

Zendoo: a zk-SNARK Verifiable Cross-Chain Transfer Protocol Enabling Decoupled and Decentralized Sidechains

Alberto Garoffolo
alberto@horizen.global
Horizen

Dmytro Kaidalov
dmytro.kaidalov@iohk.io
IOHK Research

Roman Oliynykov
roman.oliiynykov@iohk.io
IOHK Research

V.N.Karazin Kharkiv National University

January 2020

Abstract

Sidechains are an appealing innovation devised to enable blockchain scalability and extensibility. The basic idea is simple yet powerful: construct a parallel chain – sidechain – with desired features, and provide a way to transfer coins between the mainchain and the sidechain.

In this paper, we introduce *Zendoo*, a construction for Bitcoin-like blockchain systems that allows the creation and communication with sidechains of different types without knowing their internal structure. We consider a 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. We use zk-SNARKs to construct a universal verifiable transfer mechanism that is used by sidechains.

Moreover, we propose a specific sidechain construction, named *Latus*, that can be built on top of this infrastructure, and realizes a decentralized verifiable blockchain system for payments. We leverage the use of recursive composition of zk-SNARKs to generate succinct proofs of sidechain state progression that are used to generate certificates' validity proofs. This allows the mainchain to efficiently verify all operations performed in the sidechain without knowing any details about those operations.

Contents

1	Introduction	3
1.1	Related Work	4
2	Preliminaries	5
2.1	Cryptographic Definitions	5
2.2	Recursive SNARKs Composition for State Transitions	6
3	General Overview	7
3.1	Main Components of a Sidechain Design	8
4	Zendoo: a Cross-Chain Transfer Protocol for Sidechains	9
4.1	Cross-Chain Transfer Protocol	10
4.1.1	Forward Transfers	10
4.1.2	Backward Transfers	11
4.1.2.1	Mainchain Managed Withdrawals	14
4.1.2.2	Withdrawal Safeguard	16
4.1.3	Sidechain Transactions Commitment	16
4.2	Bootstrapping Sidechains	17
5	The Latus Sidechain	18
5.1	Consensus Protocol	18
5.1.1	Withdrawal Epochs	21
5.2	Accounting Model and System State	22
5.2.1	System State	22
5.3	Transactional Model	23
5.3.1	Payment Transaction	23
5.3.2	Forward Transfers Transaction	24
5.3.3	Backward Transfer Transaction	25
5.3.4	Backward Transfer Requests Transaction	26
5.4	State Transition Proof	27
5.4.1	Performance and Incentives	29
5.5	Cross-Chain Transfer Protocol	29
5.5.1	Mainchain Block Reference	29
5.5.2	Forward Transfers	31
5.5.3	Backward Transfers	31
5.5.3.1	Withdrawal Certificate	32
5.5.3.2	Backward Transfer Request	34
5.5.3.3	Ceased Sidechain Withdrawal	35
6	Conclusions	36
7	Acknowledgments	36
	Appendix A MST Delta	38

1 Introduction

Since the inception of the Bitcoin cryptocurrency in 2008 [20], the topic of decentralized ledger technology has received significant attention among experts from various areas. Bitcoin became the first decentralized payment system based on peer-to-peer networking. Its key feature – the absence of centralized control – is claimed to be the disruptive innovation that will help build more robust, fair, and transparent financial systems. Bitcoin inspired the appearance of many other systems based on the same principle of decentralization with a variety of different features.

With the increasing use of Bitcoin and similar blockchain systems, their inherent limitations became apparent: limited throughput, increased latency, reduced ability to scale and expand functionality, etc. [10]. Even more important is that such decentralized systems are challenging to update since there is no single decision-making entity. Even a small protocol change requires a cumbersome process of community agreement, which makes the introduction of new features difficult.

Sidechains, proposed by A. Back et.al. in 2014 [5], is an appealing concept that allows one to work around the constraints of a single decentralized blockchain. The basic idea is simple: to create a separate blockchain with whatever functionality is needed and provide a way to communicate with the main blockchain (Fig. 1). Communication means the ability to transfer a mainchain native asset (e.g. bitcoins) to and from a sidechain.

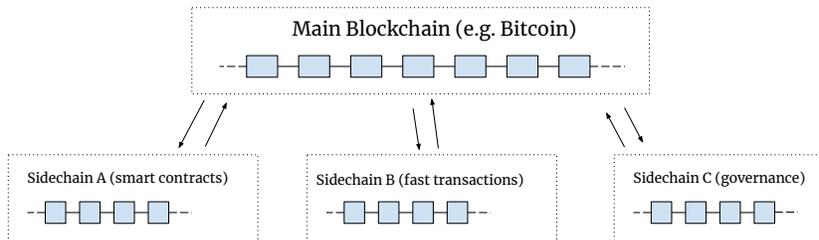


Figure 1: Sidechains. The main blockchain provides basic cryptocurrency functionality while sidechains implement specific functions.

This way, for instance, a blockchain system, like Bitcoin, can be extended with additional functionalities (such as smart contracts [3]) implemented in a separate sidechain, which uses the same native asset, hence remaining in the Bitcoin ecosystem.

In this paper, we propose *Zendoo*, a universal construction for Bitcoin-like blockchain systems that allows the creation and communication with sidechains of different types without knowing their internal structure (e.g. what consensus protocol is used, what types of transactions are supported, etc.). In fact, the sidechain may not even be a blockchain but can be any system that uses the standardized method to communicate with the mainchain.

Specifically, we consider a 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. Among other things, such certificates authorize transfers coming from sidechains. Certificate authentication and validation are achieved by using zk-SNARKs [6], which enable constant-sized proofs of arbitrary computations. The main feature of our construction is that sidechains are allowed to define their own zk-SNARKs, thus establishing their own rules for authentication and validation. The fact that all zk-SNARK proofs comply with the same verification interface used by the mainchain enables great universality as the sidechain can use an arbitrary protocol for authenticating its certificates. E.g., the sidechain may adopt a centralized solution where the zk-SNARK just verifies

that a certificate is signed by an authorized entity (like in [5]) or, for instance, a decentralized chain-of-trust model as in [13].

Moreover, we propose a specific sidechain construction, named *Latus*, that can be built on top of this infrastructure, and realizes a decentralized verifiable blockchain system. We leverage the use of recursive composition of zk-SNARKs to generate succinct proofs of sidechain state progression (as in [19]) that are used to generate certificate proofs for the mainchain. This allows the mainchain to efficiently verify all operations performed in the sidechain without knowing any details about those operations.

The paper is structured in the following way: section [2] provides basic definitions that are used throughout the paper; section [3] provides a general overview of the sidechain concept; section [4] provides details about the proposed cross-chain communication protocol Zendo and introduces basic interfaces imposed by the mainchain side; section [5] provides details of the *Latus* sidechain construction.

1.1 Related Work

The concept of sidechains was first introduced by A. Back et.al. in 2014 [5]. They introduced a general notion of a 2-way peg and described two operational modes – synchronous and asynchronous – to implement interactions between pegged chains. The synchronous mode implies that both main and side chains are aware of each other and can verify transfer transactions directly, while the asynchronous mode relies on validators to process transfers.

Notable construction of sidechains was presented in [25, 18] and called Drivechains. It aims to deploy sidechains on top of the Bitcoin network. While forward transfers (from the mainchain to a sidechain) are processed by providing SPV proofs (like the synchronous mode in [5]), backward transfers rely upon validators. Validators in Drivechains are mainchain miners who observe sidechains and endorse transfers.

The first formal treatment of sidechains was proposed by P. Gaži, A. Kiayias, and D. Zindros in [13]. In addition, they presented a sidechain construction for proof-of-stake blockchains where sidechain nodes directly observe and confirm forward transfers while backward transfers are confirmed by certifiers chosen among sidechain block forgers.

Our previous proposal on sidechains [12] presents a flexible model which allows the construction of different types of sidechains whose internal structures are unknown to the mainchain. It relies on certifiers to confirm backward transfers in the mainchain. Though, in this model certifiers are chosen randomly from a pool of certifiers registered directly in the mainchain.

Mentioned constructions differ from our current proposal in various aspects, most notably because they either assume that the mainchain observes sidechains directly or relies on some intermediary to confirm transfers. In addition, they do not provide flexibility (except [12]), which means that a sidechain construction (e.g. consensus protocol) cannot be chosen freely.

In [16], A. Kiayias and D. Zindros proposed implementation of the sidechain protocol for the proof-of-work blockchains based on smart contracts. Another notable sidechain construction that relies on smart contracts is called Plasma and was presented in [22] by J. Poon and V. Buterin. On the contrary, our construction does not rely on smart contracts.

One of the main features of the construction presented in this paper is the usage of zk-SNARKs for enabling verifiable cross-chain communication. zk-SNARK has initially been proposed as a zero-knowledge protocol which allows proving possession of some information without revealing it [6]. However, this technique is suited for more than simply securing information but also for solving scalability issues: it enables succinct constant size proofs of almost arbitrary computations. For instance, using the recursive composition of zk-SNARKs [19, 8] it is possible

to construct a succinct proof of state transition virtually for any number of transactions. We were inspired by these techniques while designing our sidechain construction.

A notable sidechain construction that also relies on zk-SNARKs is ZK Rollup [14]. It is a layer 2 solution based on smart contracts for scaling transaction throughput in Ethereum [2]. The basic idea is that transactions are carried out off-chain while the information about entailed state transitions together with a zk-SNARK proof of their validity is submitted to the contract. It still requires submission of some limited information about each transaction on-chain to prevent data availability attacks thus limiting scalability. Our construction differs from ZK Rollup in many aspects, most notably because we do not push sidechain transaction data to the mainchain.

There are many other attempts to construct cross-chain transfer mechanisms including the Liquid project [9], Polkadot [30], Interledger [27], Cosmos [1], and many others. They propose various solutions that are different from our construction.

2 Preliminaries

In this section, we introduce definitions of several cryptographic constructions that are used throughout the paper. We defer formal descriptions (especially of the recursive SNARKs composition) for a separate paper and define here only basic notations needed to describe the proposed sidechain construction.

2.1 Cryptographic Definitions

Definition 2.1. Collision-Resistant Hash Function (CRH). A hash function H is collision-resistant if the probability of finding two different input strings a and b such that $H(a) = H(b)$ is negligible (more formal definition can be found, e.g., in [15]).

Whenever we refer to a hash function, we suppose it is collision-resistant.

Definition 2.2. Merkle Hash Tree (MHT). The Merkle Hash Tree, or simply Merkle Tree (MT), is a binary tree data structure where the value of an internal node is computed as the hash of values of its children, and the value of a leaf node is the direct hash of a data block represented by this leaf (see Fig. 2) [23, 21].

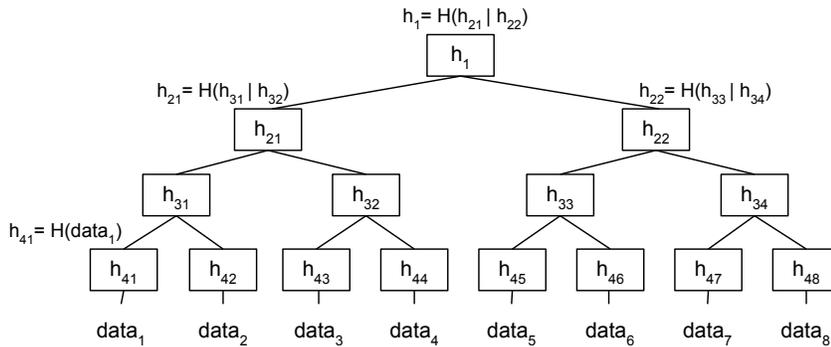


Figure 2: Merkle Hash Tree.

We call the top-level node (h_1 in Fig. 2) the root hash of the MHT. Given that a collision-resistant hash function is used to calculate tree nodes, we can consider root hash as a tree

authenticator: it is impossible to tamper even a single bit of data in the tree without also changing the root hash.

An important feature of the Merkle tree structure is that it produces a concise proof of a particular data block’s membership in a tree with the particular root hash. E.g., if one wants to prove that $data_4$ (Fig. 2) is included in the MHT tree with the root hash h_1 , they just need to provide a verifier with the data block along with a tuple of internal nodes (h_{43}, h_{31}, h_{22}) that will allow recalculating the tree root and comparing it to the provided root h_1 . We call it **Merkle proof**.

Definition 2.3. Succinct Non-Interactive Argument of Knowledge (SNARK).

A SNARK is a proving system consisting of a triplet of algorithms $(Setup, Prove, Verify)$ that allows proving satisfiability of a set of inputs to an arithmetic constraint system (see, e.g., [6, 7] for more formal definition and properties analysis).

We define an arithmetic constraint system as a set of polynomials over a finite field F in variables $(x_1, \dots, x_r, y_1, \dots, y_s)$. A satisfying assignment for the given constraint system C is an assignment of F elements to x_i and y_j such that all polynomials evaluate to zero. We indicate a satisfying assignment as $C(a, w)$, where $a = (a_1, \dots, a_r)$, $a_i \in F$ and $w = (w_1, \dots, w_s)$, $w_j \in F$. We refer to a as public input and w as witness.

