

# Trustless Cross-chain Communication for Zendoo Sidechains

Alberto Garoffolo  
alberto@horizen.global  
Horizen

Dmytro Kaidalov  
dmytro.kaidalov@iohk.io  
Input Output

Roman Oliynykov  
roman.oliynykov@iohk.io  
Input Output

V.N.Karazin Kharkiv National University

September 2022

## Abstract

In the *Zendoo* white paper [7, 1], we introduced a novel sidechain construction for Bitcoin-like blockchains, which allows a mainchain to create and communicate with sidechains of different types without knowing their internal structure. In this paper, we take a step further by introducing a comprehensive method for sidechains to communicate amongst each other. We will also discuss the details of a cross-chain token transfer protocol that extends the generic communication mechanism. With the cross-chain token transfer protocol, it can enable a broad range of new applications, such as an *exchange platform*, that allows the ability to trade tokens issued from different sidechains.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Zendoo Overview</b>	<b>3</b>
2.1	Ceased sidechain . . . . .	6
2.2	Sidechain transactions commitment . . . . .	7
<b>3</b>	<b>Cross-Sidechain Communication Protocol</b>	<b>8</b>
3.1	Message passing . . . . .	8
3.1.1	Sending a message . . . . .	10
3.1.2	Committing a message . . . . .	10
3.1.3	Receiving a message . . . . .	10
3.2	Message passing from a ceased sidechain . . . . .	11
<b>4</b>	<b>Mitto Token Transfer Protocol</b>	<b>12</b>
4.1	Tokens data structures and state . . . . .	12
4.2	Sending tokens . . . . .	15
4.3	Receiving tokens . . . . .	17
4.4	Transferring tokens from a ceased sidechain . . . . .	19
4.4.1	Withdrawing native tokens . . . . .	20
4.4.2	Withdrawing foreign tokens . . . . .	20
4.5	Design rationale . . . . .	20
4.5.1	Token wrapping . . . . .	21
<b>5</b>	<b>Conclusions</b>	<b>21</b>

# 1 Introduction

Since the appearance of Bitcoin [11], blockchain technology has gained great attention in both industry and academia. Experts from various fields have started exploring and building new platforms and applications. The distinguishing feature of Bitcoin - the absence of a centralized control - is perceived to be disruptive to the existing financial systems, thereby making them more robust, fair, and transparent. Bitcoin inspired the development of many other similar crypto-platforms with a variety of different features.

The adoption of Bitcoin and other cryptocurrencies grew significantly during the previous decade, which also exposed their limitations such as limited throughput, increased latency, and reduced ability to scale [6]. In 2014, A. Back et al. proposed a concept of sidechains [2] that has the potential to overcome these limitations. The basic premise is to create many sub-blockchains that are interoperable with the main blockchain. They operate separately but can transact with the mainchain native asset. In this way, blockchain systems like Bitcoin can extend their functionality implemented in a sidechain (e.g., introduce smart contracts [9]).

In 2020, we introduced **Zendoo** [7, 1] - a universal construction for blockchain systems that allows users to create and communicate with sidechains of different types without knowing their internal structure. Moreover, we proposed a specific sidechain construction, named **Latus**, that can be built on top of this infrastructure and can realize a decentralized, verifiable blockchain system. It leverages the use of zk-SNARKs [3, 8] to generate succinct proofs of sidechain state progression [10, 4] that are used to validate cross-chain transfers of assets between the mainchain and sidechain.

In this paper, we discuss how sidechains can interoperate with one another. The communication enhances advanced functionalities and enables a plethora of new applications. The development of cross-blockchain communication involves much complexity. Therefore, it is essential to ensure the security and effectiveness of the developed system. The cross-sidechain communication protocol (CSCP) utilizes the features of Zendoo to provide security and avoid the pitfalls of many existing bridging protocols [5]. The protocol provides a way for sidechains to transfer messages to each other.

Finally, we further build on the cross-sidechain communication protocol and extend it for transferring tokens issued on different sidechains. The main idea is that tokens can be encoded as messages and then transferred in a verifiable method using the CSCP protocol. Thus, transferring of tokens allows users to utilize and trade them on different sidechains, as well as use them in other decentralized financial applications.

The discussion of the cross-sidechain communication protocol for Zendoo is organized in the following manner:

- Section 2 provides a brief overview of Zendoo and provides a detailed description of important components for the communication scheme.
- Section 3 provides a detailed description of the generic cross-sidechain communication protocol.
- Section 4 introduces the Mitto protocol that extends CSCP and specifies secure token transfers between sidechains.

## 2 Zendoo Overview

The Zendoo construction was introduced in [7]. It allows the deployment of customizable sidechains attached to the Bitcoin-like mainchain. Here we provide a general overview and

discuss the portions relevant to the cross-sidechain communication protocol. For a full discussion of Zendo, see the white paper [1].

The main feature of Zendo is that the mainchain does not know the internal structure of the sidechains. This feature allows it to build a variety of different systems by leveraging the secure and standardized cross-chain communication.

Zendo considers the parent-child relationship between the mainchain and sidechains, where sidechain nodes directly observe the mainchain, while mainchain nodes only observe cryptographically authenticated certificates from sidechain maintainers. Certificate authentication and validation are achieved by using SNARKs [3, 8], which enable constant-sized proofs of arbitrary computations. The main feature of the construction is that sidechains are allowed to define their own SNARKs, thereby establishing their own rules for authentication and validation. All SNARK proofs comply with the same verification interface used by the mainchain which enables universality as the sidechain can use an arbitrary protocol for authenticating its certificates. This basic concept is shown in Fig. 2.1.

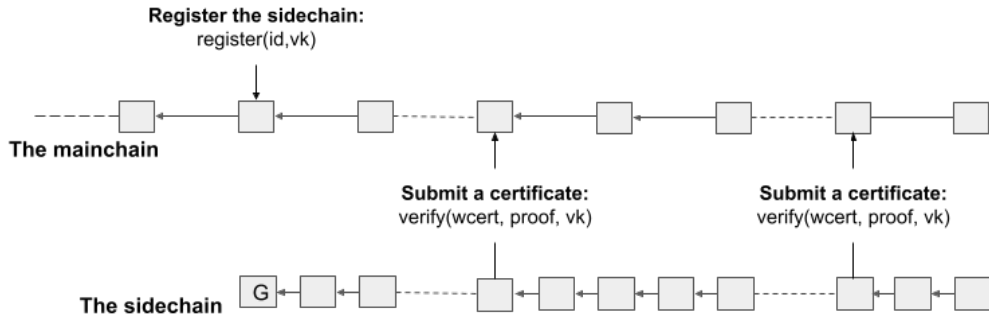


Figure 2.1: A simplified communication scheme between a sidechain and the mainchain.

