

Beldex

A Decentralized Ecosystem for Confidential Computation

Version 3.1.2

September 20, 2024

Legal Disclaimer

This document does not constitute or imply a prospectus of any sort. No wording contained herein should be construed as a solicitation for investment. Accordingly, this whitepaper does not pertain in any way to an offering of securities in any jurisdiction worldwide whatsoever. Rather, this whitepaper constitutes a technical description of the functionality of the Beldex Coin and the deployment of a user-friendly platform that will allow users to trade in cryptocurrencies.

The purpose of the document is to present the Beldex exchange project and details about BDX only. There should not be a redistribution of this without further written permission of the BELDEX team or this document and the information contained herein may not be sent or addressed wholly or in part directly or indirectly to any person in countries restricted. Nothing in this white paper shall be deemed to offer a document of any sort of a solicitation for investment or not in any way to offer any securities in any jurisdiction. The project plan was considered carefully before preparing this document. Beldex makes no representations and gives no warranties or guarantees of whatever nature in respect of this document including but not limited to the completeness of any information, facts or opinions contained therein. The founders, advisors, partners, directors, employees, and agents of Beldex and itself cannot be held liable for the use of and reliance on the content in this document.

The information set forth in this White Paper is subject to change. The information set forth below may not be exhaustive of all information that Beldex represents and does not imply any elements of a contractual relationship. This White Paper may be modified to provide more detailed information at any time.

Abstract

A complete confidential computation based ecosystem that anonymizes transactions, messages, and your online footprint by incentivizing network validators and creating an economic model that ensures sustainable development of the ecosystem. The confidential ecosystem consists of the Beldex blockchain, decentralized and confidential messenger, browser, Virtual Private Network (VPN), and wallet.

This whitepaper provides an overview of the Beldex ecosystem. Any future additions or changes to the ecosystem will be released in subsequent versions of this paper.

1 Why Beldex?

Confidentiality in communication and transactions was often overlooked as a luxury in the past. However, there is now a heightened awareness among individuals regarding their confidentiality as global digitization continues to expand. Blockchain technology is considered a potential solution to this challenge. Early adopters of cryptocurrency believed that blockchain transactions were inherently untraceable. However, this misconception was quickly dispelled with the rise of blockchain analysis firms that provide tools for tracing transactions. This issue extends to Bitcoin transactions as well [Nak08]. Most cryptocurrency transactions do not conceal the transaction amounts, sender or receiver addresses, or their balances. Cryptocurrencies such as Monero [AHJ17] and other confidentiality focused coins have sought to address the issue of confidentiality.

These blockchains were met with issues of scalability in proof-of-work consensus [MSBL⁺19]. For example, if there are no decentralized wallets, the user may have to rely on a centralized wallet to store their coins. An institution with sufficient information about the inputs and outputs of a wallet can perform chainalysis to find information pertaining to a transaction, also known as an EABE attack¹. Beldex proposes a unique, viable solution to the problems of confidential computation, scalability, and ecosystem.

2 Introduction

Beldex is a Proof-of-Stake (PoS) coin with masternodes as block proposers and validators. Beldex protocol core advantages are as follows:

- **Untraceable Roots** - Beldex uses RingCT network type but with a higher size of the Ring.
- **Confidential Transactions** - Transactions are completely unidentifiable and secure, ensuring user confidentiality.
- **Airdrop** - If an individual stakes a certain amount of crypto in the wallet they get a specific number of rewards in terms of the same coins.

¹<https://github.com/monero-project/monero/issues/1673#issuecomment-312968452>

2.1 Cryptonote Features

The CryptoNote protocol proposed a combination of ring signatures, ringCT (confidential transactions), and stealth addresses [vS13]. The protocol offers greater confidential computation to CryptoNote transactions [AHJ17]. When interactions occur across the Beldex independent layers, for example, the second and the first, to reduce the risk due to time-based analysis, an exchange medium that underpins the ecosystem is provided. This means that, when initiating a Flash transaction, users will retain the confidential computation guarantees offered by the Beldex chain and vice versa.

2.1.1 Untraceable Transactions

Untraceability and unlinkability are the two main features discussed in this section. The sender should not be required to trust the receiver or any other intermediary. Autonomy in transactions is a core feature of untraceable transactions in Beldex. Thus, each transactor independently produces cover traffic to mask their identity.

2.1.2 Definitions

Elliptic curve parameters

The fast scheme EdDSA is chosen as our base signature algorithm, which is developed and implemented by D.J. Bernstein et al. [vS13]. It is based on the elliptic curve discrete logarithm problem, which is similar to Bitcoin's ECDSA, so this model could also be applied to Bitcoin in future.

Common Parameters

q : a prime number; $q = 2^{255} - 19$;

d : an element of \mathbb{F}_q ; $d = -121665/121666$;

E : an elliptic curve equation; $-x^2 + y^2 = 1 + dx^2y^2$;

G : a base point; $G = (x, -4/5)$;

l : a prime order of the base point; $l = 2^{252} + 27742317777372353535851937790883648493$;

\mathcal{H}_s : a cryptographic hash function $\{0, 1\}^* \rightarrow \mathbb{F}_q$;

\mathcal{H}_p : a deterministic hash function $E(\mathbb{F}_q) \rightarrow E(\mathbb{F}_q)$.

2.1.3 Terminology

Enhanced confidential computation requires a new terminology which should not be confused with Bitcoin entities.

private ec-key is a standard elliptic curve private key: a number $a \in [1, l - 1]$;

public ec-key is a standard elliptic curve public key: a point $A = aG$;

one-time keypair is a pair of private and public ec-keys;

private user key is a pair (a, b) of two different private ec-keys;

tracking key is a pair (a, B) of private and public ec-key (where $B = bG$ and $a \neq b$);

public user key is a pair (A, B) of two public ec-keys derived from (a, b) ;

standard address is a representation of a public user key given into human friendly string with error correction;

truncated address is a representation of the second half (point B) of a public user key given into human friendly string with error correction.

The transaction structure remains similar to the structure in Bitcoin: every user can choose several independent incoming payments (transactions outputs), sign them with the corresponding private keys and send them to different destinations. Contrary to Bitcoin’s model, where a user possesses unique private and public key, in the proposed model a sender generates a one-time public key based on the recipient’s address and some random data. In this sense, an incoming transaction for the same recipient is sent to a one-time public key (not directly to a unique address) and only the recipient can recover the corresponding private part to redeem his funds (using his unique private key). The recipient can spend the funds using a ring signature, keeping his ownership and actual spending anonymous. The details of the protocol are explained in the next subsections.

2.1.4 Unlinkable payments

Classic Bitcoin addresses, once being published, become unambiguous identifier for incoming payments, linking them together and tying to the recipient’s pseudonyms. If someone wants to receive an untied transaction, he should convey his address to the sender by a private channel. If he wants to receive different transactions which cannot be proven to belong to the same owner he should generate all the different addresses and never publish them in his own pseudonym.