Then, the algorithms $(Setup, Prove, Verify)$ are defined such that

1. $(pk, vk) \leftarrow Setup(C, 1^\lambda)$ bootstraps SNARK for a constraint system C under security parameter λ . The bootstrapped SNARK is specified by a pair of keys (pk, vk) which are a proving key and a verification key correspondingly.
2. $\pi \leftarrow Prove(pk, a, w)$ evaluates a proof π , which confirms that (a, w) is a satisfying assignment for C .
3. $true/false \leftarrow Verify(vk, a, \pi)$ verifies that π is a valid proof attesting to the satisfying assignment (a, w) for the constraint system C .

Algorithms $(Setup, Prove, Verify)$ satisfy the following properties:

1. **Completeness.** For any constraint system C and (a, w) , if $\pi \leftarrow Prove(pk, a, w)$ is a valid proof, then $Verify(vk, a, \pi)$ is always true.
2. **Knowledge soundness.** If a pair (a, w) is not a satisfying assignment for C , then the probability of obtaining π such that $Verify(vk, a, \pi) = true$ is negligible.
3. **Succinctness.** For every constraint system C bootstrapped with (pk, vk) and every $a \in F^r$, the size of a proof and verification time is polynomial in λ .

2.2 Recursive SNARKs Composition for State Transitions

Here, we provide a high-level definition of the recursive proof composition technique that is used in our sidechain model to construct succinct proofs of state transitions. The idea of recursive proofs has been discussed, e.g., in [6, 19, 8]. What follows is based principally on the construction described in [19].

Definition 2.4. State Transition System. A state transition system is defined by a set of all possible states S , a set of all possible transitions T , and a transition function $update(t_i, s_i)$, where $s_i \in S$ and $t_i \in T$, which returns a new state s_{i+1} or \perp in case (t_i, s_i) does not constitute a valid input for the $update$ function.

Speaking informally, we would like to define a SNARK that attests to many iterative state transitions. E.g., if we have transitions (t_1, t_2, \dots, t_n) that are applied sequentially to state s_1 to

produce state s_{n+1} , we would like to have a succinct proof of the following statement: “there exist such (t_1, \dots, t_n) so that $update(t_n, update(t_{n-1}, update(\dots, update(t_1, s_1)))) = s_{n+1}$ ”.

By applying this to blockchain, we will be able to provide succinct proofs of transition between some states s_i and s_j ($i < j$). The state can be represented, for instance, as a list of unspent transaction outputs [29], while transitions are regular blockchain transactions that spend some outputs and create new ones. This construction is of great value for verifiable sidechains.

Definition 2.5. Recursive SNARKs for state transition systems. We define recursive SNARKs composition as a tuple of SNARKs $(Base, Merge)$ such that:

1. **Base** is a SNARK for a single transition that proves the existence of such t so that $s_{i+1} = update(t, s_i)$. It is defined by a triplet $(Setup, Prove, Verify)$ such that:
 - $(pk^{Base}, vk^{Base}) \leftarrow Setup(1^\lambda)$ bootstraps *Base* SNARK;
 - $\pi^{Base} \leftarrow Prove(pk^{Base}, (s_i, s_{i+1}), (t_i))$ evaluates a proof π^{Base} that confirms $s_{i+1} = update(t_i, s_i)$;
 - $true/false \leftarrow Verify(vk^{Base}, (s_i, s_{i+1}), \pi^{Base})$ verifies that π^{Base} is a valid proof attesting state transition from s_i to s_{i+1} .
2. **Merge** is a SNARK that merges two other SNARKs (either *Base* or *Merge*) proving the validity of transition between states s_i and s_j ($i + 1 < j$). It is defined by a triplet $(Setup, Prove, Verify)$ such that:
 - $(pk^{Merge}, vk^{Merge}) \leftarrow Setup(1^\lambda)$ bootstraps *Merge* SNARK;
 - $\pi^{Merge} \leftarrow Prove(pk^{Merge}, (s_i, s_j), (s_k, \pi_1^a, \pi_2^a))$ evaluates a proof π^{Merge} that confirms π_1^a, π_2^a are valid SNARKs ($a \in \{Base, Merge\}$), which attest state transitions from s_i to s_k and from s_k to s_j correspondingly. Altogether, it proves a valid transition from s_i to s_j ($i < k < j$);
 - $true/false \leftarrow Verify(vk^{Merge}, (s_i, s_j), \pi^{Merge})$ verifies that π^{Merge} is a valid proof attesting state transition from s_i to $s_j, i < j$.

We intentionally omit specifics of the recursive SNARKs composition, which in reality is more sophisticated. More details on the topic can be found, for instance, in [19, 8]. We defer the details of our construction for a separate paper. At this point, we provided only basic definitions, which allow us to describe the sidechain protocol while abstracting away the details of the SNARKs construction.

3 General Overview

This section gives an overview of the sidechain concept in general and the main components of any sidechain design. We also briefly discuss our proposed solutions.

Before going any further, we want to introduce abstract definitions of the terms *mainchain* and *sidechain* that are used throughout this paper.

Definition 3.1. Mainchain (MC). The mainchain is a blockchain system based on the Bitcoin backbone protocol model [11], which maintains a public ledger of asset-transfer transactions. Additionally, the mainchain supports a standardized mechanism to register and interact with separate sidechain systems. By interaction, we mean the cross-chain transfer protocol, which enables sending a native asset to a sidechain and receiving it back in a secure and verifiable way without the need to know anything about the internal sidechain construction or operations.

Definition 3.2. Sidechain (SC). The sidechain is a separate system attached to the mainchain by means of a cross-chain transfer protocol.

Speaking informally, we consider the mainchain to be a blockchain platform that supports basic payment functionality with some native asset *Coin* (e.g. Bitcoin [20], Horizen [26], etc.). Then, the sidechain is an attached domain-specific platform that also uses the *Coin* asset (but not limited to it). In our model, we consider a single mainchain with many sidechains attached to it.

The definition of a sidechain (as in Def. 3.2) does not imply the usage of any particular data structure or consensus algorithm. The mainchain is totally agnostic to the sidechain construction. It can be another decentralized blockchain, some centralized database maintained by the predefined authority, or more generally, an arbitrary application.

The need to introduce sidechains in a general payment-based blockchain system comes from the need to allow the creation of different blockchain applications that use the same mainchain asset. Creation of such applications directly on the mainchain is not always possible due to inherent technological limitations, such as restricted throughput, expensive storage, etc. Sidechains effectively solve these problems.

3.1 Main Components of a Sidechain Design

Analyzing existing attempts to design sidechains [12, 5, 25, 18, 16, 13, 14], we may outline three basic components that underlie any sidechain architecture:

1. **Mainchain consensus protocol (MCP).**
2. **Cross-chain transfer protocol (CCTP).**
3. **Sidechain consensus protocol (SCP).**

Depending on a specific design, these components can be highly coupled with each other or decoupled so that the mainchain is almost independent from any particular sidechain implementation.

In our construction, we aspire to multipurposeness and, thus, designing a system so that the MCP and SCP are completely decoupled. The CCTP is naturally a bridge between them and is unified and fixed by the mainchain consensus protocol. On the other end, the SCP can be freely defined by a sidechain developer. This allows a variety of different sidechains with different purposes to thrive while not requiring any changes to the mainchain.

Cross-chain transfer protocol. The CCTP protocol defines the communication between the mainchain and sidechain(s). Basically, it is a 2-way peg protocol that allows sending coins back and forth. At a high level, it defines two basic operations:

- **Forward Transfer**, and
- **Backward Transfer**.

A forward transfer sends coins from the mainchain to a sidechain. A backward transfer, correspondingly, moves coins back from the sidechain to the mainchain. These operations are the cornerstone of overall sidechain construction. The backward transfer is of particular importance since we do not want to oblige the mainchain to track sidechains and, thus, it cannot directly verify the validity of withdrawals coming from them. That is why most of the focus in developing sidechains is directed toward constructing secure and reliable backward transfers.

In our approach, Zendo, we consider a forward transfer as a special transaction on the mainchain that destroys coins and provides sidechain-specific metadata allowing a user to receive

coins in the sidechain. Implementation of forward transfers is straightforward as it does not require the mainchain to know anything about the sidechain state for validation.

A more complex procedure is required for backward transfers. They are initiated in the sidechain as special transactions, batched in a *withdrawal certificate*, and propagated to the mainchain for processing. Since we do not want the MC to follow the SC state – as this would impose enormous computational and storage burden on the MC and, thus, undermine the whole point of having sidechains – the question arises how to implement validation of backward transfers in the most efficient and secure way.

In our previous paper [12], this problem has been addressed by introducing into the system a special type of decentralized actors – *certifiers* – that were registering themselves in the MC and were responsible for signing withdrawal certificates. Although the safety of this approach has been shown, it requires certain assumptions about an honest majority of certifiers, which, in some scenarios, may not be the case.

In Zendo, we avoid direct reliance on certifiers or any other special type of actors assigned to validate withdrawal certificates. Instead, we are going to leverage SNARKs [6, 7, 19] to provide means for the mainchain to effectively validate withdrawals.

Sidechain consensus protocol. We consider the SCP as a generalized notion that encompasses all the details about a particular sidechain construction such as consensus algorithm, accounting system, types of supported transactions, incentives mechanism, a protocol for withdrawal certificate generation, etc. Also, importantly, each sidechain defines its own SNARK¹ that is used to validate withdrawal certificates. This provides flexibility to define its own rules for backward transfers. For instance, a sidechain can adopt a chain-of-trust model [13] or even the certifiers model [12]. It is completely decoupled from the mainchain consensus protocol, which will just invoke a unified verifier to validate a proof.

Even though the SCP can be designed in different ways, we propose one specific construction of a decentralized verifiable sidechain based on the Ouroboros protocol [17]. We will call this construction Latus. In short, we are going to use recursive composition of SNARKs to generate succinct proofs of sidechain state transitions. Each withdrawal certificate commits to the SC state whereas the SNARK proof validates transition between states committed by successive withdrawal certificates. Since backward transfers are a part of the sidechain state transition, they are also validated by the proof.

The following section [4] introduces Zendo, a cross-chain communication protocol for sidechains which is principally about the definition of the transfer protocol and how a new sidechain can be registered in the mainchain. It defines the sidechains design from the mainchain point of view. Then, in section [5], we describe in detail the proposed sidechain construction Latus.

4 Zendo: a Cross-Chain Transfer Protocol for Sidechains

The following section provides details about the communication protocol between the mainchain and sidechains, which is primarily represented by the cross-chain transfer protocol. We show how the CCTP protocol is integrated in the mainchain, what interfaces are provided, and how a new sidechain can be created.

¹Speaking more formally, here, we refer to an arithmetic constraint system (arithmetic circuit) that is compiled for each sidechain and defines the logic of the SNARK. Note that the interface of the verifier is defined by the mainchain so – even though the internal logic of the SNARK may be different for different sidechains – generated proofs can be verified in the standardized way.

4.1 Cross-Chain Transfer Protocol

The cross-chain transfer protocol is the cornerstone of our sidechain design as it connects the mainchain with all sidechains spawned from it. Its main function is to allow sending coins to sidechains and receiving them back in a secure and reliable way. This section provides a high-level specification of forward and backward transfers and how they are integrated into the mainchain.

4.1.1 Forward Transfers

The design of forward transfers is straightforward and similar to many existing proposals for sidechains [5, 25, 18, 16, 13] as well as to our original proposal [12].

On the mainchain side, it is implemented as a special type of operation (we will call it *Forward Transfer*) that destroys coins and provides metadata for withdrawing coins in a sidechain. Then, it is the responsibility of the sidechain to sync forward transfers from the MC and issue the corresponding amount of coins.

Definition 4.1. Forward Transfer (FT). Forward Transfer is an operation that moves coins from the original blockchain A (the mainchain) to the destination sidechain B . It is represented by a tuple of the form:

$$FT \stackrel{\text{def}}{=} (\text{ledgerId}, \text{receiverMetadata}, \text{amount}),$$

where:

- ledgerId – a unique identifier of a previously created and active sidechain to which coins are transferred;
- amount – a number of coins to transfer;
- receiverMetadata – some metadata for receiving sidechain B (e.g., a receiver’s address); its structure is not fixed in the mainchain and can consist of different variables of predefined types depending on a sidechain’s construction; its semantic meaning is not known to the mainchain.

There can be several approaches to integrate forward transfers on the mainchain side depending on its details. For instance, forward transfer can be a separate transaction type which destroys coins in the mainchain, or, in the case of a UTXO-based blockchain system (e.g. Bitcoin or Horizen), we can consider FT as a special unspendable transaction output in a regular multi-input multi-output transaction [29, 4].