A sidechain is registered to the mainchain by creating a special transaction, which sets the new sidechain id, its initial state, and more importantly, the SNARK verification key. This is used to validate the certificates. The mainchain expects certificates to be submitted in strict intervals called **withdrawal epochs**. The length of a withdrawal epoch is also defined in the sidechain creation transaction.

The basic structure of the withdrawal certificate is defined by the mainchain and this structure must be followed by all sidechains.

**Definition 2.1. Withdrawal Certificate (WCert).** A withdrawal certificate is a standardized posting that allows sidechains (SC) to communicate with the mainchain (MC). Its main functions are:

- delivering backward transfers to the MC
- serving as a heartbeat message and enabling the mainchain to identify the sidechain status.

WCert is represented by a tuple of the form:

$$WCert \stackrel{\text{def}}{=} (ledgerId, epochId, quality, BTList, proofdata, proof),$$

where,

- ledgerId* – an identifier of the sidechain for which WCert is created,
- epochId* – a number of a withdrawal epoch,
- quality* – an integer value that indicates the quality of this withdrawal certificate,
- BTLlist* – a list of backward transfers included in this withdrawal certificate,
- proofdata* – input data to a SNARK verifier,
- proof* – a SNARK proof.

As previously mentioned, the certificate is verified upon submission to the mainchain. The SNARK verification interface is the same for all sidechains:

$$\begin{aligned} true/false &\leftarrow \text{Verify}(vk_{WCert}, public\_input, proof), \\ public\_input &\stackrel{\text{def}}{=} (wcert\_sysdata, MH(proofdata)), \end{aligned}$$

where,

- vk<sub>WCert</sub>* – a SNARK verification key registered upon the sidechain creation;
- wcert\_sysdata* – a part of the public input, which is unified for all sidechains and enforced by the mainchain (explained further);
- proofdata* – a part of the input data defined by the sidechain and passed along the withdrawal certificate; it is a list of variables of predefined types whose semantics are not known to the mainchain;
- MH(proofdata)* – a root hash of a Merkle tree where leaves are variables from proofdata; it is essential for the SNARK to keep a list of public inputs short. Thus we combine them in a tree and pass the root hash only<sup>1</sup>;
- proof* – a SNARK proof itself submitted as a part of the certificate.

The *wcert\_sysdata* parameter plays an important role from a security standpoint. The idea is to allow the mainchain to verify the proof against some public input parameters defined by the protocol. For instance, the *BTLlist* and *quality* parameters that are part of the certificate must be verified before being used by the mainchain.

The *wcert\_sysdata* is represented by the tuple of the following form:

$$wcert\_sysdata \stackrel{\text{def}}{=} (quality, MH(BTLlist), B_{(last)}^i),$$

where,

- quality* – the quality parameter from the withdrawal certificate,
- MH(BTLlist)* – a root hash of a Merkle tree where leaves are backward transfers from the BTLlist provided within the certificate,
- H(B<sub>last</sub><sup>i</sup>)* – a block hash of the last mainchain block in the withdrawal epoch *i*.

The flexibility of the scheme comes from the ability of a sidechain to define *proofdata*. This can be an arbitrary set of arguments completely defined by the sidechain. For instance, it can contain a commitment of the entire state of the sidechain at the moment of creating the certificate. Then, the SNARK proof proves the state transition from the state committed in the previous certificate.

---

<sup>1</sup>A full payload of *proofdata* is provided during the proof generation as a witness.

## 2.1 Ceased sidechain

As previously stated, the sidechain is obliged to submit a withdrawal certificate in strict intervals called **withdrawal epochs**. If a withdrawal certificate has not been submitted during this time, the sidechain is considered **ceased** and no more withdrawal certificates for this sidechain will be accepted by the mainchain. However, the funds can still be withdrawn from the ceased sidechain by means of a **ceased sidechain withdrawal**. Since the certificates are not allowed for the ceased sidechains, it becomes the backup method to retrieve funds.

**Definition 2.2. Ceased Sidechain Withdrawal (CSW).** The CSW is an operation that allows the movement of coins from the ceased sidechain  $B$  to the original mainchain  $A$ . It is represented by a tuple of the following form:

$$CSW \stackrel{\text{def}}{=} (ledgerId, receiver, amount, nullifier, proofdata, proof),$$

where,

- $ledgerId$  – an identifier of the sidechain for which CSW is created,
- $receiver$  – an address of the receiver in the mainchain,
- $amount$  – the number of coins to be transferred,
- $nullifier$  – a unique identifier of claimed coins,
- $proofdata$  – input data to a SNARK verifier,
- $proof$  – a SNARK proof.

The CSW is submitted to the mainchain and has the same validation principle as the withdrawal certificate. It is validated by a SNARK proof defined by the sidechain. The verification key  $vk_{csw}$  for the proof is set upon sidechain registration. The syntax of the  $proofdata$  and proof are the same as for the withdrawal certificate. The basic interface of the SNARK verifier is the following:

$$\begin{aligned} true/false &\leftarrow Verify(vk_{csw}, public\_input, proof), \\ public\_input &\stackrel{\text{def}}{=} (csw\_sysdata, MH(proofdata)), \end{aligned}$$

where,

- $vk_{csw}$  – a SNARK verification key for the CSW registered upon the sidechain creation;
- $csw\_sysdata$  – a part of the public input, which is unified for all sidechains and enforced by the mainchain (explained further);
- $proofdata$  – a part of the input data defined by the sidechain and passed along the CSW; it is a list of variables of predefined types whose semantics are not known to the mainchain;
- $MH(proofdata)$  – a root hash of a Merkle tree where leaves are variables from proofdata;
- $proof$  – a SNARK proof itself submitted as a part of the CSW.

$csw\_sysdata$  is defined as:

$$csw\_sysdata \stackrel{\text{def}}{=} (H(B_w), nullifier, receiver, amount),$$

where  $H(B_w)$  is a block hash of the mainchain block where the latest withdrawal certificate for this sidechain has been submitted.

It is completely up to the sidechain to define the semantics of the *proofdata* (i.e., what type of data it contains) and verification logic. For instance, one might require proof that a user, who creates CSW, owns a UTXO holding the *amount* of coins in the state committed by the last certificate before the sidechain was ceased.

See the Zendoo white paper [1] for more information on ceased sidechain withdrawals.