Figure 1: Traditional Bitcoin Keys/Transaction model; Source: CryptoNote v 2.0 [vS13]

We propose a solution allowing a user to publish a single address and receive unconditional unlinkable payments. The destination of each Beldex output (by default) is a public key, derived from recipient’s address and sender’s random data. The main advantage against Bitcoin is that every destination key is unique by default (unless the sender uses the same data for each of his transactions to the same recipient). Hence, there is no such issue as address reuse by design and no observer can determine if any transactions were sent to a specific address or link two addresses together.

First, the sender performs a Diffie-Hellman exchange to get a shared secret from his data and half of the recipient’s address. Then he computes a one-time destination key, using the shared secret and the second half of the address. Two different ec-keys are required from the recipient for these two steps, so a standard Beldex address is nearly twice as large as a Bitcoin wallet address. The receiver also performs a Diffie-Hellman exchange to recover the corresponding secret key.

A standard transaction sequence goes as follows:

1. Alice wants to send a payment to Bob, who has published his standard address. She unpacks the address and gets Bob’s public key (A, B) .
2. Alice generates a random $r \in [1, l - 1]$ and computes a one-time public key $P = \mathbb{H}_s(rA)G + B$.




Figure 2: Beldex Keys/Transaction model; Source: CryptoNote v 2.0 [vS13]




Figure 3: Standard Transaction structure; Source: CryptoNote v 2.0 [vS13]

3. Alice uses P as a destination key for the output and also packs value $R = rG$ (as a part of the Diffie-Hellman exchange) somewhere into the transaction. Note that she can create other outputs with unique public keys: different recipients' keys (A_i, B_i) imply different P_i even with the same r .
4. Alice sends the transaction.
5. Bob checks every passing transaction with his private key (a, b) , and computes $P' = \mathbb{H}_s(aR)G + B$. If Alice's transaction for with Bob as the recipient was among them, then $aR = arG = rA$ and $P' = P$.
6. Bob can recover the corresponding one-time private key: $x = \mathbb{H}_s(aR) + b$, so as $P = xG$. He can spend this output at any time by signing a transaction with x .
7. As a result Bob gets incoming payments, associated with one-time public keys which are unlinkable for a spectator. Some additional notes:
 - When Bob recognizes his transactions he practically uses only half of his private information: (a, B) . This pair, also known as the tracking key, can be passed to a third party (Carol). Bob can delegate her processing of new transactions. Bob does not need to explicitly trust Carol, because she can't recover

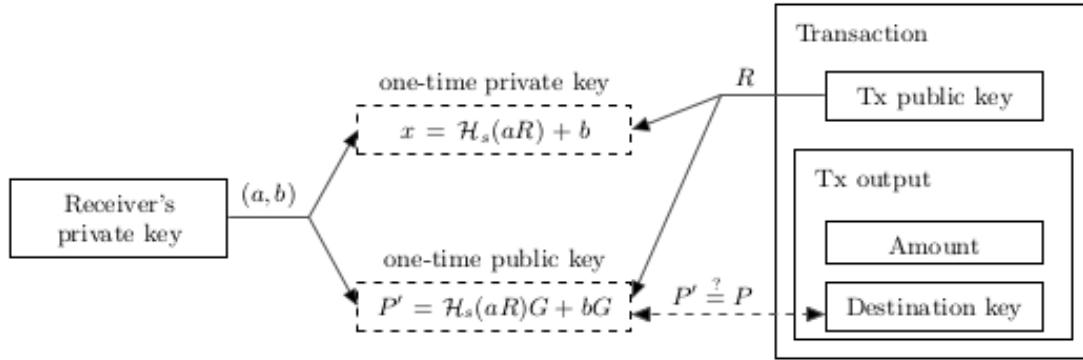


Figure 4: Incoming Transaction check; Source: CryptoNote v 2.0 [vS13]

the one-time secret key x without Bob's full private key (a, b) . This approach is useful when Bob lacks bandwidth or computation power (smartphones, hardware wallets etc.)

- In case Alice wants to prove she sent a transaction to Bob's address she can either disclose r or use any kind of zero-knowledge protocol to prove she knows r (for example by signing the transaction with r).
- If Bob wants to have an audit compatible address where all incoming transactions are linkable, he can either publish his tracking key or use a **truncated address**. That address represents only one public ec-key B , and the remaining part required by the protocol is derived from it as follows: $a = \mathbb{H}_s(B)$ and $A = \mathbb{H}_s(B)G$. In both cases every person is able to recognize all of Bob's incoming transactions, but, of course, none can spend the funds enclosed within them without the secret key b .

2.2 Ring Signatures

Ring signatures take many inputs from multiple sources and combine them with the original input from the sender. The original sender's information is concealed with a number of inputs. This masks the true history of outputs from a transaction. All Beldex transactions will employ Ring signatures and will have a fixed ring size 11 per transaction. For every unique transaction, there are ten other inputs in a ring transaction. Thus, the sender may not have to sign the transaction using their private key, but instead only the ring signature made of the private keys of ring participants.

Ring signatures are composed of a **ring** and a **signature**. Each signature is generated with a single private key and a set of unrelated public keys. Each ring is the set of public keys comprising the private key's public key, and the set of unrelated public keys. Somebody verifying a signature would not be able to tell which ring member corresponds to the private key that created it. Ring signatures were originally called **Group Signatures** because they were thought of as a way to prove a signer belongs to a group, without necessarily identifying him. In the context of Beldex they will allow for unforgeable, signer-ambiguous transactions that leave currency flows largely untraceable.

Ring signature schemes display a number of properties that will be useful for producing confidential transactions:

- **Signer Ambiguity:** An observer should be able to determine the signer must be a member of the ring (except with negligible probability), but not which member. Beldex uses this to obfuscate the origin of funds in each transaction.
- **Linkability:** If a private key is used to sign two different messages then the messages will become linked. As we will show, this property is used to prevent double-spending attacks in Beldex (except with negligible probability).
- **Unforgeability:** No attacker can forge a signature except with negligible probability. This is used to prevent theft of Beldex funds by those not in possession of the appropriate private keys.

2.2.1 One-time ring signatures

Ring signatures are composed of a *ring* and a *signature*. Each *ring* is a set of public keys, one of which belongs to the signer and the rest of which are unrelated. The *signature* is generated with that ring of keys, and anyone verifying it would not be able to tell which ring member was the actual signer.

A protocol based on one-time ring signatures allows users to achieve unconditional unlinkability. Unfortunately, ordinary types of cryptographic signatures permit to trace transactions to their respective senders and receivers. Our solution to this deficiency lies in using a different signature type than those currently used in electronic cash systems.