To be more specific and facilitate further reading, we assume that the mainchain has a UTXO-based accounting model². Then, a regular transaction with forward transfers may have the following structure:

²It is in line with our own implementation of sidechains, which is going to be deployed on the Horizen mainchain.

```

type Transaction {
  Inputs: {
    Input(addr: 0x013A..., amount: 5, signature: 0x034B..),
    Input(addr: 0x0930..., amount: 3, signature: 0x1AA1..),
    ...
  }
  Outputs: {
    Output(addr: 0x023B..., amount: 1),
    Output(addr: 0x0732..., amount: 2),
    ForwardTransfer(ledgerId: 0x300C..., receiverInfo : 0x139D..., amount: 2),
    ForwardTransfer(ledgerId: 0x300C..., receiverInfo : 0x893D..., amount: 3),
    ...
  }
}

```

Given that an FT is a non-spendable output, it basically destroys coins in the mainchain, and the amount of transferred coins is verified by the mainchain as part of the overall transaction verification.

4.1.2 Backward Transfers

The backward transfer protocol allows coins to move from a sidechain to the mainchain, and, as in [12] and [13], it relies on the idea of batched transfers. This means that all requested backward transfers submitted to the sidechain during a certain period – called the “withdrawal epoch” – are collected in a special withdrawal certificate and pushed to the mainchain for processing.

Withdrawal certificates are more than just a container for backward transfers, they are a kind of sidechain heartbeat that is periodically submitted to the mainchain even though there might be no backward transfers³.

A withdrawal epoch is defined by a range of MC blocks. Withdrawal epochs for different sidechains are not aligned and may have a different length (*epoch_len* parameter is set upon sidechain creation), and, therefore, the entire system runs asynchronously.

Let us define an MC block B that belongs to a specific epoch as B_j^{ep-id} , where ep_id is the epoch number and $j \in [0, epoch_len)$ is the serial number of the block within the epoch (Fig. 3).

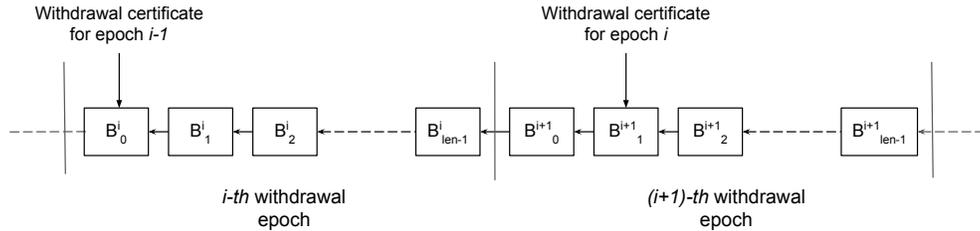


Figure 3: Withdrawal epochs in the mainchain. Note that withdrawal epochs for different sidechains may not overlap. It depends on parameters that have been set upon sidechain creation.

The sidechain is obliged to submit a withdrawal certificate for epoch i during the first *submit_len* blocks of the epoch $i + 1$ (*submit_len* is a system parameter). If a withdrawal

³Later, in section [5], we will show how this can be leveraged to construct secure verifiable sidechains by providing their state commitments as part of withdrawal certificates

certificate has not been submitted during this time, the sidechain is considered *ceased* (see Def. 4.2) and no more withdrawal certificates for this sidechain will be accepted by the mainchain (however, the funds can still be withdrawn with a ceased sidechain withdrawal [4.1.2.1]).

Definition 4.2. Ceased Sidechain. A sidechain is deemed ceased by the mainchain if a withdrawal certificate for that sidechain has not been submitted on time, i.e. a certificate for withdrawal epoch i has not been submitted during the first $submit_len$ blocks of the epoch $i+1$ ⁴.

Note that the mainchain consensus protocol does not impose any rules on how exactly a withdrawal certificate should be generated and by whom it should be submitted. It is up to the sidechain to define corresponding procedures. We only assume that it is submitted by means of a special transaction in the mainchain.

As it has been mentioned, a withdrawal certificate contains backward transfers. We may consider them as requests that are fulfilled once included in the withdrawal certificate and propagated to the mainchain. There are no restrictions for how backward transfers should be submitted and collected (e.g. it can be a separate transaction on the SC side).

Definition 4.3. Backward Transfer (BT). Backward Transfer is an operation that moves coins from the sidechain B to the original mainchain A . It is represented by a tuple of the form:

$$BT \stackrel{\text{def}}{=} (receiverAddr, amount),$$

where:

- receiverAddr* – an address in the mainchain where transferred coins should be credited;
- amount* – the number of transferred coins.

There can be different approaches to integrate backward transfers in the mainchain. Following the assumption of a UTXO-based mainchain, a BT can be represented by a special output in a transaction with a withdrawal certificate.

Definition 4.4. Withdrawal Certificate (WCert). Withdrawal certificate is a standardized posting that allows sidechains to communicate with the mainchain. Its main functions are:

- 1) delivering backward transfers to the MC; and
- 2) serving as a heartbeat message enabling the MC to identify SC status.

It 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 (explained later);
- BTList* – a list of backward transfers included in this withdrawal certificate;
- proofdata* – input data to a SNARK verifier;
- proof* – a SNARK proof.

⁴Even though this requirement may seem strong, it is necessary to provide certain properties which we discuss later. We also explore the possibility to provide more flexibility for withdrawal certificate submission

Now, we discuss in more detail the substance of certificate parameters as it is one of the most important parts of the sidechain design.

As it has been briefly outlined, the basis of the proposed construction is that there are no special entities that authorize withdrawal certificates (e.g., like certifiers in [12] or slot leaders in [13]). Instead, the certificate authorization and validation rely completely on the included SNARK proof, and the SNARK itself is defined by the sidechain. The mainchain knows only the verification key – which is registered upon sidechain creation – and the interface of the verifier, which is unified for all sidechains. If the SNARK proof and public parameters are valid, then the certificate gets included and processed in the mainchain.

Withdrawal certificate quality. It might happen that several withdrawal certificates appear for the same sidechain in the same withdrawal epoch⁵. Since only one WCert should be selected among them and given that the mainchain does not know about the sidechain consensus protocol and does not track its state, there should be a mechanism for the mainchain to decide which certificate is the best one. Such a mechanism is realized through the *quality* parameter: the mainchain adopts a certificate with the highest quality or the one that was submitted first in case there are several certificates with equal qualities. The validity of the quality parameter is enforced by the SNARK proof.

Withdrawal certificate verification. Verification of a newly submitted WCert on the mainchain is performed using the following basic rules:

WCert Verification

1. *ledgerId* should be an identifier of a currently active sidechain;
2. *epochId* should be a valid withdrawal epoch number for the *ledgerId* (remember that the certificate should be submitted during the first *submitLen* blocks of the epoch following the one, for which such certificate was created);
3. *quality* should be higher than the quality of the previously submitted withdrawal certificate for this epoch; if it is the first WCert for this epoch - any quality is accepted;
4. *proof* should be a valid SNARK proof whose verification key vk_{WCert} is set upon sidechain registration;

SNARK verification is the most essential part of the verification procedure as it encapsulates verification of backward transfers and other parameters provided within the certificate. The basic SNARK verifier interface is the following:

$$true/false \leftarrow Verify(vk_{WCert}, public_input, proof),$$

$$public_input \stackrel{\text{def}}{=} (wcert_sysdata, MH(proofdata)),$$

where:

⁵Even though it should not happen under normal operation, it may be the case, for instance, if the sidechain is a blockchain system, which experiences a continuous fork, or due to some malicious activity

- vk_{WCert} – 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 that is defined by the sidechain and passed along the withdrawal certificate; it is basically 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 only⁶;
- $proof$ – a SNARK proof itself submitted as a part of the certificate.

$wcert_sysdata$ parameter plays an important role from the security standpoint. The idea is to allow the mainchain to verify the proof against some public input parameters that are defined by the protocol. For instance, the $BTList$ and $quality$ parameters that are part of the certificate must be verified before being used by the mainchain. Another example is the mainchain block hashes of the epoch boundaries that must be verified to guarantee that the proof refers to the current epoch and the active chain.

$wcert_sysdata$ is represented by the tuple of the following form:

$$wcert_sysdata \stackrel{\text{def}}{=} (quality, MH(BTList), H(B_{last}^{i-1}), B_{last}^i),$$

where:

- $quality$ – the quality parameter from the withdrawal certificate;
- $MH(BTList)$ – a root hash of a Merkle tree where leaves are backward transfers from the BTList provided within the certificate;
- $H(B_{last}^{i-1})$ – a block hash of the last mainchain block in the withdrawal epoch $i - 1$ (given that the certificate is for the epoch i);
- $H(B_{last}^i)$ – a block hash of the last mainchain block in the withdrawal epoch i .

The generalized SNARK verifier provides flexibility to implement different SNARKs for different sidechain models. For instance, one may want to implement the sidechain with a centralized cross-chain transfer protocol where withdrawal certificates are verified by a signature from an authorized entity. Or, conversely, a completely decentralized verifiable sidechain can be constructed as will be discussed in [5 [The Latus Sidechain](#)].

Succinct proofs and constant time verification make the overall sidechain design particularly appealing as it does not impose a significant burden for the mainchain.

4.1.2.1 Mainchain Managed Withdrawals

There might be cases when a user would want to request a backward transfer directly from the mainchain rather than creating a BT in the SC. For instance, it would allow users to withdraw funds in case of a misbehaving (e.g., maliciously controlled sidechain that censors submission of backward transfers) or ceased sidechain.

Hence, we introduced two additional mechanisms that allow users to make withdrawals directly in the mainchain:

1. **Backward transfer request (BTR)**, and
2. **Ceased sidechain withdrawal (CSW)**.

⁶A full payload of $proofdata$ is provided during the proof generation as a witness.

We consider each of them as a special type of transaction. Similar to withdrawal certificates, such operations are secured by SNARK proofs.

The **BTR** is used to withdraw funds from an active sidechain if for some reason a user cannot create a backward transfer inside the sidechain. The idea is that all BTRs submitted to the mainchain will be synchronized to the sidechain and processed there to verify their legitimacy and include the corresponding backward transfers in the next WCert using the standard flow. Such processing can be enforced by the withdrawal certificate SNARK to force a maliciously controlled sidechain to process user’s withdrawals⁷. Importantly, the BTR does not lead to a direct coin transfer in the mainchain.

The **CSW** is used to withdraw funds from ceased sidechains. Since withdrawal certificates are not allowed for ceased sidechains, it becomes the only way to retrieve funds. A valid CSW makes direct payment to the submitter.

Definition 4.5. Backward transfer request (BTR). The BTR is a generic request for a backward transfer that is submitted on the mainchain. It is represented by the following tuple:

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

where:

- ledgerId* – an identifier of the sidechain, for which BTR is created;
- receiver* – an address of the receiver on 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.

As in the case with a withdrawal certificate, the SNARK for the BTR is defined by the sidechain and represented by the verification key vk_{BTR} , which 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:

$$\begin{aligned} true/false &\leftarrow \text{Verify}(vk_{BTR}, public_input, proof), \\ public_input &\stackrel{\text{def}}{=} (btr_sysdata, MH(proofdata)), \end{aligned}$$

where:

- vk_{BTR} is a SNARK verification key for the BTR registered upon the sidechain creation;
- btr_sysdata*, *proofdata*, $MH(\cdot)$, and *proof* have the same meaning as similar parameters in the withdrawal certificate.

btr_sysdata is defined as:

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

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

Definition 4.6. 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 all parameters have the same meaning as in the case of the BTR.

⁷Note that it is up to a sidechain construction to define exactly how BTRs are processed (for example, the Latus sidechain construction [5 The Latus Sidechain] introduces its own method for enforcing BTRs processing).

As it can be seen, BTR and CSW have the same structure, though conceptually they are different because CSW performs direct payment while BTR does not. The interface of the SNARK verifier for the CSW is completely the same as for the BTR.

Additionally, we discuss the role of nullifiers in both BTR and CSW. In the mainchain, a nullifier is an abstract identifier of claimed coins. The mainchain will not allow the submission of two transactions with the same nullifier. The main reason for having the nullifier is to prevent repeated submission of BTRs or CSWs that try to withdraw the same coins (thus, essentially doing double spend). Since the mainchain does not maintain the sidechain state, at the very least, for ceased sidechains, it requires some abstract double-spend prevention mechanism, which is exactly what is provided by nullifiers.

Note that both BTR and CSW are just complementary operations to allow more flexibility in some subtle use cases or in the case of a malfunctioning sidechain. It is up to the sidechain to define how they are used. For instance, one can omit defining these operations at all (e.g., by setting vk_{BTR} and vk_{CSW} to NULL), thus completely relying on the normal withdrawal procedure through withdrawal certificates.

4.1.2.2 Withdrawal Safeguard

The **safeguard** is a special feature introduced to prevent unlimited withdrawals from a sidechain to the mainchain in the case of the malicious sidechain. The essence of the safeguard function is to maintain the balance of each created sidechain and to prevent withdrawing an amount larger than what was previously transferred to that sidechain. A similar idea was introduced in [13] and [12].

Implementation of the safeguard feature is simple: for each created sidechain, a special balance variable is maintained by the mainchain. Each forward transfer increases the balance by the transferred number of coins, and each withdrawal certificate, or ceased sidechain withdrawal reduces the balance by the withdrawn amount. The WCert and CSW cannot withdraw more coins than are stored in the sidechain balance.