## 2.2 Sidechain transactions commitment

An important concept for the cross-chain communication in Zendoo is the **Sidechain Transactions Commitment (STC)**. The STC is a special value that is inserted in every mainchain block and comprises all sidechain-related actions in the MC block for all registered sidechains. The STC value is a root hash of a Merkle tree that contains all transactions related to any sidechain (see Fig. 2.2).




Figure 2.2: Sidechain transactions commitment tree.

The root hash  $h_1$  commits to all sidechain related transactions included in the MC block. All  $SCXHash$  ( $h_{2X}$ ), where  $X$  is a sidechain identifier, are ordered by its id and commit to all transactions related to the sidechain  $X$ .  $WCertHash$  commits to the  $WCert$  for the sidechain  $X$  (if present) and only one  $WCert$  is allowed for each sidechain.  $TxsHash$  commits to different types of transactions for a particular sidechain.

The STC allows access to all sidechain-related information. Recall that the interface of SNARK verification for both  $WCert$  and  $CSW$  requires that the input to the proof contains MC block hash (Def. 2.1-2.2). This means that a SNARK proof can access the entire history of all sidechain transactions by proving that a certain STC with a certain SC transaction is included in some MC block. The connection of such MC block to the one provided in the proof can be proved recursively as the MC blocks are cryptographically chained to each other.

### 3 Cross-Sidechain Communication Protocol

The Zendoo construction defines only the communication protocol between the mainchain and a sidechain. It focuses mostly on forward and backward transfers of the native MC asset. A desirable feature is the ability of disparate sidechains to communicate with one another. It allows for the creation of powerful new functionality that leverages the benefits of multiple sidechain networks, simultaneously. However, such cross-sidechain communication can not be a part of the Zendoo protocol itself because the core concept postulates that the internal structure of the sidechains is unknown to Zendoo.

Therefore, to enable cross-sidechain communication we introduce a separate **Cross-Sidechain Communication Protocol (CSCP)** that extends Zendoo and defines how the sidechains communicate with one another. The basic idea is that sidechains that want to interact with each other must implement a special generalized interface (see Fig. 3.1).

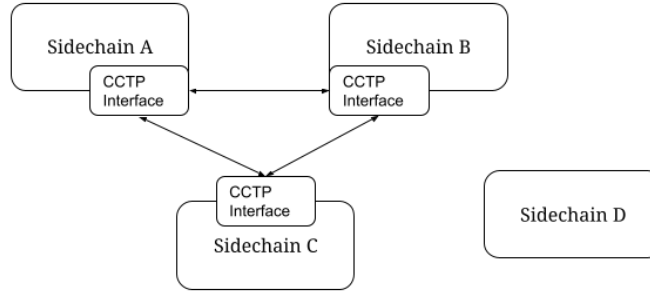


Figure 3.1: Sidechains A,B,C implement a special CCTP Interface that allows them to communicate with each other. Sidechain D does not implement the CCTP interface.

The CSCP defines an abstract mechanism of verifiable transferring of messages from one sidechain to another. It is implemented by a sidechain using customization capabilities of Zendoo. The mainchain and Zendoo itself are not affected because Zendoo already has all the necessary capabilities to create a custom logic inside the sidechains.

Then, the sidechains define particular types of messages that are exchanged along with the processing rules for them. For instance, they might implement custom token transfers based on CSCP interface (see Fig. 3.2).

#### 3.1 Message passing

In a nutshell, the CSCP protocol works the following way: given two sidechains  $SC_1$  and  $SC_2$ , if a user in  $SC_1$  wants to send a message  $M_i$  to the sidechain  $SC_2$ , he creates a transaction  $tx(M_i)$  containing  $M_i$  and submits it to  $SC_1$ . At the end of a withdrawal epoch all cross-sidechain messages are collected in a Merkle tree, the root hash of the tree becomes a part of the sidechain state that is included in the withdrawal certificate. When the certificate is confirmed in the mainchain, the state cannot be reverted. Therefore, once the state of  $SC_1$  with  $tx(M_i)$  has been committed and confirmed in the mainchain, the message can be securely confirmed in  $SC_2$ . A user in  $SC_2$  can receive a message by directly observing  $SC_1$ , extracting the message, and submitting a transaction  $tx(M_i, proof)$  to  $SC_2$  together with the proof that the message  $M_i$  has been processed in  $SC_1$  and is present in the state confirmed in the withdrawal certificate of  $SC_1$ .






Figure 3.2: A layered structure of communication protocols. Zendoo specifies the basic protocol for communication between the mainchain and a sidechain. CSCP relies on Zendoo and specifies the abstract protocol for communication amongst the sidechains. Next, the Token Transfer protocol is built on top of CSCP and specifies a concrete protocol for exchanging tokens among sidechains.

There are three main phases of the CSCP protocol:

1. **Sending** the message in the sender sidechain.
2. **Committing** of the message in the mainchain.
3. **Receiving** the message in the receiver sidechain.

The main concept is depicted in Fig. 3.3 while the following subsections provide more details about different phases of the protocol.

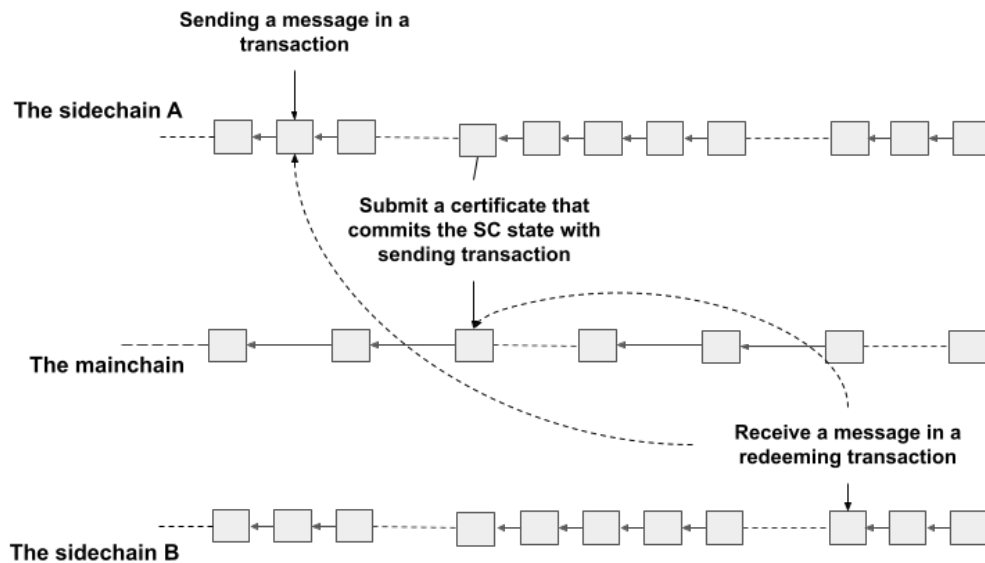


Figure 3.3: Main phases of the CSCP protocol. The receiving transaction relies on the sending transaction and commitment of the Sidechain A in the mainchain.

