamma}^Q_{aux}(id, cid).\mathsf{toExecute}$.
5. Wait until (id, cid) is marked as registered in $\mathit{\Gamma}^Q_{aux}$ and then execute $Q.\texttt{SimLocalExe}(\tau_4, id, cid, \tau_0)$.

Case P and Q are corrupt and I is honest:

Internally simulate the communication of the corrupt parties. If P starting the registration procedure for id, cid, then execute the sub-simulator $\texttt{SimRegister}(P, id, cid)$ for the case when both P and Q are corrupt and I is honest. Note that if the registration procedure is successful (a contract instance gets registered), the subsimulator $\texttt{SimRegister}_i$ instructs the ideal functionality to update the contract instance accordingly. Upon $(\text{execute}, id_P, cid_P, \texttt{ExecuteInstance}, (cid, p_n)) \overset{\tau_2}{\hookleftarrow} P$ proceed as follows:

1. Let $\alpha := \Gamma(id_P)$, $\nu_P := \alpha.\mathsf{cspace}(cid_P)$, $\gamma := \nu_P.\mathsf{storage.virtual\text{–}channel}$, $Q := \gamma.\mathsf{other\text{–}party}(P)$, $id_Q := \gamma.\mathsf{subchan}(Q)$ and $cid_Q := Q||\gamma.\mathsf{id}$. In addition, parse $p_n := (P, \tau_0, f, z, s_n)$. If $\mathtt{Vfy}_{pk_P}(cid, P, \tau_0, f, z; s_n) \neq 1$, then stop. If cid was never registered before and P did not request execution of cid_P on function $\mathtt{RegisterInstance}_i^{\mathcal{C}}$ with parameters (cid, ν^P), where $\mathtt{VerifyInstance}(id, cid, \nu^P) = 1$, then stop. Otherwise proceed as follows.
2. Send $(\text{execute}, id, cid, f, z) \overset{\tau_2}{\hookrightarrow} \mathcal{F}_{ch}(i)$ on behalf of P, instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to set $\tau := \tau_0$. Then internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_P, cid_P, \mathtt{ExecuteInstance}, (cid, p_n)) \overset{\tau_2}{\hookleftarrow} P$ and once the internal execution is completed (before round $\tau_2 + \text{TE}_{sub}$) forward the result to P.
3. In round $\tau_3 \leq \tau_2 + \text{TER}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_Q, cid_Q, \mathtt{ExecuteInstance}, (cid, p_n)) \overset{\tau_3}{\hookleftarrow} I$, in round $\tau_4 \leq \tau_3 + \text{TER}_{sub}$ send the message $(\text{execute–requested}, id_Q, cid_Q, \mathtt{ExecuteInstance}_i^{\mathcal{C}}, (cid, p_n)) \overset{\tau_4}{\hookrightarrow} Q$ and once the internal execution is completed (before round $\tau_4 + \text{TE}_{sub}$) forward the result to Q. Instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the execute requested message.
4. Wait until the registration procedure is completed and then instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to output the result of execution.

Auxiliary procedure: $T.\mathtt{SimLocalExe}(\tau, id, cid, \tau_0)$

1. Let $\gamma := \Gamma^T(id)$ and $\sigma^{(0)} := \gamma.\mathsf{cspace}(cid).\mathsf{storage}$.
2. Let $E \subseteq \Gamma^T_{aux}(id, cid).\mathsf{toExecute}$ consist of all tuples (f', T', z', τ_0'), where $\tau_0' \leq \tau_0$.
   - If $T^{(i)} \neq T^{(j)}$, then $T^{(i)} = A$ and $T^{(j)} = B$.
   - If $T^{(i)} = T^{(j)}$, then either $f_i <_{\mathtt{C}} f_j$, where $<_{\mathtt{C}}$ is total ordering of the contract functions defined by the contract code $\mathtt{C}$, or $f_i = f_j$ and $z_n^{(i)} \leq_{\text{lex}} z_n^{(j)}$, where $\leq_{\text{lex}}$ is the lexicographic ordering of binary strings.
3. For $k = 1$ to ℓ
   (a) Compute $(\sigma^{(k)}, add_L^{(k)}, add_R^{(k)}, m^{(k)}) := f(\sigma^{(k-1)}, T^{(k)}, \tau_0^{(k)}, z^{(k)})$.
   (b) Instruct the ideal functionality $\mathcal{F}_{ch}(i)$ to execute and output the result of execution of $e^{(k)}$.
   (c) Set $\Gamma^T := \mathtt{UpdateChanSpace}(\Gamma^T, id, cid, \sigma^{(k)}, \mathtt{C}, add_L^{(k)}, add_R^{(k)})$, where $\mathtt{C} := \gamma.\mathsf{cspace}(cid).\mathsf{code}$.
   (d) Delete $e^{(k)}$ from $\Gamma^T_{aux}(id, cid).\mathsf{toExecute}$.

Simulator Sim_i: Closing a virtual state channel

We use the abbreviated notation from Sec. 4.1 and Sec. 6.1. Let $\text{TER}_{sub} := \text{TimeExeReq}(\lceil i/2 \rceil)$, $\text{TE}_{sub} := \text{TimeExe}(\lceil i/2 \rceil)$, $\mathcal{F}_{ch}(i) := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i, \mathcal{C})$ and $\mathcal{F}_{ch}(i-1) := \mathcal{F}_{ch}^{\hat{\mathcal{L}}(\Delta)}(i-1, \mathtt{VSCC}_i(\mathcal{C}) \cup \mathcal{C})$.

Case A, B, I are honest

Let γ the virtual state channel to be closed. In round $\gamma.\mathsf{validity}$ proceed as follows for both $T \in \{A, B\}$.