This feature prevents possible implications of sidechain corruption. It guarantees that only the transferred number of coins can be withdrawn back to the mainchain. Even in the case of total corruption or a maliciously constructed sidechain, an adversary cannot mint coins out of thin air.

4.1.3 Sidechain Transactions Commitment

So far, we defined 4 types of actions (that are either separate transactions or outputs in a regular transaction) that determine cross-chain communication from the mainchain point of view:

1. Forward Transfer (FT).
2. Withdrawal Certificate (WCert).
3. Backward Transfer Request (BTR).
4. Ceased sidechain withdrawal (CSW).

To facilitate efficient implementation of the synchronization between the mainchain and sidechains, we modify the structure of a mainchain block header to include an additional value that commits to all sidechain-related actions in the MC block (except the CSW because it is used only when the SC is ceased). This value is a root hash of a Merkle tree that contains

all transactions or outputs related to any sidechain (see Fig. 4). We call it the *Sidechain Transactions Commitment (SCTxsCommitment)*.

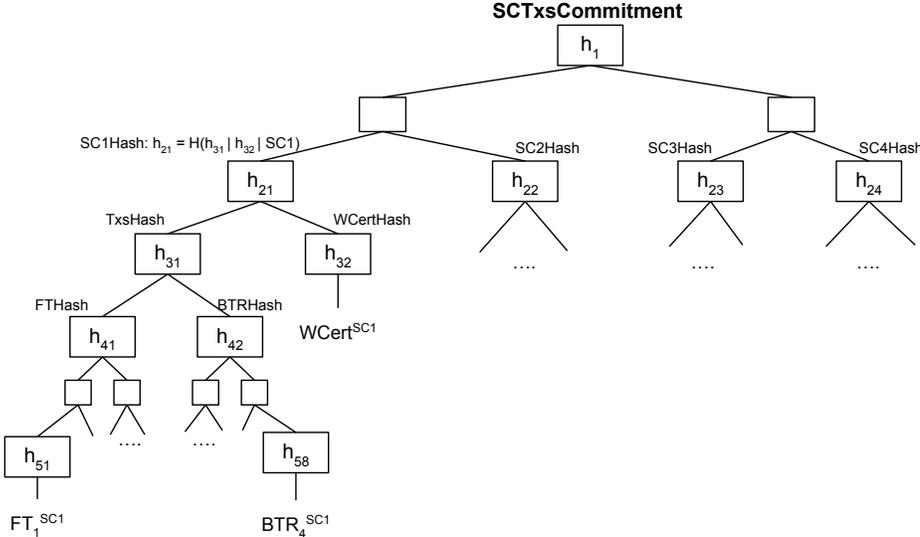


Figure 4: Sidechain transactions commitment tree. The root hash h_1 commits to all sidechain related transactions (for all sidechains) included in the MC block. All $SCXHash$, where X is a sidechain identifier, are ordered by the id and commit to all transactions related to the sidechain X . $WCertHash$ commits to the $WCert$ for the sidechain X (if present); only one $WCert$ is allowed for each sidechain. $TxsHash$ commits to FTs and BTRs.

Having *SCTxsCommitment* in the MC block header allows SC nodes to synchronize and verify SC-related transactions without the need to transmit the entire MC block. Also, it allows the construction of a SNARK proving that all SC-related transactions of the specific MC block have been processed correctly.

4.2 Bootstrapping Sidechains

We assume that the mainchain implements a special transaction that allows one to create a sidechain. Such a transaction can be submitted by anyone, and it registers the SC in the mainchain and sets its unique identifier and some system parameters. Once the sidechain is created, a schedule of withdrawal epochs is defined deterministically, and forward/backward transfers must be processed in the mainchain.

The following set of SC parameters are set upon creation:

Sidechain configuration

<i>ledgerId</i>	– a unique identifier of the sidechain that has not been used;
<i>start_block</i>	– the block number in the mainchain, from which the first withdrawal epoch begins; this parameter defines when the sidechain becomes active;
<i>epoch_len</i>	– the length of a withdrawal epoch (in MC blocks);
<i>submit_len</i>	– the period length – starting from the first block of the withdrawal epoch – when a withdrawal certificate for the previous epoch must be submitted to the mainchain;
<i>wcert_vk</i>	– a SNARK verification key vk_{WCert} for WCert proofs;
<i>btr_vk</i>	– a SNARK verification key vk_{BTR} for BTR proofs;
<i>csw_vk</i>	– a SNARK verification key vk_{CSW} for CSW proofs;
<i>wcert_proofdata</i>	– the definition of the <i>proofdata</i> structure for the withdrawal certificate; it defines the number and types of included data elements;
<i>btr_proofdata</i>	– the definition of the <i>proofdata</i> structure for the BTR;
<i>csw_proofdata</i>	– the definition of the <i>proofdata</i> structure for the CSW.

Customizable parameters give flexibility in choosing those, which are suitable for a particular sidechain. The triplet (cer_vk, btr_vk, csw_vk) is especially important as it defines how the mainchain verifies backward communication from the sidechain. These keys define SNARKs for corresponding operations eventually enabling different designs for sidechains.

5 The Latus Sidechain

In the previous section, we described the general sidechain design. Mostly, it was about defining the cross-chain transfer protocol, which provides a communication interface with the mainchain. In this section, we focus on a specific sidechain construction. We give an example of how a decentralized verifiable sidechain can be built on top of the given CCTP.

The general idea is to utilize a recursive composition of SNARKs to construct a succinct proof of the sidechain state progression for the period of a withdrawal epoch. Then, a SNARK for a withdrawal certificate is constructed so that it proves correct sidechain state transition for the whole epoch and validates backward transfers. This allows the mainchain to efficiently verify the sidechain without having to rely on any intermediary – such as certifiers [12] – and still be oblivious to the sidechain construction and interactions within.

In this section, we provide details of the proposed Latus sidechain that implements decentralized permissionless blockchain \mathcal{B}_{SC} with a proof-of-stake based consensus protocol. We consider \mathcal{B}_{SC} as a simple ledger of payment transactions. We assume that \mathcal{B}_{SC} does not possess its own native asset and, instead, uses only *Coin* asset transferred from the mainchain by means of the CCTP. Additionally, we assume that the mainchain is a classical proof-of-work based blockchain system with Nakamoto consensus [20] (e.g., Horizen [26]).

5.1 Consensus Protocol

We use a similar consensus protocol as in our previous proposal [12] with some minor adjustments. It is based on a modified version of the Ouroboros proof-of-stake consensus protocol [17].

In Ouroboros, time is divided into epochs with a predefined number of slots. Each slot is assigned with a slot leader who is authorized to generate a block during this slot. Slot leaders of a particular epoch are chosen randomly before the epoch begins from the set of all sidechain stakeholders (Fig. 5). The protocol operates in a synchronous environment where each slot takes a specific amount of time (e.g., 20 seconds).

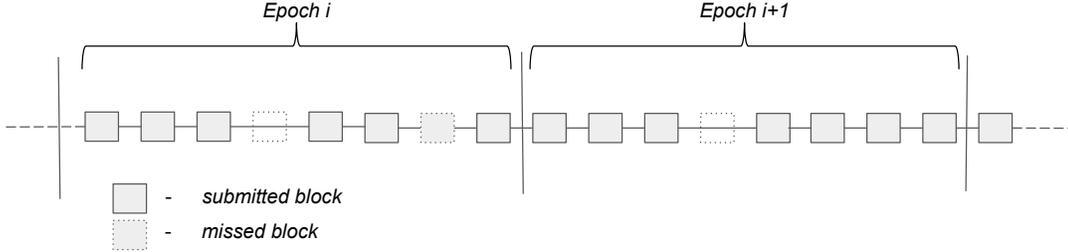


Figure 5: A general scheme of an epoch. Note that even though there is an assigned slot leader for each slot, the leader may skip block generation, and in this case, the slot remains empty.

Epoch. An epoch is a sequence of the k successive slots $Ep_i = (sl_i^0, sl_i^1, \dots, sl_i^{k-1})$, where k is the predefined length of the epoch and i is the epoch sequence number.

Slot. A slot is a specific period in time during which a slot leader is authorized to issue a block. Each slot has the corresponding slot leader who is chosen randomly before the epoch begins. A slot leader may skip generating a block, in this case, the following block will refer to the latest generated block.

Slot Leader. The slot leader of the slot sl_i^j is a stakeholder who was authorized by the Slot Leader Selection Procedure to forge a block at slot sl_i^j .

Slot Leader Selection Procedure. The slot leader selection procedure $Select(SD_{Ep_i}, rand)$ is a procedure that selects all slot leaders of the epoch Ep_i according to the fixed stake distribution SD_{Ep_i} and some random value $rand$. The stake distribution SD_{Ep_i} is fixed before the epoch Ep_i begins. The randomness $rand$ is revealed only after the stake distribution is fixed.

In our construction, we additionally introduce binding with the mainchain. This implies that sidechain blocks contain references to mainchain blocks so that their history is preserved in the sidechain. The chain resolution algorithm is altered to enforce that the sidechain follows the longest mainchain branch.

As a “mainchain block reference”, we consider a whole mainchain block header together with transactions related to the referencing sidechain.

Sidechain block forgers are obliged to keep mainchain references consistent and ordered when included in SC blocks. A sidechain block SB_j can contain a reference to the mainchain block B_i if and only if

1. the block B_i is a valid mainchain block, and
2. references to all previous mainchain blocks B_k , $k \in \{\eta, \eta + 1, \dots, i - 1\}$ have been already included in sidechain blocks (also considering the current one, as a sidechain block may contain more than one reference), where η is the genesis reference (Fig. 6).

Even though it is not mandatory for the block forgers to include mainchain references, we assume that honest block forgers will do this to support the cross-chain transfer protocol between chains. It is also possible to construct an incentive mechanism for block forgers who include references. For instance, users who initiate forward/backward transfers may pay some fee from

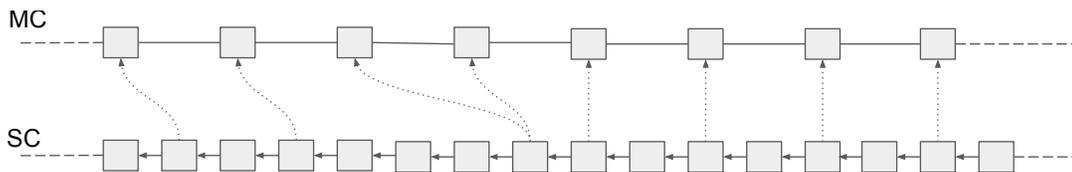


Figure 6: An example of the sidechain binding to the mainchain.

each transaction. The incentive mechanism is beyond the scope of the current paper as we only provide an example of a sidechain consensus protocol.

The binding to the mainchain provides two important properties of our sidechain construction:

1. **Deterministic synchronization between the MC and the SC.** When the sidechain block SB_i refers to the mainchain block B_j , it explicitly acknowledges all transactions included in the block B_j . It means that if B_j contains any transactions related to this sidechain (by transactions, we mean forward transfers and backward transfer requests), such transactions are immediately included in the sidechain (see Fig. 7).

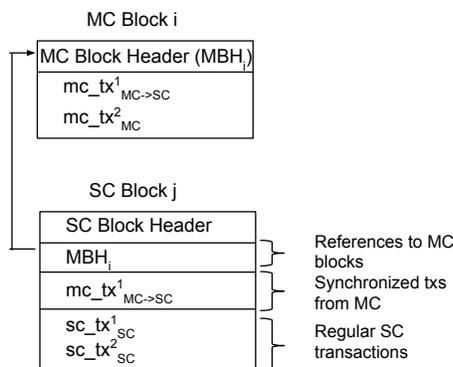


Figure 7: An example of transaction synchronization between the mainchain and the sidechain: MC block B_i contains one SC-related transaction $mc.tx^1_{MC \rightarrow SC}$, which is also included in the SC block SB_j because it refers to B_i .

2. **Mainchain forks resolution.** It is known that Nakamoto consensus does not provide finality on a chain of blocks [24]. It means that there is always a non-zero probability that some sub-chain of MC blocks will be reverted and substituted by another sub-chain with the more cumulative work. Such behaviour is normally handled by the mainchain but may be disastrous for the sidechain because $MC \rightarrow SC$ transactions that are already confirmed in the sidechain may be reverted in the mainchain. The binding eliminates such situations because in the case of a fork in the MC, SC blocks that refer to forked blocks in the MC would also be reverted.

Security. The standard procedure for proving blockchain consensus protocol security requires demonstrating the ability of the protocol to satisfy two fundamental properties of a distributed ledger: **liveness** and **persistence** [11]. Liveness ensures that transactions broadcasted by honest parties will be eventually included in the ledger, and persistence ensures that once a transaction is confirmed by one honest node, it will also be confirmed by all other honest

nodes (so that eventually it becomes final and immutable). Such properties are usually proven under certain assumptions, such as honest majority among protocol participants, etc. We refer the interested readers to the original Ouroboros paper [17] for an exhaustive list of assumptions and properties analysis.

Since the proposed consensus protocol also incorporates binding with the mainchain, it implies an additional assumption of the honest hashing power majority in the mainchain.

We suppose that under these assumptions the proposed protocol derives security guarantees provided by original Ouroboros and Nakamoto consensus protocols.