### 3.1.1 Sending a message

The CSCP message is defined as follows:

```
type CSCP_Message {
  sendingScId: Int,
  receivingScId: Int,
  msgType: Int,
  senderId: ByteArray,
  receiverId: ByteArray,
  payloadHash: ByteArray
}
```

where,

- sendingScId* – the id of the sending sidechain as it was registered in the mainchain;
- receivingScId* – the id of the receiving sidechain as it was registered in the mainchain;
- msgType* – the type of the message (e.g., token);
- senderId* – the unique id of the sender. The semantics of this id is not known to the CSCP protocol and is defined by the implementation of the sending sidechain (e.g., public key);
- receiverId* – the unique id of the receiver. The semantics of this id is defined by the receiving sidechain (e.g., public key);
- payloadHash* – the hash of the message itself.

The structure of the CSCP message is generalized, where the payload of the message is represented as a hash, which means that the message can be anything depending on a particular instantiation of the CSCP protocol. For instance, the message can be some entity like a token, transaction output, or a text string.

We assume that on the sending sidechain, a user can create a special transaction which contains *CSCP\_Message*. For instance, in the UTXO-based sidechain the message can be represented as a special output.

### 3.1.2 Committing a message

At the end of the withdrawal epoch in the sending sidechain, all CSCP messages are combined into a Merkle tree (Fig. 3.4).

The messages ( $msg_i$ ) are inserted in the tree in the order they are submitted to the sidechain. The tree contains messages from different users and to different sidechains. Next, the root hash of the tree is inserted in the withdrawal certificate at the end of the epoch as a custom field in the *proofdata* section. The certificate proof enforces the validity of the CSCP messages tree.

### 3.1.3 Receiving a message

To receive a message, a user on the receiving sidechain has to create a special redeem transaction. To do this, the user directly observes the sending sidechain for messages directed to him, extracts the message, and constructs a special transaction that includes both the message and the proof of its validity:

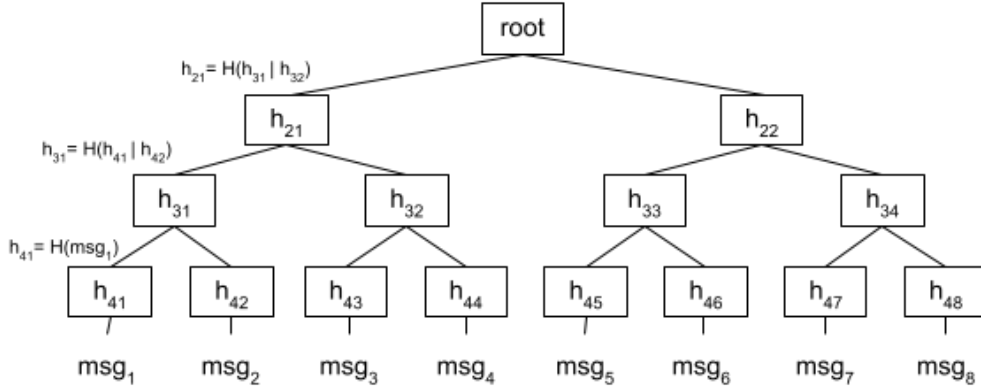


Figure 3.4: An example of the Merkle tree with CSCP messages. It contains eight messages, but the tree expands depending on the actual number of messages in a particular epoch.

```

CSCP_Redeem_Tx: {
  msg:      CSCP_Message,
  payload:  ByteArray,
  proof:    ByteArray
}

```

where *msg* is the original CSCP message from the sending sidechain, the *payload* is a data represented by *payloadHash*, and *proof* is some abstract proof that verifies the following statements about message validity:

- The message *msg* has been accepted and confirmed in the sending sidechain.
- The *payload* corresponds to *payloadHash* included in *msg*.
- The hash of a Merkle tree of CSCP messages containing *msg* is committed in the withdrawal certificate of the sending sidechain and the certificate has been confirmed in the mainchain.
- The message *msg* has not been redeemed in the receiving sidechain.
- The user who redeems the message, has rights to do so.

The CSCP protocol does not define the exact instantiation of the proof. Depending on the sidechain implementation there might be different ways to implement it. For instance, if the sidechain is based on Latus construction [1], then its blocks contain references to the MC blocks, which as discussed in section 2.2, contain *Sidechain Transaction Commitment Tree* (STC). Given the access to MC blocks, the sidechain can access STC values, which commit to all SC-related transactions in all sidechains. Therefore, it can also prove cryptographically that the certificate of the sending sidechain with the message was included in the mainchain.

The side effect of processing a received message is defined completely by the SC logic.

### 3.2 Message passing from a ceased sidechain

Under certain circumstances there might be a need to send a message from a ceased sidechain. For instance, if the sidechain has tokens that a user wants to recover in another sidechain, then a

message can be passed. The CSCP protocol defines a mechanism based on the *Ceased Sidechain Withdrawal* (recall CSW from Def. 2.2) that allows “to withdraw message” from the ceased sidechain.

In this case, the protocol works in the following way: let there be a ceased sidechain  $SC_1$  and a user  $U_1$  who owns some entity  $E_1$  (e.g., a token) that is withdrawable by CSW according to the rules of the sidechain  $SC_1$ . The CSW SNARK proof of  $SC_1$  is implemented in the way that allows verification of the validity of the *CSCP\_Message* provided in *proofdata* as a custom field. A user creates the CSW on the mainchain to withdraw the entity  $E_1$  and inserts the message  $M_1$  in *proofdata*. Once the CSW is confirmed in the mainchain, the message can be safely received in  $SC_2$ . A user in  $SC_2$  can receive a message by directly observing the mainchain, extracting the message, and submitting a transaction  $tx(M_1, proof)$  to  $SC_2$  together with the proof that the message  $M_1$  has appeared in the CSW for  $SC_1$ .

The primary objective of the CSW mechanism, as introduced in the Zendo white paper [1], is to recover MC native assets that have been sent to the sidechain. However, it can also be used for recovering abstract messages as well. In this case, the *amount* field of the CSW has to be zero, which means that there is a recovery of something other than the native asset.