We will first provide a general description of our algorithm with no explicit reference to electronic cash. A one-time ring signature contains four algorithms: (GEN, SIG, VER, LNK):

- **GEN:** Takes public parameters and outputs an ec-pair (P, x) and a public key I .
- **SIG:** Takes a message m , a set S' of public keys $\{P_i\}_{i \neq s}$, a pair (P_s, x_s) and outputs a signature σ and a set $S = S' \cup \{P_s\}$.
- **VER:** Takes a message m , a set S , a signature σ and outputs *true* or *false*.
- **LNK:** takes a set $I = \{I_i\}_i$, a signature σ and outputs *linked* or *indep*.

The idea behind the protocol is fairly simple: a user produces a signature which can be checked by a set of public keys rather than a unique public key. The identity of the signer is indistinguishable from the other users whose public keys are in the set until the owner produces a second signature using the same keypair.

- **GEN:** The signer picks a random secret key $x \in [1, l - 1]$ and computes the corresponding public key $P = xG$. Additionally he computes another public key $I = x\mathcal{H}_p(P)$ which we will call the *key image*.




Figure 5: Ring Signature Anonymity; Source: CryptoNote v 2.0 [vS13]

- **SIG:** The signer generates a one-time ring signature with a non-interactive zero-knowledge proof using the techniques from [CDS94]. He selects a random subset S' of n from the other users' public keys P_i , his own keypair (x, P) and key image I . Let $0 \leq s \leq n$ be signer's secret index in S (so that his public key is P_s). He picks a random $\{q_i | i = 0, \dots, n\}$ and $\{w_i | i = 0, \dots, n, i \neq s\}$ from $(1, \dots, l)$ and applies the following transformations:

$$L_i = \begin{cases} q_i G & \text{if } i = s \\ q_i G + w_i P_i & \text{if } i \neq s \end{cases}$$

$$R_i = \begin{cases} q_i \mathcal{H}_p(P_i) & \text{if } i = s \\ q_i \mathcal{H}_p(P_i) + w_i I & \text{if } i \neq s \end{cases}$$

The next step is getting the non-interactive *challenge*:

$$c = \mathcal{H}_s(m, L_i, \dots, L_n, R_1, \dots, R_n)$$

Finally the signer computes the *response*:

$$c_i = \begin{cases} w_i & \text{if } i \neq s \\ c - \sum_{i=0}^n c_i \text{ mod } l & \text{if } i = s \end{cases}$$

$$r_i = \begin{cases} q_i & \text{if } i \neq s \\ q_s - c_s x \text{ mod } l & \text{if } i = s \end{cases}$$

The resulting signature is $\sigma = (I, c_1, \dots, c_n, r_1, \dots, r_n)$

- **VER:** The verifier checks the signature by applying the inverse transformations:

$$\begin{cases} L'_i = r_i G + c_i P_i \\ R'_i = r_i \mathcal{H}_p(P) + c_i I \end{cases}$$

Finally, the verifier checks if

$$\sum_{i=0}^n c_i \stackrel{?}{=} \mathcal{H}_s(m, L'_1, \dots, L'_n, R'_1, \dots, R'_n) \bmod l$$

If this equality is correct, the verifier runs the algorithm **LNK**. Otherwise the verifier rejects the signature.

- **LNK**: The verifier checks if I has been used in past signatures (these values are stored in the set I). Multiple uses imply that two signatures were produced under the same secret key.

The meaning of the protocol: by L -transformations the signer proves that he knows such x that at least one $P_i = xG$. To make this proof non-repeatable we introduce the key image as $I = x\mathcal{H}_p(P)$. The signer uses the same coefficients (r_i, c_i) to prove almost the same statement: he knows such x that at least one $\mathcal{H}_p(P_i) = Ix^{-1}$.

If the mapping $x \rightarrow I$ is an injection:

1. Nobody can recover the public key from the key image and identify the signer.
2. The signer cannot make two signatures with different I 's and the same x .

2.2.2 Linkable Spontaneous Anonymous Group (LSAG) signatures

Originally, group signature schemes required the system be set up, and in some cases managed, by a trusted person in order to prevent illegitimate signatures, and, in a few schemes, adjudicate disputes. Relying on a group secret is not desirable since it creates a disclosure risk that could undermine anonymity. Moreover, requiring coordination between group members (i.e. for setup and management) is not scalable beyond small groups or within companies.

Liu et al. presented a more interesting scheme [LWW04] in building on the work of Rivest et. al. in [RST01]. The authors detailed a group signature algorithm characterized by three properties: *anonymity*, *linkability*, and *spontaneity*. In other words, the owner of a private key could produce one anonymous signature by selecting any set of co-signers from a list of candidate public keys, without needing to collaborate with anyone.

Ring signature schemes are very Schnorr-like. In fact, our signature scheme can be considered a one-key ring signature. We get to two keys by, instead of defining r right away, generating a decoy r' and creating a new challenge to define r with.

Let \mathbf{m} be the message to sign, $\mathcal{R} = \{K_1, K_2, \dots, K_n\}$ a set of distinct public keys (a group/ring), and k_π the signer's private key corresponding to his public key $K_\pi \in \mathcal{R}$, where π is a secret index.

- **Signature**

1. Generate random number $\alpha \in_R \mathbb{Z}_l$ and fake responses $r_i \in_R \mathbb{Z}_l$ ($i = 1, 2, \dots, n$) excluding $i = \pi$.
2. Calculate

$$c_{\pi+1} = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [\alpha G]) \tag{1}$$

3. For $i = \pi + 1, \pi + 2, \dots, n, 1, 2, \dots, \pi - 1$ calculate, replacing $n + 1 \rightarrow 1$,

$$c_{i+1} = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [r_i G + c_i K_i]) \quad (2)$$

4. Define the real response r_π such that $\alpha = r_\pi + c_\pi k_\pi \pmod{l}$

The ring signature contains the signature $\sigma(\mathbf{m}) = (c_1, r_1, r_2, \dots, r_n)$ and the ring \mathcal{R} .

• **Verification**

1. For $i = 1, 2, \dots, n$ iteratively compute, replacing $n + 1 \rightarrow 1$

$$c'_{i+1} = \mathcal{H}_n(\mathcal{R}, \tilde{K}, \mathbf{m}, [r_i G + c_i K_i]) \quad (3)$$

2. If $c'_1 = c_1$, then the signature is valid. Note that c'_1 is the last term calculated.

In this scheme we store $(1 + n)$ integers, have one EC key image, and use n public keys.