1. Set $id := \gamma.\mathsf{id}$, $id_T := \gamma.\mathsf{subchan}(\gamma.\mathsf{id})$, $cid_T := T||\gamma.\mathsf{id}$.
2. If there is cid such that $\Gamma^T(id).\mathsf{cspace}(cid) \neq \perp$ and (id, cid) is not marked as registered in Γ^T_{aux}, then run $\mathtt{SimRegister}_i(T, id, cid)$.
3. In round $\gamma.\mathsf{validity} + \text{TER}_{sub} + \text{TE}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving $(\text{execute}, id_T, cid_T, \mathtt{Close}_i^{\mathcal{C}}, \emptyset) \hookleftarrow T$.
4. After the internal execution, set $\Gamma^T(id) := \perp$.

Case A, B are honest and I is corrupt

In round $\gamma.\mathsf{validity}$ for both $T \in \{A, B\}$ proceed as follows.

1. Set $id := \gamma.\mathsf{id}$, $id_T := \gamma.\mathsf{subchan}(T)$, $cid_T := T||id$.
2. If $\Gamma^T(id) = \bot$ and $\Gamma(id_T).\mathsf{cspace}(cid_T) = \bot$, then stop.
3. If $\Gamma^T(id) = \bot$ but $\Gamma(id_T).\mathsf{cspace}(cid_T) \neq \bot$, then goto step 5.
4. If $\Gamma^T(id) \neq \bot$ and there is cid such that $\Gamma^T(id).\mathsf{cspace}(cid) \neq \bot$ and (id, cid) is not marked as registered in Γ^T_{aux}, then run the subsimulator $\mathtt{SimRegister}_i(T, id, cid)$.
5. In $\tau_1 := \gamma.\mathsf{validity} + \mathrm{TER}_{sub} + \mathrm{TE}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ on receiving (execute, id_T, cid_T, $\mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookleftarrow{\tau_1} T$ and in round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ send the message (execute–requested, id_T, cid_T, $\mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookrightarrow{\tau_2} I$. Once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to I and set $\Gamma^T(id) := \bot$.

Case A, I are honest and B is corrupt

In round $\gamma.\mathsf{validity}$, set $id := \gamma.\mathsf{id}$, $id_A := \gamma.\mathsf{subchan}(A)$, $cid_A := A||id$, $id_B := \gamma.\mathsf{subchan}(B)$, $cid_B := B||id$. Then proceed as follows.

1. If $\Gamma^A(id) = \bot$, then stop.
2. If $\Gamma^A(id) \neq \bot$, and there is cid such that $\Gamma^A(id).\mathsf{cspace}(cid) \neq \bot$ and (id, cid) is not marked as registered in Γ^A_{aux}, then run the subsimulator $\mathtt{SimRegister}_i(A, id, cid)$.
3. In round $\tau_1 := \gamma.\mathsf{validity} + \mathrm{TER}_{sub} + \mathrm{TE}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, id_A, $cid_A, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookleftarrow{\tau_1} A$. After the internal execution is completed, set $\Gamma^A(\gamma.\mathsf{id}) := \bot$.
4. If (execute, $id_B, cid_B, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookleftarrow{\tau_1} B$, then internally simulate the functionality $\mathcal{F}_{ch}(i-1)$ upon receiving this message and after the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to B.
5. Otherwise (i.e. if B does not initiate the execution of cid_B in round τ_1), then in round $\tau_2 := \tau_1 + \mathrm{TE}_{sub}$ internally simulate the functionality $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, id_B, $cid_B, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookleftarrow{\tau_2} I$, in round $\tau_3 := \tau_2 + \mathrm{TER}_{sub}$ send the message (execute–requested, $id_B, cid_B, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookrightarrow{\tau_3} B$ and once the internal execution is completed (before round $\tau_2 + \mathrm{TE}_{sub}$), forward the result to B.

Case B, I are honest and A is corrupt

Analogous to the previous case.

Case A is honest and I, B are corrupt

In round $\gamma.\mathsf{validity}$, set Set $id := \gamma.\mathsf{id}$, $id_A := \gamma.\mathsf{subchan}(A)$, $cid_A := A||id$ and then proceed as follows.

1. If $\Gamma^A(id) = \bot$ and $\Gamma(id_A).\mathsf{cspace}(cid_A) = \bot$, then stop.
2. If $\Gamma^A(id) = \bot$ but $\Gamma(id_A).\mathsf{cspace}(cid_A) \neq \bot$, then goto step 4.
3. If $\Gamma^A(id) \neq \bot$, and there is cid such that $\Gamma^A(id).\mathsf{cspace}(cid) \neq \bot$ and (id, cid) is not marked as registered in Γ^A_{aux}, then run the subsimulator $\mathtt{SimRegister}_i(A, id, cid)$.
4. In round $\tau_1 := \gamma.\mathsf{validity} + \mathrm{TER}_{sub} + \mathrm{TE}_{sub}$, internally simulate $\mathcal{F}_{ch}(i-1)$ upon receiving (execute, id_A, $cid_A, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookleftarrow{\tau_1} A$ and in round $\tau_2 \leq \tau_1 + \mathrm{TER}_{sub}$ send the message (execute–requested, id_A, $cid_A, \mathtt{Close}^{\mathcal{C}}_i, \emptyset) \xhookrightarrow{\tau_2} I$. Once the internal execution is completed (before round $\tau_1 + \mathrm{TE}_{sub}$), forward the result to I and set $\Gamma^A(id) := \bot$.

Case B is honest and I, A are corrupt

Analogous to the previous case.