## 4 Mitto Token Transfer Protocol

One of the key concepts, when considering blockchain interoperability, is the ability to transfer assets among different chains. There are many considerations, but the main use cases are:

- allow the ability to trade assets from different chains,
- enable users to borrow assets on one sidechain by providing tokens issued on another sidechain as collateral,
- enable transfer of liquidity among different sidechains.

In this section we introduce *Mitto* - an extension to the CSCP protocol that enables transferring of tokens issued on different sidechains. The basic idea is that tokens can be encoded as messages and transferred in a verifiable way using the CSCP protocol. The sent tokens are burned on the sending side, while on the receiving side they are redeemed when receiving the message.

One of the main design principles of the Mitto Token Transfer Protocol is that tokens are tracked by the issuing sidechain when they are sent. Therefore, it is impossible for a malicious sidechain to send back forged tokens that have not been previously transferred there. This also implies that the tokens can only be sent from or to the original sidechain, but not among “foreign” sidechains.

Sidechains that enable token transfer protocol must implement certain interfaces and adhere to rules of processing tokens that are specified in this section. The Mitto protocol does not rely on a specific sidechain construction and can be integrated into different sidechains given that they follow the general rules of processing tokens. The Mitto protocol can be viewed as something similar to ERC-20 or ERC-721 standards in Ethereum. For unification purposes, the protocol supports both fungible and non-fungible tokens.

### 4.1 Tokens data structures and state

The core of the *Mitto* Token Transfer Protocol is the concept of *TokenInstance*. It is an entity that is used to represent tokens inside the sidechain. The protocol does not specify how exactly

token instances are managed by the sidechain, it is only assumed that every sidechain that enables *Mitto* holds tokens as *TokenInstances* and follows certain rules when transferring them. The *TokenInstance* has the following structure:

**Definition 4.1. Token Instance.** *TokenInstance* is a data structure that represents a single non-fungible token or some amount of fungible tokens from the same set identified by a name. It is defined as follows:

```

type TokenInstance {
  tokenName:   String,
  fungibility : Boolean
  Union {
    tokenId:   Int,    // for non-fungible tokens
    amount:   Int,    // for fungible tokens
  }
  issuerSidechainID: ByteArray,
  ownerPubKey:   ByteArray,
  dataHash:      ByteArray,
}

```

where,

<i>tokenName</i>	– a unique name that identifies a set of fungible or non-fungible tokens;
<i>fungibility</i>	– indicates whether the tokens specified by <i>tokenName</i> are fungible or not. Depending on this flag, the <i>TokenInstance</i> structure contains either <i>tokenId</i> (for non-fungible tokens) or <i>amount</i> field (for fungible tokens). The <i>fungibility</i> flag is the same for all token instances with the same <i>tokenName</i> ;
<i>tokenId</i>	– if the token is non-fungible, then the <i>tokenId</i> defines a unique identifier of a particular token in the set defined by the <i>tokenName</i> . In this case <i>TokenInstance</i> represents a single token with a unique <i>tokenId</i> ;
<i>amount</i>	– if the token is fungible, then the <i>amount</i> represents an amount of fungible tokens held by the token instance;
<i>issuerSidechainID</i>	– the unique identifier of the sidechain where the tokens with <i>tokenName</i> were originally issued;
<i>ownerPubKey</i>	– the public key of the owner of a particular token instance;
<i>dataHash</i>	– the hash of other token-related data that is not relevant to the <i>Mitto</i> protocol. Recall that the <i>TokenInstance</i> is an abstract data structure which defines the basic information needed to manage tokens. A specific token type is defined by the sidechain logic and contains additional data representing the token.

The tokens can be fungible and non-fungible. A fungible token can be represented by several token instances that have the same *tokenName*, but different *amounts*. A set of non-fungible tokens can be represented by several token instances with the same *tokenName*, but different *tokenIDs*, where each token instance represents a unique non-fungible token from the set. The sidechain can issue several different sets of fungible or non-fungible tokens.

The *TokenInstance* structure basically resembles the concept of an output in a UTXO based blockchain. But in *Mitto* it is represented abstractly so that the protocol can be implemented also in a non-UTXO based blockchain.




Figure 4.1: Example of sets of fungible (wBTC) and non-fungible (Cars) tokens. **TI** is a *TokenInstance*. A sidechain may issue different sets of token instances, it also can contain token instances issued on other sidechains.

The *SentRecord* data structure is also required to manage tokens in sidechains. This data structure keeps track of tokens that are sent out.

**Definition 4.2. Sent Record.** *SentRecord* is a data structure that keeps track of the tokens that were originally issued on the sidechain but then have been sent to some other sidechain. It is represented as follows:

```

type SentRecord {
  receiverSidechainID: Int,
  tokenName:          String,
  fungibility :       Boolean
  Union {
    tokenID:   Int,    // for non-fungible tokens
    amount:    Int,    // for fungible tokens
  }
}

```

where,

- receiverSidechainID* – unique identifier of the sidechain where the tokens of type *tokenName* were sent;
- tokenName* – a unique name that identifies a set of fungible or non-fungible tokens;
- fungibility* – defines whether the tokens specified by *tokenName* are fungible or not. Depending on this flag, the *SentRecord* structure contains either *tokenId* (for non-fungible tokens) or *amount* field (for fungible tokens);
- tokenId* – if the token is non-fungible, *tokenId* is an identifier of a particular token that has been sent outside the sidechain of origin. In this case *SentRecord* represents a single token with a unique *tokenId*;
- amount* – if the token is fungible, *amount* represents an amount of fungible tokens of type *tokenName* that have been sent to sidechain *receiverSidechainID*.

The *TokenInstance* and *SentRecord* are basic structures used for token management in the *Mitto* token transfer protocol. It is assumed that a sidechain keeps track of both *TokenInstances*, that represent the tokens that a sidechain currently possesses, and *SentRecords*, that represent

the tokens that have been sent outside the sidechain where they were issued. Let formalize the sidechain state as follows:

$$SC_i^{state} = \{S_i^{tk_s}, S_i^{sent}, \dots\},$$

where,

- $S_i^{tk_s}$  is a set of token instances that are currently owned by the sidechain  $SC_i$ ;
- $S_i^{sent}$  is a set of *SentRecords* to keep track of tokens that were issued on sidechain  $SC_i$  but were sent outside it. Each *SentRecord* from the set keeps track of the amount of tokens of a particular type that were sent to a particular sidechain.

Note that  $S_i^{tk_s}$  may contain token instances representing tokens issued on other sidechains (i.e.,  $issuerSidechainID \neq SC_i$ ).

The following sections formalize the rules for transferring tokens using the model described above.