• **Why it works?** We can informally convince ourselves the algorithm works by going through an example. Consider ring $R = K_1, K_2, K_3$ with $k_\pi = k_2$. First the signature:

- Generate the random numbers $\alpha, r_1, r_2 \in \mathbb{Z}_l$.
- Calculate

$$c_3 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [\alpha G])$$

- For $i = 1, 2$

$$c_1 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [r_3 G + c_3 K_3])$$

$$c_2 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [r_1 G + c_1 K_1])$$

- close the loop by responding with $r_2 = \alpha - k_2 c_2 \pmod{l}$

We can substitute α into c_3 to see where the word *ring* comes from:

$$c_3 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [(r_2 + k_2 c_2) G])$$

$$c_3 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [r_2 G + c_2 K_2])$$

Then verification using \mathcal{R} , and $\sigma(\mathbf{m}) = (c_1, r_1, r_2, r_3)$:

- We use r_1 and c_1 to compute

$$c'_2 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [r_1 G + K_1 G])$$

- From when we made the signature, we see $c'_2 = c_2$. With r_2 and c'_2 we compute

$$c'_3 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [(r_2 G + K_2 c'_2)])$$

- We can easily see that $c'_3 = c_3$ by substituting c_2 for c'_2 . Using r_3 and c'_3 we get

$$c'_1 = \mathcal{H}_n(\mathcal{R}, \mathbf{m}, [(r_3 G + K_3 c'_3)])$$

$c'_1 = c_1$ if we substitute c_3 for c'_3 .

2.2.3 Back's Linkable Spontaneous Anonymous Group (bLSAG)

In the LSAG signature scheme [LWW04], the owner of a private key could produce one anonymous unlinked signature per ring.⁸ In this section we present an enhanced version of the LSAG algorithm where linkability is independent of the ring's decoy members.

Let \mathbf{m} be the message to sign, $\mathcal{R} = \{K_1, K_2, \dots, K_n\}$ a set of distinct public keys (a group/ring), and k_π the signer's private key corresponding to his public key $K_\pi \in \mathcal{R}$ existence of two hash functions, \mathcal{H}_n and \mathcal{H}_p , mapping to integers from 1 to l , and curve points EC, respectively.

• Signature

1. Compute key image $\tilde{K} = k_\pi \mathcal{H}_p(K_\pi)$
2. Generate random number $\alpha \in_R \mathbb{Z}_l$ and fake responses $r_i \in_R \mathbb{Z}_l$ ($i = 1, 2, \dots, n$) excluding $i = \pi$.
3. Calculate

$$c_{\pi+1} = \mathcal{H}_n(\mathbf{m}, [\alpha G], [\alpha \mathcal{H}_p(\mathcal{R})]) \quad (4)$$

4. For $i = \pi + 1, \pi + 2, \dots, n, 1, 2, \dots, \pi - 1$ calculate, replacing $n + 1 \rightarrow 1$,

$$c_{i+1} = \mathcal{H}_n(\mathbf{m}, [r_i G + c_i K_i], [r_i \mathcal{H}_p(K_i)] + c_i \tilde{K}) \quad (5)$$

5. Define the real response r_π such that $\alpha = r_\pi + c_\pi k_\pi \pmod{l}$

The ring signature contains the signature $\sigma(\mathbf{m}) = (c_1, r_1, r_2, \dots, r_n)$, the key image \tilde{K} and the ring \mathcal{R} .

• Verification

1. Check $l\tilde{K} \stackrel{?}{=} 0$.
2. For $i = 1, 2, \dots, n$ iteratively compute, replacing $n + 1 \rightarrow 1$

$$c'_{i+1} = \mathcal{H}_n(\mathbf{m}, [r_i G + c_i K_i], [r_i \mathcal{H}_p(K_i)] + c_i \tilde{K}) \quad (6)$$

3. If $c'_1 = c_1$, then the signature is valid. Note that c'_1 is the last term calculated.

In this scheme we store $(1 + n)$ integers, have one EC key image, and use n public keys.

We must check $l\tilde{K} \stackrel{?}{=} 0$ because it is possible to add an EC point from the subgroup of size h (the cofactor) to \tilde{K} and, if all c_i are multiples of h (which we could achieve with automated trial and error using different α and r_i values), make h unlinked valid signatures using the same ring and signing key. This is because an EC point multiplied by its subgroup's order is zero.

To be clear, given some point K in the subgroup of order l , some point K^h with order h , and an integer c divisible by h :

$$\begin{aligned} c * (K + K^h) &= cK + cK^h \\ &= ck + 0 \end{aligned}$$

- **Linkability** Given two valid signatures that are different in some way (e.g. different fake responses, different messages, different overall ring members),

$$\begin{aligned}\sigma(\mathbf{m}) &= (c_1, r_1, \dots, r_n) \text{ with } \tilde{K}, \text{ and} \\ \sigma'(\mathbf{m}') &= (c'_1, r'_1, \dots, r'_n) \text{ with } \tilde{K}'\end{aligned}$$

if $\tilde{K} = \tilde{K}'$ then clearly both signatures come from the same private key.

While an observer could link σ and σ' , he wouldn't necessarily know which K_i in \mathcal{R} or \mathcal{R}' was the culprit unless there was only one common key between them. If there was more than one common ring member, his only recourse would be solving the DLP or auditing the rings in some way (such as learning all k_i with $i \neq \pi$, or learning k_π).

2.2.4 Multilayer Linkable Spontaneous Anonymous Group (ML-SAG) signatures

In order to sign transactions, one has to sign with multiple private keys. In [NMRL16], Shen Noether et al. describe a multi-layered generalization of the bLSAG signature scheme applicable when we have a set of $n \cdot m$ keys; that is, the set

$$\mathcal{R} = \{K_{i,j}\} \text{ for } i \in \{1, 2, \dots, n\} \text{ and } j \in \{1, 2, \dots, m\}$$

where we know the m private keys $\{k_{\pi,j}\}$ corresponding to the subset $\{K_{\pi,j}\}$ for some index $i = \pi$. Such an algorithm would address our needs if we generalize the notion of linkability.

- **Linkability** If any private key $k_{\pi,j}$ is used in 2 different signatures, then those signatures will be automatically linked.
- **Signature:**

1. Calculate key images $\tilde{K}_j = k_{\pi,j} \mathcal{H}_p(K_\pi, j)$
2. Generate random number $\alpha_j \in_R \mathbb{Z}_l$ and fake responses $r_{i,j} \in_R \mathbb{Z}_l$ for $i \in \{1, 2, \dots, n\}$ (excluding $i = \pi$) and $j = \{1, 2, \dots, m\}$.
3. Calculate

$$c_{\pi+1} = \mathcal{H}_n(\mathbf{m}, [\alpha_1 G], [\alpha_1 \mathcal{H}_p(K_{\pi,1})], \dots, [\alpha_m \mathcal{H}_p(K_{\pi,m})]) \quad (7)$$

4. For $i = \pi + 1, \pi + 2, \dots, n, 1, 2, \dots, \pi - 1$ calculate, replacing $n + 1 \rightarrow 1$,

$$\begin{aligned}c_{i+1} &= \mathcal{H}_n(\mathbf{m}, [r_{i,1} G + c_i K_{i,1}], [r_{i,1} \mathcal{H}_p(K_{i,1}) + c_i \tilde{K}_1], \dots, [r_{i,m} G + c_i K_{i,m}], \\ &\quad [r_{i,m} \mathcal{H}_p(K_{i,m}) + c_i \tilde{K}_m])\end{aligned}$$

5. Define the real response r_π such that $r_{\pi,j} = \alpha_j - c_\pi k_{\pi,j} \pmod{l}$

The ring signature contains the signature $\sigma(\mathbf{m}) = (c_1, r_{1,1}, \dots, r_{1,m}, \dots, r_{n,1}, \dots, r_{n,m})$, the key images $\tilde{K}_1, \dots, \tilde{K}_m$ and the ring \mathcal{R} .