We want to emphasize that different types of sidechains may adopt different consensus protocols that better suit specific use cases (e.g., fast coin transferring support). A sidechain consensus protocol (including the one described in this section) is not the focus of this research and needs further analysis.

5.1.1 Withdrawal Epochs

As it has been described in section [4.1.2 Backward Transfers], the Cross-Chain Transfer Protocol introduces the notion of a **withdrawal epoch**⁸ (WE), which is defined as a fixed-length range of MC blocks (length is set upon SC creation). The concept of withdrawal epochs is essential for commanding backward transfers.

Following this design, we also introduce withdrawal epochs in a sidechain which coincide with the mainchain withdrawal epochs. A WE is defined as a range of SC blocks where the first and last blocks of the range are determined by references to the first and last MC blocks in the corresponding withdrawal epoch in the MC (see Fig. 8).

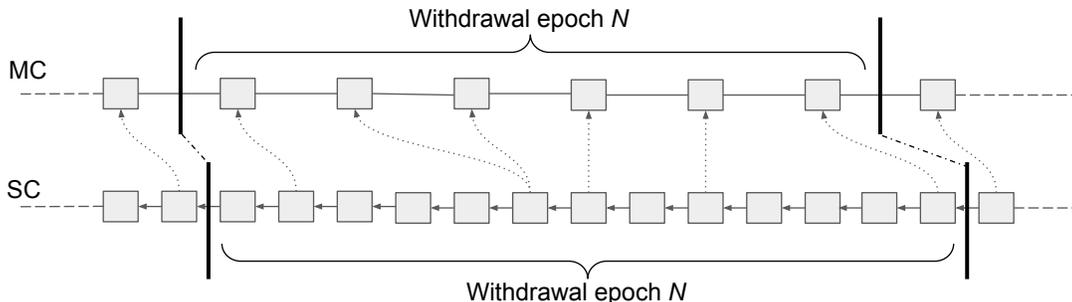


Figure 8: An example of a withdrawal epoch in the sidechain.

Even though a withdrawal epoch in the SC may have variable length (as it depends on when corresponding MC blocks will be referenced), the binding between chains allows to deterministically define the boundaries of the WE in the sidechain.

More formally, if the withdrawal epoch WE_i^{MC} of size len in the MC is defined by a sequence of blocks $WE_i^{MC} = (B_i^0, B_i^1, \dots, B_i^{len-1})$, then the corresponding withdrawal epoch in the SC can be determined as:

$$WE_i^{SC} = (SB_i^0, SB_i^1, \dots, SB_i^k),$$

where:

- SB_i^0 is an immediate descendant of the block SB_{i-1}^n which refers to the MC block B_{i-1}^{len-1} (the last one in the withdrawal epoch WE_{i-1}^{MC}); and
- SB_i^k is the block that refers to B_i^{len-1} .

⁸Note that withdrawal epochs are independent from epochs in the Ouroboros consensus protocol.

Note that to simplify implementation, it might be needed to restrict SC blocks to not refer to several MC blocks on the boundaries of the withdrawal epoch (i.e., if the SC block refers to B_i^{len-1} it cannot also refer to the next MC block B_{i+1}^0).

It is important to restate that the notion of the withdrawal epoch is independent from epochs in the Ouroboros consensus protocol.

5.2 Accounting Model and System State

The Latus blockchain adopts the UTXO-based accounting model [21] where the state is represented by a set of unspent outputs combined into a fixed-size Merkle tree (see Fig. 9). We call such a tree a **Merkle State Tree (MST)**. Lowercase mst_t stands to denote the root hash of the MST_t tree at the moment t .

The depth D_{MST} of the MST tree is a fixed system parameter that also constrains the total number of UTXOs that can exist in the system to be at most $2^{D_{MST}}$ (see Fig. 9).

We consider each leaf of the MST as a UTXO slot that can be “occupied” or “empty” at a given moment. We introduce the deterministic function $MST_Position(utxo_i)$ that returns the position of some unspent output $utxo_i$ if it is included in the tree. Note that the $utxo$ position does not depend on the current state of the MST.

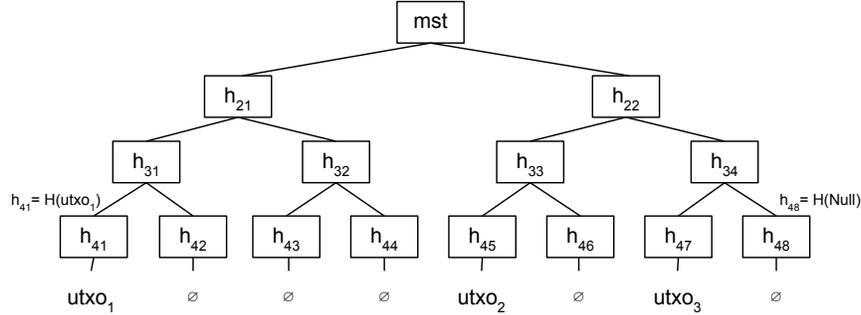


Figure 9: An example of the Merkle State Tree with $D_{MST} = 3$. A leaf of the tree is either an unspent output or $Null$ value. The tree contains 3 occupied and 5 empty slots. The function $MST_Position$ returns the position of a given $utxo$ in the tree, e.g., $MST_Position(utxo_2) = 4$.

The unspent transaction output (UTXO) is defined as a tuple $(addr, amount, nonce)$ where

- $addr$ is an address of the UTXO owner who possesses the corresponding private key that allows to spend it;
- $amount$ is the number of coins secured by the UTXO; and
- $nonce$ is a unique identifier of the UTXO.

5.2.1 System State

Provided with the MST structure, which is the core of a sidechain state, we define an overall SC system state at the moment t as a tuple:

$$state_t \stackrel{\text{def}}{=} (MST_t, backward_transfers_t),$$

where $backward_transfers$ is a list of backward transfers initiated in the current withdrawal epoch. $backward_transfers$ is transient and reset every new withdrawal epoch.

5.3 Transactional Model

There are 4 types of transactions defined in the Latus sidechain that realize basic payment functionality and cross-chain transfer protocol. To simplify the model, we consider them as logical transactions, though we stress that a real-world implementation can be optimized so that a single transaction on the blockchain may combine several logical transactions (even with different types).

The transactions are the following:

1. **Payment (PTx)** – transfers coins within the sidechain.
2. **Backward Transfer (BTTx)** – initiates transfer of coins from the SC to the MC.
3. **Forward Transfers (FTTx)** - receives coins transferred from the mainchain.
4. **Backward Transfer Requests (BTRTx)** - initiates coin transfer from the SC to the MC. In contrast to *BTTx*, *BTRTx* contains BTRs initially submitted in the mainchain and then synchronized to the SC.

Whereas *PTx* and *BTTx* are inherently SC-defined transactions (thus, submitted and processed in the sidechain), *FTTx* and *BTRTx* are MC-defined transactions (they encapsulate FTs and BTRs that are initially submitted to the MC). We describe each type in detail in the following sections.

5.3.1 Payment Transaction

We define a regular payment as a multi-input multi-output transaction [29]:

```
type PaymentTx {
  inputs: List[UTXO];
  signatures: List[Signature];
  outputs: List[UTXO];
}
```

where:

1. *inputs* are some unspent outputs from previous transactions, spending of which are authorized by *signatures*, and
2. the total coin's value of *inputs* is equal or greater than the total coin's value of *outputs*.

The state transition function *update* for the payment transaction is defined in the following way:

$$state_{i+1} = update(tx_{pay}, state_i),$$

where $state_{i+1}[backward_transfers]$ is unchanged and $state_{i+1}[MST]$ is derived from $state_i[MST]$ by

1. removing all UTXOs that are inputs in tx_{pay} and substituting them with *Null* to produce MST_i^- ; and
2. sequentially adding to MST_i^- all UTXOs that are outputs in tx_{pay} according to $MST_Position(utxo_j)$.

5.3.2 Forward Transfers Transaction

Forward transfers allow one to send coins from the mainchain to a sidechain. As such, FTs are first submitted to the MC and processed there (destroying coins) and then, by means of deterministic synchronization, are included and processed in the sidechain. Recall from [4.1.1 Forward Transfers] the basic structure of a forward transfer on the mainchain side:

$$FT \stackrel{\text{def}}{=} (\text{ledgerId}, \text{receiverMetadata}, \text{amount}).$$

receiverMetadata is defined by the sidechain construction and in Latus it is just a receiver address and a payback address on the MC needed in case of transfer failure:

$$\text{receiversMetadata} \stackrel{\text{def}}{=} (\text{receiverAddr}, \text{paybackAddr}).$$

A single MC block may contain several forward transfers related to different sidechains. The sidechain will synchronize FTs present in the referenced MC block by including a special *ForwardTransfers* transaction (FTTx) in the SC block. Such FTTx specifies all forward transfers from the referenced MC block that are related to this specific sidechain. From the sidechain perspective, we can consider FTTx as a coinbase transaction (the one that creates new coins [28]) that is authorized by the mainchain.

We assume that a particular forward transfer may fail so that coins cannot be received by the sidechain. In this case, coins are sent back to the mainchain by creating a corresponding backward transfer. It is done automatically upon FTTx execution in the sidechain.

The reasons for FT failure can be different. For instance, FT's *receiverMetadata* may be malformed (recall that the MC does not validate semantics of FT's *receiverMetadata*) or some other sidechain-specific failures occur (e.g., it may happen that $MST_Position(\text{output}_{new})$ maps newly created output to an already occupied slot in MST_i , thus causing a collision).

The basic structure of the *ForwardTransfers* transaction is the following:

```
type ForwardTransfersTx (mcid: BlockID, ft: List[FT]) {
  outputs: List [UTXO];
  rejectedTransfers : List [BackwardTransfer]
}
```

where:

- mcid* – an identifier of the MC block whose forward transfers are synchronized;
- ft* – a list of forward transfers from the MC block *mcid* related to the sidechain where FTTx occurs;
- outputs* – outputs created for the transferred coins; each valid forward transfer spawns a corresponding output with the same amount of coins;
- rejectedTransfers* – a list of backward transfers for failed forward transfers; each failed forward transfer spawns a corresponding backward transfer with the same amount of coins.

The state transition function update for the FTTx is defined in the following way:

$$\text{state}_{i+1} = \text{update}(tx_{FT}, \text{state}_i),$$

where:

1. $state_{i+1}[MST]$ is derived from $state_i[MST]$ by sequentially adding all UTXOs that are outputs in tx_{FT} according to $MST_Position(utxo_j)$;
2. $state_{i+1}[backward_transfers]$ is derived from $state_i[backward_transfers]$ by appending $rejectedTransfers$ from tx_{FT} .

Note that failed forward transfers are recovered with the backward transfer mechanism through a withdrawal certificate at the end of the epoch. Recall that the MC knows nothing about the SC state and cannot know that an FT is failed; thus, we use the standard mechanism to reclaim coins in the MC.

5.3.3 Backward Transfer Transaction

A backward transfer transaction (BTTx) allows one to create a request for a backward transfer in the sidechain that will be included in the next withdrawal certificate and then passed and processed in the mainchain.

```

type BackwardTransferTx {
  inputs: List [UTXO];
  signatures: List [Signature];
  backwardTransfers: List [BackwardTransfer];
}

```

where:

1. $inputs$ are some unspent outputs from previous transactions (spending of which is authorized by $signatures$);
2. $backwardTransfers$ are data about receivers of coins on the mainchain side; recall from [4.1.2 Backward Transfers] that the basic structure of a backward transfer is $BT \stackrel{\text{def}}{=} (receiverAddr, amount)$;
3. the total coin value of $inputs$ is equal or greater to the total coin value of $backwardTransfers$.

The state transition function $update$ for a backward transfer transaction is defined as:

$$state_{i+1} = update(tx_{BT}, state_i),$$

where:

1. $state_{i+1}[MST]$ is derived from $state_i[MST]$ by removing all UTXOs that are inputs in tx_{BT} ; and
2. $state_{i+1}[backward_transfers]$ is derived from $state_i[backward_transfers]$ by appending $backwardTransfers$ from tx_{BT} .

Essentially, we can consider $backwardTransfers$ in tx_{BT} as specialized outputs that are unspendable on the sidechain but used to reclaim coins in the mainchain (when transferred by means of a withdrawal certificate). In this respect, BTTx transaction is a special case of regular payment transaction where all outputs are backward transfers.

More details about the entire backward transfer flow can be found in [5.5.3 Backward Transfers].

5.3.4 Backward Transfer Requests Transaction

The *Backward Transfer Request* (BTR), which is submitted to the MC, is similar to BTTx in the sense that it allows one to create a request that will result in a backward transfer in the next withdrawal certificate if the request is legitimate (e.g., claimed coins were present at the moment of BTR inclusion in the SC block). The difference from BTTx is that BTR is submitted in the mainchain and is used in situations when BTTx cannot be used for some reason. Recall from [4.1.2.1 Mainchain Managed Withdrawals] the basic structure of the BTR on the mainchain side:

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

proofdata and SNARK *proof* are defined by sidechain construction. In Latus, *proofdata* contains an unspent output that should be consumed in the SC to provide coins for transferring:

$$proofdata = \{utxo\}.$$

The spending right for the *utxo* should be enforced by the *proof* which is validated upon submission in the MC.