## 4.2 Sending tokens

The tokens are sent in the form of CSCP messages which wrap token instances. To send tokens from  $SC_a$  to  $SC_b$  a user creates a special sending transaction  $tx_{SC_a}^{send}(TI_t \rightarrow M_t)$  on  $SC_a$  which burns a token instance  $TI_t$  and produces an CSCP message  $M_t$  with the following data:

$$M_t: \text{CSCP\_Message} = \{$$

sendingScId	=	$SC_a$ ,
receivingScId	=	$SC_b$ ,
msgType	=	$TokenTransfer$ ,
senderId	=	$TI_t.ownerPubKey$ ,
receiverId	=	receiverPubKey,
payloadHash	=	$Hash(TI_t)$

$$\}$$

where *receiverPubKey* is the receiver's public key in  $SC_b$  specified by a user. Note that the validity of the transaction (i.e., that the user is the owner of  $TI_t$ ) is verified upon its processing by the sidechain.

We formalize the process of sending tokens by defining the validation rules for the  $tx_{SC_a}^{send}(TI_t \rightarrow M_t)$  and the side effects of it to the state of  $SC_a$  in case the transaction is valid and accepted by the sidechain  $SC_a$ .

### Validation rules

1. Verify  $TI_t \in S_a^{tkS}$  (i.e., the token instance exists in the state of  $SC_a$ ).
2. If  $TI_t.issuerSidechainId \neq SC_a$ , then verify that  $M_t.receivingScId == TI_t.issuerSidechainId$  (i.e., verify that foreign tokens are sent only to their original blockchain and not elsewhere).
3. Verify semantic validity of  $M_t$ :
  - (a)  $M_t.sendingScId == SC_a$ ;
  - (b)  $M_t.receivingScId \neq SC_a$ ;
  - (c)  $M_t.msgType == TokenTransfer$ ;
  - (d)  $M_t.senderId == TI_t.ownerPubKey$ ;
  - (e)  $M_t.payloadHash == Hash(TI_t)$ .
4. Verify that the transaction is authorized by the owner of  $TI_t$  (i.e., there is a signature on  $tx_{SC_a}^{sent}(TI_t \rightarrow M_t)$  under the public key  $TI_t.ownerPubKey$ ).

### Side effects on the state of $SC_a$

1.  $TI_t$  is removed from  $S_a^{tkS}$ ;
2. If  $TI_t.issuerSidechainId == SC_a$  (i.e., the token issued on this sidechain is being sent), then:
  - (a) If  $TI_t.fungibility == false$ , then create a sent record
$$SR_t: \text{SentRecord} = \{ \begin{array}{l} \text{receiverSidechainID} = M_t.receivingScId, \\ \text{tokenName} = TI_t.tokenName, \\ \text{fungibility} = false, \\ \text{tokenId} = TI_t.tokenId \end{array}$$
and add  $SR_t$  to  $S_a^{sent}$  (the sidechain keeps track of all tokens that were issued by it and sent elsewhere);
  - (b) If  $TI_t.fungibility == true$ , then check if  $S_a^{sent}$  already contains some sent record  $SR'_t$  such that  $SR'_t.receivingSidechainID == M_t.receivingScId$  and  $SR'_t.tokenName == TI_t.tokenName$ . If yes, just update  $SR'_t.amount = SR'_t.amount + TI_t.amount$ , otherwise create a sent record
$$SR_t: \text{SentRecord} = \{ \begin{array}{l} \text{receiverSidechainID} = M_t.receivingScId, \\ \text{tokenName} = TI_t.tokenName, \\ \text{fungibility} = false, \\ \text{amount} = TI_t.amount \end{array}$$
and add  $SR_t$  to  $S_a^{sent}$  (the sidechain keeps track of the total amount of tokens that were sent to every other sidechain).

At the end of the epoch all CSCP messages are collected in the Merkle tree and the hash of



the tree is included in the withdrawal certificate as a custom field, so that they can be received in another sidechain.

### 4.3 Receiving tokens

To receive a token from  $SC_a$  in  $SC_b$  a user creates a special redeem transaction  $tx_{SC_b}^{redeem}(M_t \rightarrow TI_t, proof)$  on  $SC_b$  that includes the message  $M_t$ , that has been previously sent from  $SC_a$ , the corresponding token instance  $TI_t$ , and a proof of the message validity. The redeem transaction is an extension of the *CSCP\_Redeem\_Tx* (see section 3.1.3). It has the following structure:

```
type Mitto_Redeem_Tx {
  M_t:  CSCP_Message,
  proof: ByteArray,
  TI_t: TokenInstance
}
```

where  $M_t$  is the message that has been previously confirmed in  $SC_a$  and that wraps the token instance  $TI_t$ . The proof proves the validity of the message.

As in the case with tokens sending, we formalize the process of receiving by defining the validation rules for the redeem transaction and the side effects to the state of  $SC_b$  in case the transaction is valid and accepted by the sidechain  $SC_b$ .

### Validation rules

1. Verify  $TI_t.issuerSidechainId == M_t.sendingScId$  or  $TI_t.issuerSidechainId == SC_b$  (it is allowed to receive tokens only from the sidechains that issued them or the tokens that were created by the receiving sidechain).
2. If  $TI_t.issuerSidechainId == SC_b$ , then verify that:
  - (a) If  $TI_t.fungibility == true$ , then there exists some sent record  $SR'_t \in S_b^{sent}$  such that:
    - i.  $SR'_t.receivingSidechainID == M_t.sendingScId$ , and
    - ii.  $SR'_t.tokenName == TI_t.tokenName$ , and
    - iii.  $SR'_t.fungibility == true$ , and
    - iv.  $SR'_t.amount \geq TI_t.amount$  (i.e., if we receiving back the native token then it should be present in  $S_b^{sent}$  and the *amount* should not be more than has been originally sent to that sidechain).If such  $SR'_t$  does not exist, reject the transaction.
  - (b) If  $TI_t.fungibility == false$ , then there should exist some sent record  $SR'_t \in S_b^{sent}$  such that:
    - i.  $SR'_t.receivingSidechainID == M_t.sendingScId$ , and
    - ii.  $SR'_t.tokenName == TI_t.tokenName$ , and
    - iii.  $SR'_t.fungibility == false$ , and
    - iv.  $SR'_t.tokenId == TI_t.tokenId$  (i.e., if we receiving back an NFT token from the *tokenName* set, then it should be present in  $S_b^{sent}$ ).If such  $SR'_t$  does not exist, reject the transaction.
3. If  $TI_t.fungibility == false$  then verify that  $TI_t \notin S_b^{tkts}$  (if it is an NFT, it can not already be present in the receiving sidechain).
4. Verify semantic validity of  $M_t$ :
  - (a)  $M_t.sendingScId \neq SC_b$ ;
  - (b)  $M_t.receivingScId == SC_b$ ;
  - (c)  $M_t.msgType == TokenTransfer$ ;
  - (d)  $M_t.senderId == TI_t.ownerPubKey$ ;
  - (e)  $M_t.payloadHash == Hash(TI_t)$ .
5. Verify that the message  $M_t$  is authorized by the sender (previous owner of  $TI_t$ ) (i.e., there is a signature on  $M_t$  under the public key  $M_t.senderId$ ).
6. Verify that the redeem transaction is authorized by the receiver (the new owner of  $TI_t$ ) on  $SC_b$  (i.e., there is a signature on  $tx_{SC_b}^{redeem}(TI_t \rightarrow M_t, proof)$  under the public key  $M_t.receiverId$ ).
7. Verify the proof, which should check the validity of the message according to CSCP rules (section 3.1.3).