- **Verification:**

1. For $j \in \{1, 2, \dots, m\}$, Check $l\tilde{K}_j \stackrel{?}{=} 0$.
2. For $i = 1, 2, \dots, n$ iteratively compute, replacing $n + 1 \rightarrow 1$

$$c'_{i+1} = \mathcal{H}_n(\mathbf{m}, [r_{i,1}G + c_i K_{i,1}], [r_{i,1}\mathcal{H}_p(K_{i,1}) + c_i \tilde{K}_1], \dots, [r_{i,m}G + c_i K_{i,m}], [r_{i,m}\mathcal{H}_p(K_{i,m}) + c_i \tilde{K}_m])$$

3. If $c'_1 = c_1$, then the signature is valid. Note that c'_1 is the last term calculated.

- **Why it works ?** Just as with the original LSAG algorithm, we can readily observe that

1. If $i \neq \pi$, then clearly the values c'_{i+1} are calculated as described in the signature algorithm.
2. If $i = \pi$ then, since $r_{\pi,j} = \alpha_j - c_\pi k_{\pi,j}$ closes the loop

$$r_{\pi,j}G + c_\pi k_{\pi,j}G = (\alpha_j - c_\pi k_{\pi,j})G + c_\pi k_{\pi,j}G = \alpha_j G$$

and

$$r_{\pi,j}\mathcal{H}_p(K_{\pi,m}) + c_\pi \tilde{K}_j = (\alpha_j - c_\pi k_{\pi,j})\mathcal{H}_p(K_{\pi,j}) + c_\pi \tilde{K}_j = \alpha_j \mathcal{H}_p(K_{\pi,j})$$

In other words, it holds also that $c'_{\pi+1} = c_{\pi+1}$.

- **Linkability:** If a private key $k_{\pi,j}$ is re-used to make any signature, the corresponding key image \tilde{K}_j supplied in the signature will reveal it. This observation matches our generalized definition of linkability².
- **Space Requirements** In this scheme we store $(1 + m * n)$ integers, have m EC key images, and use $m * n$ public keys.

2.3 Stealth Addresses

For any crypto transaction, the receiver would have to share their address, either directly to the sender, or on a public forum, thus revealing the details of their transaction. Beldex uses a form of intermediary addresses called stealth addresses to obfuscate the identity of the receiver. Whenever a sender initiates a transaction, the receiver generates a one-time intermediary address and funds are sent to this address. Funds sent to the intermediary address always reflect on the public address. The Diffie-Hellman key exchange enables the receiver to generate a private spend key for their intermediary address. So, they can claim ownership of the funds transferred to this intermediary address [DH76]. Thus, the receivers in a Beldex transaction are protected by a one-time expendable stealth address.

²As with bLSAG, linked MLSAG signatures do not indicate which public key was used to sign it. However, if the linking key image's sub-loops' rings have only one key in common, the culprit is obvious. If the culprit is identified, all other signing members of both signatures are revealed since they share the culprit's indices.

2.4 RingCT

Protecting transactional confidentiality doesn't stop with obfuscating the sender and receiver identity. To prevent any time-based analysis of transactions using the amount transferred, The paper [NMRL16], first proposed Ring Confidential Transactions (RingCT) as a means to obfuscate the amount in a transaction. RingCT uses range proofs which employ Pedersen commitments to ensure that the amount transferred can be verified without knowing the actual value, and this range is between 0 and 2^{64} . If Alice sends an amount x to Bob, she commits a value y and hashes y to the secret x such as $H(x||y)$. Once Bob receives the amount, she can then reveal her secret and Bob can verify that Alice indeed committed the correct outcome y . Thus, with Pedersen commitments, one can ensure that only non-negative values are sent and the actual value is masked. The size of crypto transactions that use traditional range proofs are large, thus, as an alternative, bulletproofs is used to greatly reduce the size of the transaction [BBB⁺18]. Therefore, to improve transaction scalability, Beldex will utilise a form of bulletproofs to reduce the payload that nodes and masternodes are required to store and transmit.

2.4.1 Pedersen Commitments

Generally speaking, a cryptographic commitment scheme is a way of committing to a value without revealing the value itself. For example, in a coin-flipping game Alice could privately commit to one outcome (i.e. 'call it') before Bob flips the coin by publishing her committed value hashed with secret data. After Bob flips the coin, Alice could declare which value she committed to and prove it by publishing her secret data. Bob could then verify her claim.

In other words, assume that Alice has secret string *blah* and the value she wants to commit to is *heads*. She could simply hash $h = \mathcal{H}(\textit{blah}, \textit{heads})$ and give h to Bob. Bob flips a coin, Alice tells Bob about *blah* and informs him she committed to *heads*, and then Bob calculates $h' = \mathcal{H}(\textit{blah}, \textit{heads})$. If $h' = h$, then he knows Alice called *heads* before the coin flip. Alice uses the so-called 'salt' *blah* so that Bob can't just guess $\mathcal{H}(\textit{heads})$ and $\mathcal{H}(\textit{tails})$ flip, and figure out she committed to *heads*.

A Pedersen commitment is a commitment that has the property of being additively homomorphic. If $C(a)$ and $C(b)$ denote the commitments for values a and b respectively, then $C(a + b) = C(a) + C(b)$. This property is useful when committing transaction amounts, as one could prove, for instance, that inputs equal outputs, without revealing the amounts at hand. Fortunately, Pedersen commitments are easy to implement with elliptic curve cryptography, as the following holds trivially

$$aG + bG = (a + b)G$$

Clearly, by defining a commitment as simply $C(a) = aG$, we could easily create cheat tables of commitments to help us recognize common amounts a .

To attain information-theoretic hiding, one needs to add a secret blinding factor and another generator H , such that it is unknown for which value of γ the following holds: $H = \gamma G$. The hardness of the discrete logarithmic problem ensures calculating γ from H is infeasible. In the case of Beldex, $H = \mathcal{H}_p(G)$. We can then define the commitment to an amount a as $C(x, a) = xG + aH$, where x is a blinding factor that prevents observers

from guessing a (for example: if you just commit $C(a = 1)$, it is trivial to guess and check). Commitment $C(x, a)$ is information-theoretically private because there are many possible combinations of x and a that would output the same C . If x is truly random, an attacker would have literally no way to figure out amount a .