Similar to forward transfers, a single MC block may contain several BTRs. The sidechain synchronizes BTRs by including in the SC block a special *BackwardTransferRequests* transaction (BTRTx) that contains all BTRs relevant to this sidechain from the referenced MC block. From a sidechain perspective, we can consider BTRTx as an aggregated transaction where each BTR represents a separate backward transfer.

Some BTRs from BTRTx may be invalid when they are synced to the sidechain (e.g., a malicious user may try to spend the same utxo directly in the sidechain before BTR is synced (double-spend problem). Such BTRs are rejected by the sidechain (rejection means that they do not spawn corresponding backward transfers and do not affect the state).

The basic BTRTx structure on the SC side is the following:

```
type BackwardTransferRequestsTx (mcid: BlockId, btr: List[BTR]) {
  inputs: List[UTXO];
  backwardTransfers: List[BackwardTransfer];
}
```

where:

- mcid* – an identifier of the MC block whose BTRs are synchronized;
- btr* – a list of backward transfer requests from the MC block *mcid* related to this sidechain;
- inputs* – a combined list of UTXOs derived from *btr.proofdata* of each valid BTR;
- backwardTransfers* – a list of backward transfers for valid BTRs.

The state transition function update for a BTRTx transaction is:

$$state_{i+1} = update(tx_{BTR}, state_i),$$

where:

1. $state_{i+1}[MST]$ is derived from $state_i[MST]$ by removing all UTXOs that are inputs in tx_{BTR} ;
2. $state_{i+1}[backward_transfers]$ is derived from $state_i[backward_transfers]$ by appending *backwardTransfers* from tx_{BTR} .

Note that correct processing of BTRs in the sidechain is to be enforced by a withdrawal certificate SNARK proof.

5.4 State Transition Proof

In [5.3 Transactional Model], we defined four types of transactions that represent basic state transitions in our sidechain system. Given that all transactions are applied sequentially in an order defined by blocks containing them, we can consider a merged state transition for a sequence of transactions from several blocks:

$$state_{i+k} = update([tx_1, \dots, tx_k], state_i) = update(tx_k, update(tx_{k-1}, \dots, update(tx_1, state_i))).$$

In particular, we are interested in merging transitions for the whole withdrawal epoch and proving that the top-level merged transition is correct. It can be accomplished using the recursive SNARKs composition for state transitions defined in [Def. 2.5].

The main idea is to construct a single SNARK proof of transition for the whole withdrawal epoch which then can be attached to a withdrawal certificate proving to the mainchain the validity of everything that has happened in the sidechain – including certificate backward transfers – without actually revealing any details except state snapshots (in a form of simple hashes) before and after transition.

We do not go deeply into the details of the SNARKs architecture which is quite sophisticated in this case and requires separate writing to be properly explained; instead, we are going to provide the basic idea of constructing such proof and how it is going to be used.

Let us denote by $s_i = H(state_i)$ the hash value that represents $state_i$. Note that it must be an efficient hashing procedure as it should be implemented for a SNARK arithmetic constraint system. For instance, we can consider $H(\cdot)$ as a root hash of a Merkle tree that contains all the data from $state_i$.

Let us assume that for each basic state transition (represented by tx_{pay} , tx_{FT} , tx_{BT} , and tx_{BTR}) we have a corresponding *Base* SNARK [Def. 2.5] which proves the correct state transition for a single tx_a , $a \in \{pay, FT, BT, BTR\}$:

$$\begin{aligned} \pi_a^{Base} &\leftarrow Prove(pk_a^{Base}, (s_i, s_{i+1}), (tx_a)), \\ true/false &\leftarrow Verify(vk_a^{Base}, (s_i, s_{i+1}), \pi_a^{Base}). \end{aligned}$$

Also, let us assume that we have a *Merge* SNARK which takes two proofs of adjacent state transitions (*Base* or *Merge*) and combines them into a single proof:

$$\begin{aligned} \pi^{Merge} &\leftarrow Prove(pk^{Merge}, (s_i, s_{i+k}), (s_{i+j}, \pi_1^b, \pi_2^c)), \\ true/false &\leftarrow Verify(vk^{Merge}, (s_i, s_{i+k}), \pi^{Merge}). \end{aligned}$$

where:

- $b, c \in \{Base, Merge\}$;
- π_1^b proves that there exist such tx_1, \dots, tx_j so that $state_{i+j} = update([tx_1, \dots, tx_j], state_i)$;
- π_2^c proves that there exist such tx_{j+1}, \dots, tx_k so that $state_{i+k} = update([tx_{j+1}, \dots, tx_k], state_{i+j})$.

Provided with this construction, we can recursively build a single SNARK proof of state transition for a whole withdrawal epoch from the sequence of basic transitions. This process is visualized in figures 10 and 11.

Figure 10 demonstrates the recursive construction of a state transition proof for a single sidechain block. Note that the scheme is simplified; in reality, the SNARKs composition is more sophisticated and the proof itself attests not only for the correctness of basic transitions but also for the validity of the SC block, the validity of included MC block references, their contiguity, etc.

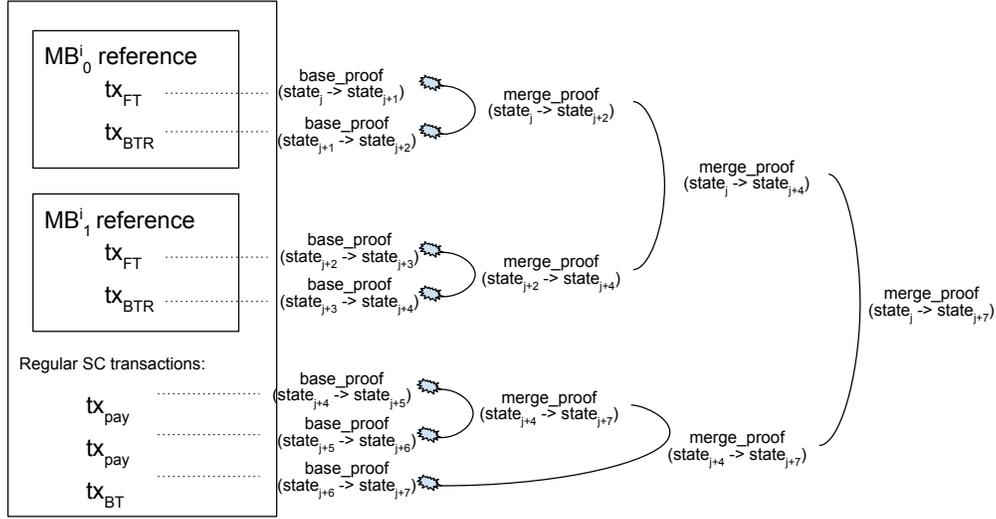


Figure 10: Recursive composition of state transition proofs for the whole SC block. At the bottom level, there are proofs for basic transitions (represented by transactions included in the block) which are then recursively merged into a single proof.

Figure 11 demonstrates the recursive construction of a state transition proof for an entire withdrawal epoch. Provided with the proofs of state transitions for blocks from the previous step, now they are merged to generate a single proof for the whole epoch which is used to construct a final proof for a withdrawal certificate.

In a nutshell, each withdrawal certificate $WCert_i$ for epoch i commits to the new state $state_{len}^i$ produced by applying all blocks belonging to epoch i and proves correct transition from the $state_{len}^{i-1}$ committed by the previous withdrawal certificate. This also involves proving that all MC blocks belonging to the withdrawal epoch are referenced and all MC transactions related to this sidechain are processed. As forward and backward transfers are among basic transitions they will also be proven.

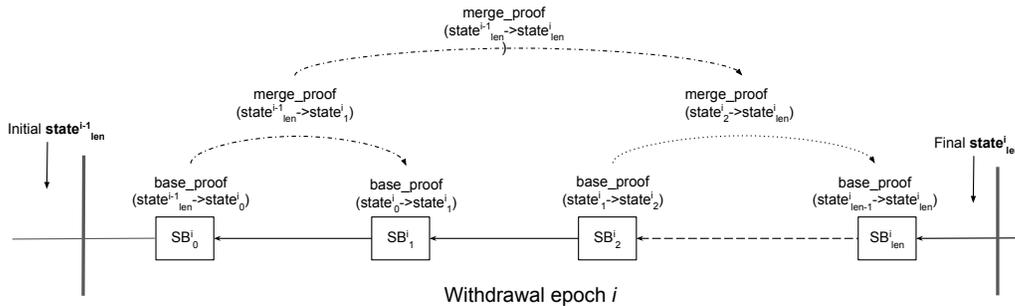


Figure 11: Recursive composition of state transition proofs for the whole withdrawal epoch. State transitions for SC blocks are considered as base transitions though they themselves are recursively constructed from basic transitions (see Fig. 10).

Again, we stress that this description is greatly simplified just to show the basic idea of recursive SNARKs composition for state transitions.

5.4.1 Performance and Incentives

Generating a SNARK proof for each basic transition and then merging them together requires a significant amount of computation. This task cannot be solely levied upon forgers or WCert issuers. Currently, we are investigating different approaches.

One of the possible solutions is to introduce a special dispatching scheme that assigns generation of proofs randomly to interested parties who then do these tasks in parallel and submit generated proofs to the blockchain. An incentive scheme provides a reward for each valid submission.

We consider this as a separate topic of research which we do not elaborate in this paper.

5.5 Cross-Chain Transfer Protocol

In the previous sections, we described the consensus protocol, accounting model, and transactional model of the proposed sidechain construction. In this section, we will focus on the structure of the cross-chain transfer protocol on the sidechain side which is based on those components.

5.5.1 Mainchain Block Reference

In [5.1 Consensus Protocol], we briefly described the synchronization procedure between the mainchain and sidechain, which relies on MC block referencing. Here, we describe the reference structure in a more detailed way.

Recall that an MC block header contains the **SCTxsCommitment** [4.1.3 Sidechain Transactions Commitment] field that commits to all SC-related transactions/outputs in that block:

```
type MCBlockHeader {
  prevBlock: BlockId
  height: Int
  ...
  scTxsCommitment: Hash
  ...
}
```

ScTxsCommitment is a root hash of a Merkle tree where one of the subtrees is the Merkle tree of transactions related to the sidechain that referenced the block (see Fig. 12).

The structure of the mainchain block reference is the following:

```
type MCBlockReference {
  header: MCBlockHeader
  mproof: Option[MerkleProof]
  proofOfNoData: Option[MerkleProof[]]
  forwardTransfers: Option[FTTx]
  btRequests: Option[BTRTx]
  wcert: Option[WCert]
}
```

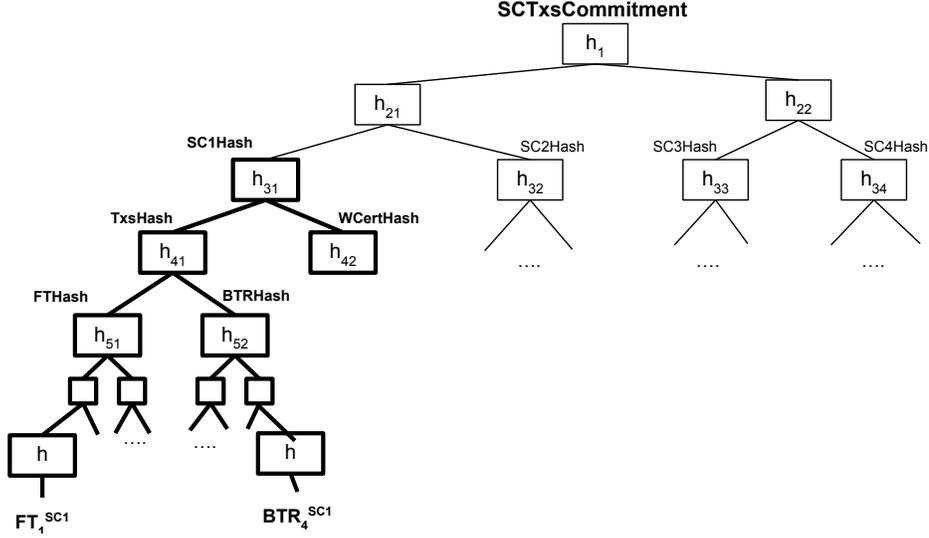


Figure 12: An example of the sidechain transactions commitment tree. One of the intermediate nodes (**SC1Hash**) is a root hash of the subtree that commits to all transactions related to the sidechain *SC1*.

where:

- header* – a header of the MC block that is referenced;
- mproof* – optional field: in case the MC block includes at least one transaction related to this SC, *mproof* will contain a Merkle proof [Def. 2.2] for the intermediate node in the sidechain transactions commitment tree that is a root of a subtree of transactions related to this sidechain (by the example in Fig. 12: the subtree root for the sidechain *SC1* is h_{31} and the corresponding Merkle proof is the tuple of nodes $\{h_{32}, h_{22}\}$); in case the MC block has no transactions related to this SC, the *mproof* must be *Null*;
- proofOfNoData* – optional field: in case the MC block has no transactions related to this SC, *proofOfNoData* contains the Merkle proof(s) necessary to prove that this *ledgerId* was not part of the *SCTxsCommitment* tree;
- forwardTransfers* – optional field: it is either a *ForwardTransfers* transaction [5.3.2 [Forward Transfers Transaction](#)] (if the MC block contains at least one forward transfer to this sidechain) or otherwise *Null*;
- btRequests* – optional field: it is either a *BackwardTransferRequests* transaction [5.3.4 [Backward Transfer Requests Transaction](#)] (in case the MC block contains at least one backward transfer request to this sidechain) or otherwise *Null*;
- wcert* – optional field: it is either a withdrawal certificate (in case the MC block contains the withdrawal certificate related to this sidechain) or otherwise *Null*.