Note that the *Mitto* token transfer protocol does not specify precise rules for the message

validity proof. Similarly to the CSCP protocol, we view *Mitto* as an abstract protocol that can be integrated into sidechains with different architecture, different state management, etc. That is why at this level of abstraction it is impossible to precisely define the verification logic. For instance, for the Latus-based sidechains the proof can be a SNARK proof that verifies existence of the message in the confirmed certificate of  $SC_a$ .

**Side effects on the state of  $SC_b$**

1. Let  $TI_t^{new}$  be a copy of  $TI_t$  except that  $TI_t^{new}.ownerPubKey = M_t.receiverId$ .
2.  $TI_t^{new}$  is added to  $S_b^{tk}$ .
3. If  $TI_t^{new}.issuerSidechainID == SC_b$  (i.e., we received back the token that was issued in this sidechain), then:
  - (a) If  $TI_t^{new}.fungibility == false$ , then find such  $SR_t \in S_b^{sent}$  such that  $SR_t.receivingSidechainID == M_t.sendingScId$ , and  $SR_t.tokenName == TI_t.tokenName$ , and  $SR_t.tokenId == TI_t.tokenId$  and remove it from  $S_b^{sent}$ ;
  - (b) If  $TI_t^{new}.fungibility == true$ , then find such  $SR_t \in S_b^{sent}$  such that  $SR_t.receivingSidechainID == M_t.sendingScId$  and  $SR_t.tokenName == TI_t.tokenName$ . Then update  $SR_t.amount = SR_t.amount - TI_t.amount$ . If, after update,  $SR_t.amount == 0$ , then remove  $SR_t$  from  $S_b^{sent}$ .

#### 4.4 Transferring tokens from a ceased sidechain

By transferring tokens from a ceased sidechain we understand the case when some sidechain  $SB_a$  has been ceased while holding some token instances  $TI_i \in S_a^{tk}$  or storing sent records  $SR_i \in S_a^{sent}$  for tokens issued by  $SC_a$  and sent elsewhere, and a user wants to transfer such tokens. In a nutshell, this can be done by utilizing the Ceased Sidechain Withdrawal mechanism. A user should create a CSW containing a message with the transferred token instance. The SNARK proof of the CSW should be designed in a way that allows to prove the existence of such a TI in the state of the sidechain at the moment of ceasing. Then, the token can be redeemed in the receiving sidechain by creating a special CSW redeem transaction:

```

type Mitto_CSW_Redeem_Tx {
   $M_t$ : CSCP_Message,
  proof: ByteArray,
   $TI_t$ : TokenInstance,
   $CSW_t$ : CSW
}

```

The validation rules and side effects are almost similar to the one described in section 4.3 for regular redemption. The only exception is the proof, which in this case should verify that the message has been committed in the  $CSW_t$  rather than in a withdrawal certificate. The side effects remain completely the same as in the case of *Mitto\_Redeem\_Tx*.

While receiving is relatively simple, the token sending from a ceased sidechain is more complex. We consider two basic scenarios of withdrawing tokens from a ceased sidechain:

1. **Withdrawing native tokens** (i.e., tokens that were issued on the ceased sidechain).
2. **Withdrawing foreign tokens** (i.e., tokens that were not issued on the ceased sidechain, but were held there at the moment of ceasing).

#### 4.4.1 Withdrawing native tokens

Let assume that the sidechain  $SC_a$  has ceased. There are two cases to consider:

1. At the moment of ceasing the sidechain  $SC_a$  was the owner of the token  $TI_i$  to be withdrawn. I.e., there is a token instance  $TI_i \in S_a^{tk_s}$ , where  $S_a^{tk_s}$  is a set of owned tokens at the moment of ceasing. In this case, it is needed to create a CSW that proves the existence of  $TI_i$  in  $S_a^{tk_s}$  and the right to withdraw it. The CSW is submitted to the mainchain. Once it is confirmed in the mainchain, *Mitto-CSW-Redeem-Tx* can be used to redeem the token in another sidechain.
2. At the moment of ceasing the sidechain  $SC_a$  did not hold the token to be withdrawn. I.e., there is a sent record  $SR_t \in S_a^{sent}$ , that represents a token(s) that were originally issued in  $SB_a$  but currently is owned by some sidechain  $SC_c$ . We assume that  $SC_c$  is active. Such tokens can not be directly transferred from  $SC_c$  to another active sidechain  $SC_b$ , because according to the *Mitto* protocol foreign tokens can be sent only to the sidechain where they have been originally issued. Therefore, to transfer such tokens,  $SC_c$  first has to send them back to the ceased sidechain  $SC_a$  using the standard sending flow. Then, they can be withdrawn using the CSW mechanism. The CSW should be designed to prove the following:
  - (a) there exist a valid message  $M_t$  that were sent by  $SC_c$  to  $SC_a$ , the message has been committed to the confirmed withdrawal certificate of  $SC_c$ ;
  - (b) the message  $M_t$  transfers the token instance  $TI_t$ , which represent a token(s) that were originally issued by  $SC_a$ ;
  - (c) there is a sent record  $SR_t \in S_a^{sent}$  that confirms that the token(s) have been sent to  $SC_c$  previously and that it has not been withdrawn yet.

Once such CSW is confirmed in the mainchain, *Mitto-CSW-Redeem-Tx* can be used to redeem the token in the sidechain  $SC_b$ .