2.4.2 Amount Commitments

Owning cryptocurrency is not like a bank account, where a person's balance exists as a single value in a database. Rather, a person owns a bunch of transaction outputs. Each output has an 'amount', and the sum of amounts in all unspent outputs owned is considered a person's balance. To send cryptocurrency to someone else, we create a transaction. A transaction references old outputs as inputs and addresses new outputs to recipients. Once an output has been referenced, i.e. 'spent', it can never be referenced/spent again. Since it is rare for input amounts to equal intended output amounts, most transactions include 'change', an output that sends excess back to the sender.

In Beldex, transaction amounts are hidden using a technique called RingCT, first implemented in September, 2018. While transaction verifiers don't know how much Beldex is contained in each input and output, they still need to prove the sum of input amounts equals the sum of output amounts. In other words, if we had a transaction with inputs containing amounts a_1, \dots, a_m and outputs with amounts b_1, \dots, b_p , then an observer would justifiably expect that:

$$\sum_j a_j - \sum_t b_t = 0$$

Since commitments are additive and we don't know γ , we could easily prove our inputs equal outputs to observers by making the sum of commitments to input and output amounts equal zero

$$\sum_j C_{j,in} - \sum_t b_{t,out} = 0$$

To avoid sender identifiability, Shen Noether [NMRL16] proposed verifying that commitments sum to a certain non-zero value:

$$\begin{aligned} \sum_j C_{j,in} - \sum_t b_{t,out} &= zG \\ \sum_j (x_j G + a_j H) - \sum_t (y_t G + b_t H) &= zG \\ \sum_j x_j - \sum_t y_t &= zG \end{aligned}$$

The reasons why this is useful will become clear in next chapter when we discuss the structure of transactions. Note that $C = zG$ is called a commitment to *zero* because we can make a signature with z , which proves that there is no H component to C (assuming γ is unknown). In other words $C = zG + 0H$.

2.4.3 Range proofs

One problem with additive commitments is that, if we have commitments $C(a_1)$, $C(a_2)$, $C(b_1)$, and $C(b_2)$ and we intend to use them to prove that $(a_1 + a_2) - (b_1 + b_2) = 0$, then those commitments would still apply if one value in the equation were 'negative'.

For instance, we could have $a_1 = 6$, $a_2 = 5$, $b_1 = 21$, and $b_2 = -10$.

$$(6 + 5) - (21 + -10) = 0$$

Where,

$$21G + -10G = 21G + (l - 10)G = (l + 11)G = 11G$$

The solution addressing this issue in Beldex is to prove each output amount is in a certain range using the Borromean signature scheme. This method treats each bit in the binary representation of the amount as a separate amount which we can commit to zero.

Given a commitment $C(b)$ with blinding factor for amount b , use the binary representation $(b_0, b_1, \dots, b_{k-1})$ such that

$$b = b_0 2^0 + b_1 2^1 + \dots + b_{k-1} 2^{k-1}$$

Generate random numbers $y_0, y_1, \dots, y_{k-1} \in \mathbb{Z}_l$ to be used as blinding factors. Define also Pedersen commitments for each b_i , $C_i = y_i G + b_i 2^i H$ and derive public keys $\{C_i, C_i - 2^i H\}$. Clearly one of those public keys will equal $y_i G$:

- If $b_i = 0$, then $C_i = y_i G + 0H$
- If $b_i = 1$, then $C_i = 2^i H = y_i G + 2^i H - 2^i H = y_i G$

In other words, a blinding factor y_i will always be the private key corresponding to one of the points $\{C_i, C_i - 2^i H\}$. The point that equals $y_i G$ is a commitment to zero, because either $b_i 2^i = 0$ or $b_i 2^i - 2^i = 0$. We can prove a transaction output's amount b is in the range $[0, \dots, 2 - 1]$ by signing it using the Borromean Ring Signature scheme with the ring of public keys:

$$\{\{C_0, C_0 - 2^0 H\}, \dots, \{C_{k-1}, C_{k-1} - 2^{k-1} H\}\}$$

Where, we know the private keys $\{y_0, \dots, y_{k-1}\}$ corresponding to each pair. Resulting in a signature

$$\sigma = (c_1, r_{0,1}, r_{0,2}, \dots, r_{k_1,1}, r_{k_1,2})$$

2.4.4 Range proofs in a blockchain

In the context of Beldex we will use range proofs to prove there are valid amounts in the outputs of each transaction. Transaction verifiers will have to check that the sum of each output's range proof commitments C_i equals its amount commitment C_b . For this to work we need to modify our definition of the blinding factor y_b : set $y_0, \dots, y_{k-1} \in \mathbb{Z}_l$ and define $y_b = \sum_{i=0}^{k-1} y_i$. The following equation now holds

$$\sum_{i=0}^{k-1} C_b$$

We will store only the range proof commitments/keys C_i , the output amount's commitment C_b , and the signature σ 's terms in the blockchain. The mining community can easily calculate $C_i - 2^i H$ and verify the Borromean ring signature for each output.

It will not be necessary for the receiver nor any other party to know the blinding factors y_i , as their sole purpose is proving a new output's amount is in range. Since the Borromean signature scheme requires knowledge of y_i to produce a signature, any third party who verifies one can convince himself that each loop contains a commitment to zero, so total amounts must fall within range and money is not being artificially created.

In this scheme we store $(1 + 2 * k)$ integers, use k public keys (commitments), and verify with:

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[2 * k]$.

Verifiers need to calculate $C_i - 2^i H$ and $\sum_i C_i$, so verification times include:

PA Point addition for some points $A, B : A + B[k - 1]$

PS Point subtraction for some points $A, B : A - B[k]$

KBSM Known-base scalar multiplications for some integer $a : aG[k]$

Beldex uses a special function to compute $C_i - 2^i H$ and $\sum_i C_i$ that is supposed to be faster:

VRSF Verify range proof special function $[k]$

For a total verification requirement:

VRSF Verify range proof special function $[k]$

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[2 * k]$

2.5 Beldex Transactions

2.5.1 User Keys

Unlike Bitcoin, Beldex users have two sets of private/public keys, (k^v, K^v) and (k^s, K^s) as described in public key cryptography. The address of a user is the pair of public keys (K^v, K^s) . Her private keys will be the corresponding pair (k^v, k^s) .

Using two sets of keys allows function segregation. The rationale will become clear later in this chapter, but for the moment let us call private key k^v is *the view key*, and k^s is *the spend key*. A person can use their view key to determine if an output is addressed to them, and their spend key will allow them to send that output in a transaction (and retroactively figure out it has been spent).

2.5.2 One-time addresses

To verify a transfer of money, observers need to know the money starts off being owned by the spender. How do they own it in the first place? Every Beldex user has a public address which they may distribute to other users, who can then send money to that address via transaction outputs.