Provided with *mproof*, *forwardTransfers*, *btRequests*, and *wcert* fields, the *SCTxsCommitment* can be reconstructed and verified against the *scTxCommitment* field included in the MC block

header. It allows to verify that all SC-related transactions were correctly synchronized from the MC block without the need to download and verify its body. Moreover, we can construct a SNARK proving that the MC block reference has been correctly processed and that all SC-related transactions have been applied – it is an essential part of constructing a state transition proof for a withdrawal epoch [5.4 State Transition Proof].

5.5.2 Forward Transfers

In [4.1.1 Forward Transfers] and [5.3.2 Forward Transfers Transaction], we have already discussed most of the details related to the forward transfer design both on the mainchain and sidechain sides. Here we combine everything.

In general, it looks as follows: an MC to SC transfer is represented by a pair of transactions which we can consider as “*sending*” and “*receiving*”. “*Sending*” is done on the mainchain side by means of the forward transfer defined in [4.1.1 Forward Transfers] and “*receiving*” is done on the sidechain side by means of aggregated *ForwardTransfers* transaction defined in [5.3.2 Forward Transfers Transaction]. While “*sending*” destroys coins in the mainchain, “*receiving*” creates the corresponding number of coins in the sidechain.

Forward transfers submitted to the mainchain become available in the sidechain at the moment the MC block containing them is referenced in the sidechain (see Fig. 13). With the MC block reference [5.5.1 Mainchain Block Reference], a *ForwardTransfers* transaction (FTTx) is included in the SC block (if there are any FTs).

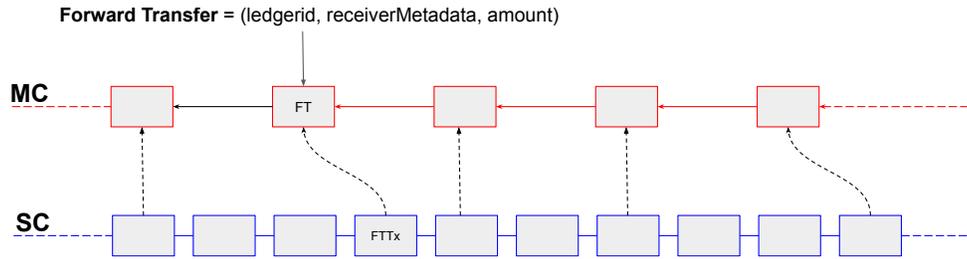


Figure 13: Forward transfers syncing from the mainchain to the sidechain.

The consistency of forward transfers included in a sidechain FTTx is verified by recalculating the *FTHash* (Fig. 12) and checking the **SCTxsCommitment** following the procedure described in [5.5.1 Mainchain Block Reference].

5.5.3 Backward Transfers

In general, backward transfer is a transfer of coins in the opposite direction: from the sidechain to the mainchain. This operation is more sophisticated and thus requires several sub-protocols to provide sufficient security and reliability.

There are three ways to withdraw coins from the sidechain to the mainchain:

1. **Regular withdrawal** is a standard mechanism that is used under normal conditions. It implies the usage of a backward transfer transaction [5.3.3 Backward Transfer Transaction] and a withdrawal certificate to transfer coins to the mainchain.
2. **Backward transfer request** is similar to the regular withdrawal with that difference that it is initially submitted to the mainchain [4.1.2.1 Mainchain Managed Withdrawals] and then synchronized to the sidechain by means of [5.3.4 Backward Transfer Requests Transaction]. The coins are transferred to the mainchain with a withdrawal certificate.

3. **Ceased sidechain withdrawal** is a mechanism that is used when the sidechain is no longer operating. This type of withdrawal does not use withdrawal certificates and supposes direct handling by the mainchain.

The first two types of withdrawals (regular and BTR) use the standard mechanism for backward transfers - **withdrawal certificate**. Most of the details related to their submission and processing have already been discussed in [4.1.2.1 Mainchain Managed Withdrawals], [5.3.3 Backward Transfer Transaction], and [5.3.4 Backward Transfer Requests Transaction]. The basic principle is summarized in figure 14.

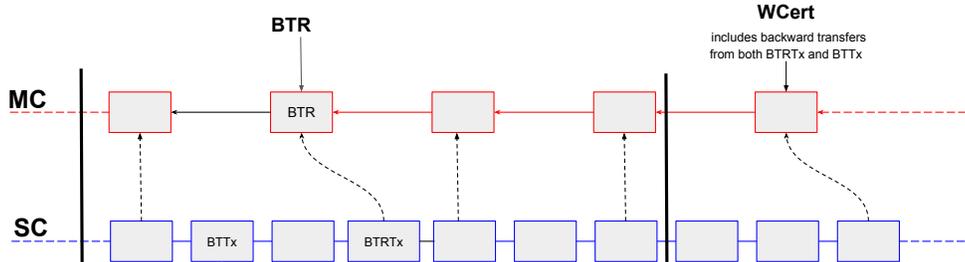


Figure 14: Withdrawing coins with BT and BTR transactions.

Regular withdrawal. A special *BackwardTransfer* transaction is submitted to the sidechain by a user who wants to transfer coins. This transaction destroys coins in the sidechain. At the end of the withdrawal epoch, all backward transfers are collected in a withdrawal certificate which is submitted to the mainchain where it is processed, and the corresponding number of coins is created in the mainchain.

Backward transfer request. BTRs are submitted to the mainchain and synchronized to the sidechain by means of a *BackwardTransferRequests* transaction (analogously to forward transfers). The consistency of BTRs included in the sidechain is verified by recalculating the *BTRHash* (Fig. 12) and checking its presence in the *SCTxsCommitment* tree following the procedure described in [5.5.1 Mainchain Block Reference]. After the BTR is synchronized to the sidechain, it is processed as regular withdrawal through a withdrawal certificate.

In the following sections, we will discuss more deeply the structure and generation of a withdrawal certificate and BTR. We will also separately discuss CSW as it is conceptually different from the first two withdrawal methods.

5.5.3.1 Withdrawal Certificate

Withdrawal certificate is a pivotal component of the backward transfer flow. Recall the basic structure of a withdrawal certificate that is defined by the mainchain [4.1.2 Backward Transfers]:

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

While *ledgerId* and *epochId* are global parameters known to the mainchain, the semantics of *quality*, *proofdata*, and *proof* are defined by the sidechain. The withdrawal certificate is created once per a withdrawal epoch and includes all backward transfers that have been submitted during the epoch.

Quality. The quality parameter is used by the mainchain to determine what WCert should be

adopted in case several have been submitted for the same epoch. From the MC side, this is just an integer value which can be compared with quantities from other certificates. In the Latus sidechain, we define the quality to be the height of the blockchain up until which the WCert proves state transition.

Backward Transfers List. *BTList* is a list of backward transfers collected during a withdrawal epoch for which the certificate is created:

$$BTList = state_i[backward_transfers],$$

where $state_i$ is the state of the sidechain after applying the last block in the withdrawal epoch.

Withdrawal certificate proof

Withdrawal certificate proof is a SNARK proof that validates compliance of the certificate with a set of predefined rules.

As it is defined in [Def. 2.3], a SNARK is a proving system. Its particular instantiation is specified by a set of arithmetic constraints defining the verification rules. Each sidechain specifies its own set of constraints for the withdrawal certificate SNARK, thus establishing its own rules.

The basic interface for the SNARK prover and verifier is the following:

$$\begin{aligned} proof &\leftarrow Prove(pk_{WCert}, public_input, witness), \\ true/false &\leftarrow Verify(vk_{WCert}, public_input, proof). \end{aligned}$$

A particular instantiation of the SNARK proving system is determined by a pair of keys – proving key pk_{WCert} and verifying key vk_{WCert} . Verifying key is registered upon sidechain creation and cannot be changed during the SC lifetime. It completely defines the rules of the withdrawal certificate validation (including the semantics of the public input and witness for the prover and verifier).

Recall from [Def. 4.4] that *public_input* for the WCert SNARK is comprised of two parts:

$$public_input \stackrel{\text{def}}{=} (wcert_sysdata, MH(proofdata)),$$

where *wcert_sysdata* is a set of arguments enforced directly by the mainchain⁹:

$$wcert_sysdata \stackrel{\text{def}}{=} (quality, MTHash(BTList), H(B_{len}^{i-1}), H(B_{len}^i));$$

and *proofdata* is a set of arguments defined by the sidechain construction and passed along the withdrawal certificate. In the Latus sidechain, it is defined as follows:

$$proofdata \stackrel{\text{def}}{=} (H(SB_{last}^i), H(state_{SB_{last}^i}[MST]), mst_delta),$$

where

- $H(SB_{last}^i)$ – a hash of the last sidechain block in the epoch i for which the certificate is created;
- $H(state_{SB_{last}^i}[MST])$ – a root hash of the MST tree derived after applying SB_{last}^i ; note that by including $H(state_{SB_{last}^i}[MST])$ in *proofdata*, the withdrawal certificate commits to the updated sidechain state;
- mst_delta – a bit vector of MST modifications; given that the MST is a fixed size Merkle tree, *mst_delta* is also a fixed-size bit vector where each bit represents a particular leaf in the tree; the bit is set to “1” if the MST leaf has been modified at least once during the epoch, otherwise it is “0” (see example in [Appendix A]).

⁹These arguments are explained in [4.1.2 Backward Transfers]

mst_delta is used for proving that some $utxo$ has not been spent since some moment in the past (this is particularly useful for preventing data availability attacks as it allows creating mainchain managed withdrawals without knowing the current sidechain state). E.g., to prove this, one would need to provide a $utxo$ together with a Merkle proof of its inclusion in some $state_{SB_{last}^k} [MST]$ committed in one of the previous certificates and a list of mst_delta 's from the following certificates where the corresponding bit has not been triggered to "1".

In the Latus sidechain construction, a withdrawal certificate proof enforces the following rules:

WCert SNARK Statement

- SB_{last}^i is the last block of the withdrawal epoch i for which the certificate is created.
- SB_{last}^i is connected to the SB_{last}^{i-1} from the previous withdrawal certificate by a valid chain of blocks.
- $H(state_{SB_{last}^i} [MST])$ is a valid root of the MST for $state_{SB_{last}^i}$.
- Assuming that after applying the block SB_{last}^{i-1} the sidechain state is $state_{SB_{last}^{i-1}}$ and after the block SB_{last}^i the state is $state_{SB_{last}^i}$, the proof verifies correct transition from $state_{SB_{last}^{i-1}}$ to $state_{SB_{last}^i}$ which means that all transactions from the subchain $[SB_0^i, \dots, SB_{last}^i]$ are correctly processed according to the rules from [5.3 Transactional Model].
- MC blocks from range $[B_0^i, \dots, B_{last}^i]$ are referenced from the sidechain blocks $[SB_0^i, \dots, SB_{last}^i]$ (this also implies that all SC-related transactions from these blocks have been processed).
- $BTLlist$ is a valid list of backward transfers that corresponds to $state_{SB_{last}^i} [backward_transfers]$.
- $quality$ parameter is the height of the block SB_{last}^i .
- mst_delta is a bit vector that reflects changes in MST between $state_{SB_{last}^{i-1}} [MST]$ and $state_{SB_{last}^i} [MST]$.

In general, a withdrawal certificate proof validates correct transition for a range of blocks that belongs to the withdrawal epoch and that this range is adjacent to the range committed in the previous withdrawal certificate. This includes proving the correctness of backward transfers.

Given that all state transitions are proved, it becomes infeasible to create a malicious backward transfer (without creating a corresponding transaction in the sidechain), and it is infeasible to create new coins on the sidechains without real forward transfers.

5.5.3.2 Backward Transfer Request

In [5.3.4 Backward Transfer Requests Transaction], we have already discussed how BTRs are submitted and processed in the sidechain. Here, we only provide details about the SNARK proof included in a BTR.

Recall that the BTR structure has been defined as follows [Def. 4.5]:

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

proofdata is defined by the Latus construction as:

$$proofdata = \{utxo\},$$

where *utxo* is an unspent output that holds coins that a user wants to withdraw. The basic idea is that the *proof* should validate the user's right to withdraw this *utxo* and that this *utxo* is present in the sidechain state MST committed by the last withdrawal certificate included in the mainchain.