#### 4.4.2 Withdrawing foreign tokens

Let assume that the sidechain  $SC_a$  has been ceased and it owns a token instance  $TI_t \in S_A^{tk_s}$ , such that  $TI_t.issuerSidechainId \neq SC_a$ . In this case it is needed to create a CSW that proves the existence of  $TI_t$  in  $S_a^{tk_s}$  and the right to withdraw it. The CSW will include a corresponding message  $M_t$  that wraps  $TI_t$  as a custom field. Then,  $TI_t$  can be redeemed with *Mitto-CSW-Redeem-Tx* in sidechain  $TI_t.issuerSidechainId$  (and only there).

### 4.5 Design rationale

One of the main design principles of the *Mitto* token transfer protocol is that tokens (in the form of token instances) can only be sent from the sidechain where they have been originally issued to some other sidechain and back. But it is disallowed to send tokens from the sidechain where they have not been issued to another sidechain, which is also not their original place of issuance.

More formally: if there is a message  $M_t$  that sends a token instance  $TI_t$  from sidechain  $SC_a$  to  $SC_b$ , then

- either  $SC_a == TI_t.issuerSidechainId$ ,
- or  $SC_b == TI_t.issuerSidechainId$ .

The reason for such a design decision is security. Recall that a sidechain that issues some set of tokens keeps track (through *SentRecord* mechanism) of tokens that have been sent outside. And when these tokens return back they are checked through *SentRecords*. In this case, it is impossible for some corrupted sidechain to send back more tokens than was initially sent there by the issuer. This feature is similar to the *Withdrawal Safeguard* feature introduced in Zendoo (see section 4.1.2.2 in [1]).

Let's consider this situation in more detail and assume that it is allowed to send foreign tokens from the sidechain where they have not been issued to another third-party sidechain. Let there be several fungible token instances  $TI_i$  of type  $TI_i.tokenName = SuperCoin$ , such that  $TI_i.issuerSidechainId = SC_a$ . Let assume they are held by  $SC_b$  at the moment (i.e.,  $TI_t \in S_b^{tkss}$ ) and  $SC_b$  sends some of them to  $SC_c$ . But given that this operation is outside  $SC_a$ , it can not track this transfer and update sent records accordingly (decrease balance of  $SC_b$  and increase balance of  $SC_c$ ). Possible solutions to this problem can be:

1. Eliminate recording the receiving sidechain in the sent record and just save the amount of tokens of a particular type that have been sent outside. But this creates a vulnerability: if the sidechain  $SC_c$  becomes corrupted, such that it can forge the sending transaction to pretend to transfer more tokens that it really has, then the original  $SC_a$  has no way to detect this and will accept a malicious transfer that may drain  $SR.amount$  to zero and prevent  $SC_b$  from transferring back the valid  $SC_a$  tokens it holds.
2. Eliminate the concept of sent records at all, but then a malicious sidechain would be able to forge any amount of tokens and send them back to the original sidechain without any control.
3. Notify the original sidechain  $SC_a$  every time when tokens are transferred somewhere. I.e., if  $SC_b$  sends  $TI_t$  to  $SC_c$ , then an additional transaction is created on  $SC_a$  that updates  $SR$  accordingly. Theoretically, it solves the problem, but it will require to design and implement complex set of proofs and it would be basically equivalent to sending tokens through  $SC_a$ . We do not consider this option due to its complexity.

There are some other solutions as well, but they are even more complex.

#### 4.5.1 Token wrapping

An interesting concept that could also serve as a workaround for the mentioned above problem is token wrapping. For instance, if a sidechain  $SC_b$  holds some tokens  $TI_i$  issued in  $SC_a$ , where  $TI_i.tokenName = ACoin$  and  $TI_i.issuerSidechainId = SC_a$ , it can issue a new set of tokens  $TI_j$ , such that  $TI_j.tokenName = wACoin$  and  $TI_j.issuerSidechainId = SC_b$ . The idea is that  $wACoin$  tokens can serve as a wrapper for  $ACoin$  tokens. Then,  $SC_b$  can send  $TI_j$  to other sidechains.

Wrapping may be particularly interesting for the native mainchain asset. Given that every sidechain may contain a certain amount of native MC assets, representing them in the form of wrapped tokens can simplify interaction among different sidechains.

## 5 Conclusions

We presented a generic protocol for cross-sidechain communication for Zendoo sidechains. It allows different chains to exchange messages with one another in a verifiable way. Moreover, we presented an extension of the protocol that allows to transfer tokens issued on different sidechains.

We view our work as a base for building interoperable Zendoo sidechains. The protocol does not impose any rules on the sidechains structure and can be used by different systems with different security models.

## References

- [1] Roman Oliynykov Alberto Garoffolo, Dmytro Kaidalov. Zendoo: a zk-snark verifiable cross-chain transfer protocol enabling decoupled and decentralized sidechains. Cryptology ePrint Archive: Report 2020/123. <https://eprint.iacr.org/2020/123>.
- [2] Adam Back, Matt Corallo, Luke Dashjr, Mark Friedenbach, Gregory Maxwell, Andrew Miller, Andrew Poelstra, Jorge Timón, and Pieter Wuille. Enabling blockchain innovations with pegged sidechains, 2014. <https://blockstream.com/sidechains.pdf>.
- [3] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. Cryptology ePrint Archive, Report 2013/879, 2013.
- [4] Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019. <https://eprint.iacr.org/2019/1021>.
- [5] Vitalik Buterin. Skeptical of cross-chain bridges. online article, 2022. <https://cryptobriefing.com/vitalik-buterin-skeptical-of-cross-chain-bridges/>.
- [6] Croman K. et al., editor. *On Scaling Decentralized Blockchains*, volume 9604 of *Financial Cryptography and Data Security, Lecture Notes in Computer Science*. Springer, 07 2016.
- [7] Alberto Garoffolo, Dmytro Kaidalov, and Roman Oliynykov. Zendoo: a zk-snark verifiable cross-chain transfer protocol enabling decoupled and decentralized sidechains. In *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)*, pages 1257–1262, 2020.
- [8] Ulrich Haböck, Alberto Garoffolo, and Daniele Di Benedetto. Darlin: Recursive proofs using marlin. Cryptology ePrint Archive, Paper 2021/930, 2021. <https://eprint.iacr.org/2021/930>.
- [9] Sergio Demian Lerner. Rootstock: Bitcoin powered smart contracts, 2019. <https://www.rsk.co/Whitepapers/RSK-White-Paper-Updated.pdf>.
- [10] Izaak Meckler and Evan Shapiro. Coda: Decentralized cryptocurrency at scale, 2018. <https://cdn.codaprotocol.com/v2/static/coda-whitepaper-05-10-2018-0.pdf>.
- [11] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. <https://bitcoin.org/bitcoin.pdf>.