The public address is never used directly. Instead, a Diffie-Hellman-like exchange is applied to it, creating a unique *one-time address* for each transaction output to be paid to the user. In this way, even external observers who know all users' public addresses will not be able to identify which user received any given transaction output. A recipient can spend their received outputs by signing a message with the one-time addresses, thereby

proving they know the private keys and must therefore own what they are spending. We will gradually flesh out this concept.

Let's start with a very simple transaction, containing exactly one input and one output a payment from Alice to Bob.

Bob has private/public keys (k_B^v, k_B^s) and (K_B^v, K_B^s) and Alice knows his public keys. A transaction could proceed as follows:

1. Alice generates a random number $r \in \mathbb{Z}_l$, and calculates the one-time public key

$$K^o = \mathcal{H}_n(rK_B^v)G + K_B^s \quad (8)$$

2. Alice sets K^o as the addressee of the payment, adds the value rG to the transaction data, and submits it to the network.
3. Bob receives the data and sees the values rG and K^o . He can calculate $k_B^v rG = rK_B^v$. He can then calculate $K_B^{s'} = K^o - \mathcal{H}_n(rk_B^v)G$. When he sees that $K_B^{s'} = K_B^s$, he knows the transaction is addressed to him.

The private key k_B^s is called the view key because anyone who has it (and Bob's public spend key) can calculate $K_B^{s'}$ for every transaction output in the blockchain (record of transactions), and 'view' which ones are addressed to Bob.

4. The one time keys for the output are:

$$\begin{aligned} K^o &= \mathcal{H}_n(rK_B^v)G + K_B^s = (\mathcal{H}_n(rK_B^v) + k_B^s)G \\ k^o &= \mathcal{H}_n(rK_B^v) + k_B^s \end{aligned}$$

While Alice can calculate the public key K^o for the one-time address, she cannot compute the corresponding private key k^o , since it would require either knowing Bob's spend key k_B^s , or solving the discrete logarithm problem for $K_B^s = k_B^s G$, which we assume to be hard. As will become clear later in this chapter, without k^o Alice can't compute the output's key image, so she can never know for sure if Bob spends the output she sent him.

A third party with Bob's view key can verify an output is addressed to Bob, yet without knowledge of the spend key this third party would not be able to spend that output nor know when it has been spent. Such a third party could be a trusted custodian, an auditor, a tax authority, etc. Somebody who could be allowed read access to the user's transaction history, without any further rights. This third party would also be able to decrypt the output's amount (to be explained in Section ??).

2.5.3 Multi-output transactions

Most transactions will contain more than one output. If nothing else, to transfer 'change' back to the sender. Beldex senders generate only one random value r . The value rG is normally known as the transaction public key and is published in the blockchain.

To ensure that all output addresses in a transaction with p outputs are different even in cases where the same addressee is used twice, Beldex uses an output index. Every output from a transaction has an index $t \in \{1, \dots, p\}$. By appending this value to

the shared secret before hashing it, one can ensure the resulting one-time addresses are unique:

$$\begin{aligned} K_t^o &= \mathcal{H}_n(rK_t^v, t)G + K_t^s = (\mathcal{H}_n(rK_t^v, t) + k_t^s)G \\ k_t^o &= \mathcal{H}_n(rK_t^v, t) + k_t^s \end{aligned}$$

2.5.4 Sub-addresses

Beldex users can generate subaddresses from each address. Funds sent to a subaddress can be viewed and spent using its main address' view and spend keys. By analogy: an online bank account may have multiple balances corresponding to credit cards and deposits, yet they are all accessible and spendable from the same point of view - the account holder. Subaddresses are convenient for receiving funds to the same place when a user doesn't want to link his activities together by publishing/using the same address. As we will see, an observer would have to solve the DLP in order to determine a given subaddress is derived from any particular address.

They are also useful for differentiating between received outputs. For example, if Alice wants to buy an apple from Bob on a Tuesday, Bob could write a receipt describing the purchase and make a subaddress for that receipt, then ask Alice to use that subaddress when she sends him the money. This way Bob can associate the money he receives with the apple he sold. We explore another way to distinguish between received outputs in the next section.

Bob generates his i^{th} subaddress ($i = 1, 2, \dots$) from his address as a pair of public keys ($K^{v,i}, K^{s,i}$):

$$\begin{aligned} K^{s,i} &= K^s + \mathcal{H}_n(k^v, i)G \\ K^{v,i} &= k^v K^{s,i} \end{aligned}$$

So,

$$\begin{aligned} K^{v,i} &= k^v(k^s + \mathcal{H}_n(k^v, i))G \\ K^{s,i} &= (k^s + \mathcal{H}_n(k^v, i))G \end{aligned}$$

Sending to a subaddress:

Let's say Alice is going to send Bob money via his subaddress ($K_B^{v,1}, K_B^{s,1}$) simple one-input, one-output transaction.

1. Alice generates a random number $r \xleftarrow{R} \mathbb{Z}_l$, and calculates the one-time public key

$$K^o = \mathcal{H}_n(rK_B^{v,1})G + K_B^{s,1}$$

2. Alice sets K^o as the addressee of the payment, adds the value $rK_B^{s,1}$ to the transaction data, submits it to the network.
3. Bob receives the values K^o and $rK_B^{s,1}$, and he can calculate $k_B^v rK_B^{s,1} = rK_B^{v,1}$. He can then calculate $K_B^{s',1} = K^o - \mathcal{H}_n(rK_B^{v,1})$. When he sees that $K_B^{s',1} = K_B^{s,1}$, he knows that the transaction is addressed to him.

Bob only needs his private view key k^v and public spend key $K_B^{s,1}$, to find the transaction outputs sent to his subaddress.

4. The one-time keys for the output are:

$$K^o = \mathcal{H}_n(rK_B^{v,1})G + K_B^{s,1} = (\mathcal{H}_n(rK_B^{v,1}) + k_B^{s,1})G$$

$$k^o = \mathcal{H}_n(rK_B^{v,1}) + k_B^{s,1}$$

Now, Alice's transaction public key $rK_B^{s,1}$ is particular to Bob (instead of rG). If she creates a p -output transaction with at least one output intended for a subaddress, Alice needs to make a unique transaction public key for each output $\in \{1, 2, \dots, p\}$. In other words, if Alice is sending to Bob's subaddress $(K_B^{v,1}, K_B^{s,1})$ and Carol's address (K_C^v, K_C^s) , she will put two transaction public keys $(r_1K_B^{s,1}, r_2G)$ in the transaction data.

2.5.5 Integrated addresses

In order to differentiate between the outputs they receive, a recipient can request senders include a payment ID in transaction data. For example, if Alice wants to buy an apple from Bob on a Tuesday, Bob could write a receipt describing the purchase and ask Alice to include the receipt's ID number when she sends him the money. This way Bob can associate the money he receives with the apple he sold.