Note that the BTR SNARK *proof* is validated by the mainchain upon BTR submission. Even though it verifies that the withdrawn *utxo* has been present in the last committed SC state, it cannot guarantee that it will remain valid at the moment BTR will be synchronized to the sidechain. This proof serves more like a pre-validation for the BTR in the mainchain to impede submission of wittingly invalid requests.

The basic interface for the SNARK prover and verifier is the following [Def. 4.5]:

$$\begin{aligned} proof &\leftarrow Prove(pk_{BTR}, public_input, witness), \\ true/false &\leftarrow Verify(vk_{BTR}, public_input, proof). \end{aligned}$$

The verifying key vk_{BTR} is registered upon sidechain creation. It defines the rules of the BTR validation (including the semantics of the public input and witness for the prover and verifier).

The *public_input* comprises two parts [Def. 4.5]:

$$\begin{aligned} public_input &\stackrel{\text{def}}{=} (btr_sysdata, MH(proofdata)), \\ btr_sysdata &\stackrel{\text{def}}{=} (H(B_w), nullifier, receiver, amount), \end{aligned}$$

where $H(B_w)$ is the hash of the MC block with the latest withdrawal certificate (at the moment when BTR is included in the mainchain), *receiver*, *amount*, and *nullifier* are taken from the BTR itself. Note that *btr_sysdata* is enforced by the mainchain so its parameters cannot be manipulated by the BTR issuer.

A BTR proof enforces the following rules:

BTR SNARK Statement

- $H(B_w)$ is the hash of the mainchain block where the last certificate $WCert_w$ has been submitted for this sidechain.
- $utxo \in state_w[MST]$, where $state_w[MST]$ has been committed in $WCert_w$.
- The BTR issuer has rights to spend this *utxo* (i.e., possesses the corresponding private key).
- *amount* is equal to the *utxo.amount*.
- *nullifier* is the hash of the *utxo*.
- *receiver* is the address of the receiver in the mainchain.

5.5.3.3 Ceased Sidechain Withdrawal

CSWs are used to allow sidechain stakeholders to withdraw coins from a ceased sidechain.

As it has been defined in [Def. 4.6], a ceased sidechain withdrawal is submitted to the mainchain as a special transaction and performs a direct payment in the mainchain. Recall the basic structure of the ceased sidechain withdrawal that is defined by the mainchain [Def. 4.6]:

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

The main prerequisite for CSW validity is the existence of the claimed coins in the sidechain state committed by the last withdrawal certificate. A sidechain user should point to the specific unspent output from $state_{SB^i}^{last}[MST]$ and authorize its spending. Basically, it is the same SNARK that is used for the BTR [5.5.3.2 Backward Transfer Request]; the difference is that now it authorizes direct payment in the mainchain, whereas in the BTR, it is essentially a pre-validation.

We will not dive deeply into the SNARK construction for the CSW as technically it is completely the same as for the BTR.

In general, the CSW proof validates that a submitter owns the utxo with a particular amount of coins at the moment of the sidechain halt. Also, it enforces a nullifier which is a unique identifier of the withdrawn utxo. Nullifiers are tracked by the mainchain to prevent withdrawal of the same coins twice.

6 Conclusions

The concept of sidechains has been acknowledged as an appealing solution for enhancing existing blockchain systems. It allows creating platforms and applications that are bound to the mainchain without imposing significant burden. Yet, we have not seen wide adoption of this concept. We believe that the value of sidechains as a scalability solution is underestimated and seek to develop this area.

In this paper, we introduced Zendoo, a universal construction for blockchain systems that enables the creation and communication with different sidechains without knowing their internal structure. We also provided a specific sidechain construction, Latus, that leverages zk-SNARK techniques to establish decentralized and verifiable cross-chain transfers.

We consider this as a research paper whose subject is still under ongoing research. In future publications, we plan to uncover more details about specific components and properties of the proposed sidechain construction.

7 Acknowledgments

We would like to express great appreciation to Maurizio Binello and Andrey Sobol for participating in technical discussions.

We would also like to thank Rob Viglione, Daniele Di Benedetto, Marcelo Kaihara, Luca Cermelli, and Lyudmila Kovalchuk for reviewing and providing valuable comments.

References

- [1] Cosmos network, 2018. <https://cosmos.network/docs/>.
- [2] Ethereum. a next-generation smart contract and decentralized application platform., 2018. <https://github.com/ethereum/wiki/wiki/White-Paper>.

- [3] Rootstock: smart contracts on bitcoin network, 2018. <https://www.rsk.co/>.
- [4] Andreas M. Antonopoulos. Mastering bitcoin (second edition). O'Reilly Media, Inc., 2017.
- [5] A. Back, M. Corallo, L. Dashjr, M. Friedenbach, G. Maxwell, A. Miller, A. Poelstra, J. Timón, and P. Wuille. Enabling blockchain innovations with pegged sidechains, 2014. <https://blockstream.com/sidechains.pdf>.
- [6] 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.
- [7] Sean Bowe and Ariel Gabizon. Making groth's zk-snark simulation extractable in the random oracle model. IACR Cryptology ePrint Archive, 2018:187, 2018.
- [8] 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>.
- [9] J. Dilley, A. Poelstra, and J. Wilkins. Strong federations: An interoperable blockchain solution to centralized third party risks. arXiv:1612.05491, 2016. <https://arxiv.org/abs/1612.05491>.
- [10] 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.
- [11] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. *Advances in Cryptology - EUROCRYPT 2015. Lecture Notes in Computer Science*, vol 9057. Springer, Berlin, Heidelberg, 2015.
- [12] Alberto Garoffolo and Robert Viglione. Sidechains: Decoupled consensus between chains. arXiv:1812.05441, 2018. <https://arxiv.org/abs/1812.05441>.
- [13] P. Gazi, A. Kiayias, and D. Zindros. Proof-of-stake sidechains. *Proceedings of the IEEE Symposium on Security & Privacy*. IEEE Computer Society Press, 2019.
- [14] Alex Gluchowski. Zk rollup: scaling with zero-knowledge proofs. Matter Labs, 2019. <https://pandax-statics.oss-cn-shenzhen.aliyuncs.com/statics/1221233526992813.pdf>.
- [15] Oded Goldreich. *The foundations of cryptography - volume 1, basic techniques*. Cambridge University Press, 2001.
- [16] A. Kiayias and D. Zindros. Proof-of-work sidechains. Cryptology ePrint Archive, Report 2018/1048, 2018. <https://eprint.iacr.org/2018/1048>.
- [17] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. *CRYPTO 2017, Part I*, volume 10401 of LNCS, pages 357–388. Springer, Heidelberg, 2017.
- [18] S. Lerner. Drivechains, sidechains and hybrid 2-way peg designs., 2016. https://docs.rsk.co/Drivechains_Sidechains_and_Hybrid_2-way_peg_Designs_R9.pdf.
- [19] Izaak Meckler and Evan Shapiro. Coda : Decentralized cryptocurrency at scale, 2018. <https://cdn.codaprotocol.com/v2/static/coda-whitepaper-05-10-2018-0.pdf>.

- [20] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008. <https://bitcoin.org/bitcoin.pdf>.
- [21] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. Bitcoin and cryptocurrency technologies: A comprehensive introduction. Princeton University Press, Princeton, NJ, USA, 2016.
- [22] J. Poon and V. Buterin. Plasma: Scalable autonomous smart contracts. <http://plasma.io/>.
- [23] Merkle R.C. A digital signature based on a conventional encryption function. Advances in Cryptology — CRYPTO 1987. Lecture Notes in Computer Science, vol 293. Springer, Berlin, Heidelberg, 1988.
- [24] Nicholas Stifter, Aljosha Judmayer, Philipp Schindler, Alexei Zamyatin, and Edgar Weippl. Agreement with satoshi – on the formalization of nakamoto consensus. Cryptology ePrint Archive, Report 2018/400, 2018. <https://eprint.iacr.org/2018/400>.
- [25] P. Sztorc. Drivechain - the simple two way peg, november 2015., 2015. <http://www.truthcoin.info/blog/drivechain/>.
- [26] The Horizen Team and Community. Horizen: A blockchain platform for fully customizable decentralized applications, 2019. <https://www.horizen.global/assets/files/Horizen-White-Paper.pdf>.
- [27] S. Thomas and E. Schwartz. A protocol for interledger payments, 2016. <https://interledger.org/interledger.pdf>.
- [28] Bitcoin Wiki. Coinbase. <https://en.bitcoin.it/wiki/Coinbase>.
- [29] Bitcoin Wiki. Transaction. <https://en.bitcoin.it/wiki/Transaction>.
- [30] Gavin Wood. Polkadot:vision for a heterogeneous multi-chain framework, 2016. <https://polkadot.network/Polkadot-lightpaper.pdf>.

Appendix A MST Delta

Here, we provide an example of how the *mst_delta* value from a withdrawal certificate [5.5.3.1 [Withdrawal Certificate](#)] is calculated and give some explanations on why it is needed. Note that this relates only to the Latus sidechain construction.

In general, *mst_delta* shows which leaves have been changed between two Merkle state trees [5.2 [Accounting Model and System State](#)] MST_i and MST_j , $i < j$ (e.g., in case of a withdrawal certificate, these are MSTs committed by the previous certificate and the current one which shows how the system state changed during the epoch). *mst_delta* is a bit vector that shows what leaves of the MST_i have been changed in MST_j .

Let us consider the MST of depth $D_{MST} = 3$ which has an initial state MST_0 (see Fig. 15).

The MST can contain up to eight unspent outputs (equal to the number of leaves). At the moment MST_0 , the tree contains three UTXOs $\{utxo_1, utxo_2, utxo_3\}$ which are assigned to leaf nodes 0, 4, 6 correspondingly.

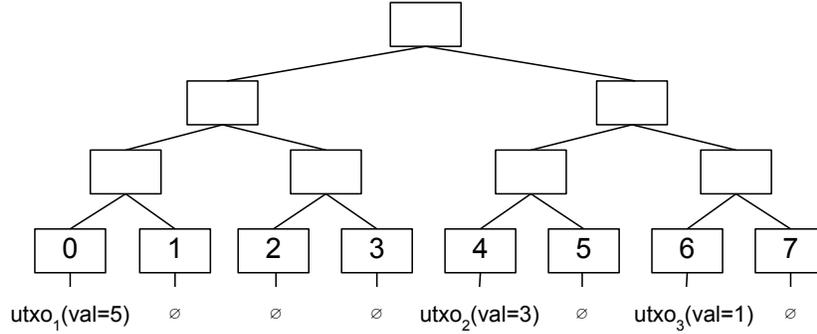


Figure 15: Merkle state tree MST_0 .

Let us assume that we have two transactions tx_1 and tx_2 such that:

```

tx1 = {
  inputs: {utxo1}
  outputs: {utxo4(val=2), utxo5(val=3)}
}
tx2 = {
  inputs: {utxo4}
  outputs: {utxo6(val=2)}
}

```

Assuming that $MST_Position(utxo_4) = 1$, $MST_Position(utxo_5) = 2$, and $MST_Position(utxo_6) = 7$, applying transactions tx_1 and tx_2 to the state MST_0 will provide the following MST_1 :

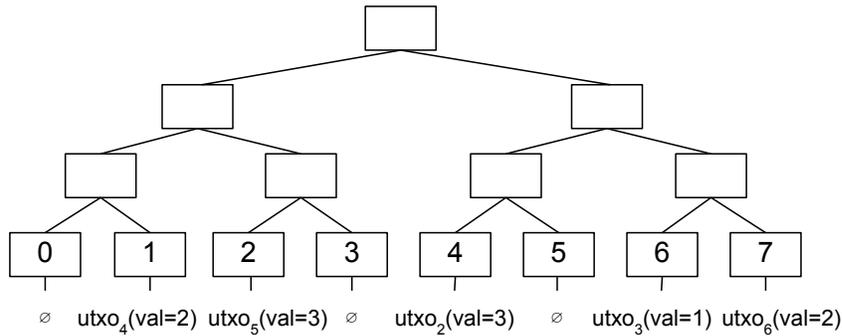


Figure 16: Merkle state tree MST_1 .

It can be seen that during the transition from MST_0 to MST_1 , the leaves 0, 1, 2, 7 have been modified. Thus, the mst_delta reflects these modifications in the bit vector:

$$mst_delta = (11100001),$$

where each bit represents whether a corresponding leaf node has been modified.

Having mst_delta in each withdrawal certificate allows to prove that some $utxo_a$ is contained in MST_k committed by the latest certificate, by providing proof of inclusion in some MST_t , $t < k$,

committed by the certificate in the past, and verifying that the bit $MST_Position(utxo_a)$ is zero for all mst_delta 's on the way from MST_t to MST_k .

This feature is of great value for circumventing data availability attacks, e.g., when a compromised sidechain (where the majority of stakeholders is adversarial) submits a withdrawal certificate to the mainchain that commits to some MST_k while not revealing to the public the MST_k tree itself. Having mst_delta in place, a user will be able to create proof of utxo ownership by using some previous MST_k . This mechanism is used for proving utxo ownership in mainchain managed withdrawals in the Latus sidechain construction ([5.5.3.2 Backward Transfer Request], [5.5.3.3 Ceased Sidechain Withdrawal]).