Senders can communicate payment IDs in clear text, but manually including the IDs in transactions is inconvenient for recipients, who might inadvertently expose their activities. In Beldex, recipients can integrate payment IDs into their addresses, and provide those integrated addresses, containing $(K^v, K^s, \textit{payment ID})$, to senders. Payment IDs can technically be integrated into any kind of address, including normal addresses, subaddresses, and multisignature addresses.

Senders addressing outputs to integrated addresses can encode payment IDs using the shared secret rk_t^v and a XOR operation, which recipients can then decode with the appropriate transaction public key and another XOR procedure. Encoding payment IDs in this way allows senders to prove they made particular transaction outputs (i.e. for audits, refunds, etc.).

2.5.6 Binary operator XOR

The binary operator XOR evaluates two arguments and returns true if one, but not both, of the arguments is true. Here is its truth table:

A	B	A XOR B
F	F	F
F	T	T
T	F	T
T	T	F

Table 1: A XOR B

In the context of computer science, XOR is equivalent to bit addition modulo 2. For example, the XOR of two bit pairs:

$$XOR(\{1, 1\}, \{1, 0\}) = \{1 + 1, 1 + 0\} = \{0, 1\}$$

Examining the previous example, each of these produce the same output: $XOR(\{1, 1\}, \{1, 0\})$, $XOR(\{0, 0\}, \{0, 1\})$, $XOR(\{1, 0\}, \{1, 1\})$, or $XOR(\{0, 1\}, \{0, 0\})$. There are 2^b combinations of XOR inputs (with b bits each) for the same output, so if input $A \in_R \{1, \dots, 2\}$, an observer who learned $C = XOR(A, B)$ could not gain any information about B .

At the same time, anyone who knows two of the elements A, B, C , where $C = XOR(A, B)$, can calculate the third element, such as $A = XOR(B, C)$. XOR indicates if two elements are different or the same, so knowing C and B is enough to expose A . A careful examination of the truth table reveals this to be true.

2.5.7 Encoding

The sender encodes the payment ID for inclusion in transaction data

$$\begin{aligned} k_{mask} &= \mathcal{H}_n(rK_t^v, pid_tag) \\ k_{payment_ID} &= k_{mask} \rightarrow \text{reduced to bit length of payment ID} \\ \text{encoded payment ID} &= XOR(k_{payment_ID}, \text{payment ID}) \end{aligned}$$

The output of a cryptographic hash function \mathcal{H} is uniformly distributed across the range of possible outputs. In other words, for some input A , $\mathcal{H}(A) \in_R^D \mathcal{S}_H$. where, \mathcal{S}_H is the set of possible outputs from \mathcal{H} . We use \in_R^D to indicate the function is deterministically random. $\mathcal{H}(A)$ produces the same thing every time, but its output is equivalent to a random number.

We include `pid_tag` to ensure k_{mask} is different from the component $\mathcal{H}_n(rK_t^v, t)$ in one-time output addresses.

2.5.8 Decoding

Whichever recipient t the payment ID was created for can find it using his view key and the transaction public key rG :

$$\begin{aligned} k_{mask} &= \mathcal{H}_n(rK_t^v, pid_tag) \\ k_{payment_ID} &= k_{mask} \rightarrow \text{reduced to bit length of payment ID} \\ \text{payment ID} &= XOR(k_{payment_ID}, \text{encoded payment ID}) \end{aligned}$$

Similarly, senders can decode payment IDs they had previously encoded by recalculating the shared secret $k_{mask} = \mathcal{H}_n(rK_t^v)$.

2.5.9 Transaction types

Transaction structures, protocols, and cryptographic schemes are always prone to evolving as new objectives or threats are found. In this report we have focused our attention on Ring Confidential Transactions (RingCT). RingCT is mandatory for all new transactions, so we will not describe any deprecated transaction types, even if they are still partially supported.

The transaction types we will describe in this chapter are `RCTTypeFull` and `RCTTypeSimple`. The former category 2.5.10 closely follows the ideas explained by Shen Noether et al. in [NMRL16]. At the time that paper was written, the authors most likely intended to fully replace the original CryptoNote transaction scheme.

However, for multi-input transactions, the signature scheme formulated in that paper was thought to entail a risk on traceability. This will become clear when we supply technical details, but in short: if one spent output became identifiable, the rest of the spent outputs would also become identifiable. This would have an impact on the traceability of currency flows, not only for the transaction originator affected, but also for the rest of the blockchain. To mitigate this risk, yet different signature scheme for multi-input transactions was proposed. The transaction type `RCTTypeSimple` 2.5.17 is used in those situations. The main difference, as we will see later, is that each input is signed independently.

2.5.10 Ring Confidential Transactions of type `RCTTypeFull`

Our perception is that the decision to limit `RCTTypeFull` transactions to one input was rather hastily taken, and that it might change in the future, perhaps if the algorithm to select additional mix-in outputs is improved and ring sizes are increased. Also, Shen Noether’s original description in Ring Confidential Transactions did not envision constraints of this type. At any rate, it is not a hard constraint. An alternative wallet might choose to sign transactions using either scheme, independently of the number of inputs involved.

We have therefore chosen to describe the scheme as if it were meant for multi-input transactions. An actual example of a `RCTTypeFull` transaction, with all its components, can be inspected in Appendix A.

2.5.11 Amount commitments

Recall from Section 2.4.2 that we had defined a commitment to an output’s amount b as:

$$c(b) = yG + bH$$

In the context of Beldex, output recipients should be able to reconstruct the amount commitments. This means the blinding factor y and amount b must be communicated to the receiver.

The solution adopted in Beldex is a Diffie-Hellman shared secret rK_B^v . For any given transaction in the blockchain, each of its outputs $t \in \{1, \dots, p\}$ has two associated values called mask and amount satisfying

$$\begin{aligned} \text{mask}_t &= y_t + \mathcal{H}_n(\mathcal{H}_n(rK_B^v, t)) \\ \text{amount}_t &= b_t + \mathcal{H}_n(\mathcal{H}_n(\mathcal{H}_n(rK_B^v, t))) \end{aligned}$$

The receiver Bob will be able to calculate the blinding factor y_t and the amount b_t using the transaction public key rG and his view key K_B^v . He can also check that the commitment $C(y_t, b_t)$ provided in the transaction data, henceforth denoted, corresponds to the amount at hand. More generally, any third party with access to Bob’s view

key could decrypt his output amounts, and make sure they agree with their associated commitments.

2.5.12 Commitments to zero

Assume a transaction sender has previously received various outputs with amounts a_1, \dots, a_m addressed to one-time addresses $K_{\pi,1}^o, \dots, K_{\pi,m}^o$ and with amount commitments $C_{\pi,1}^a, \dots, C_{\pi,m}^a$. This sender knows the private keys $k_{\pi,1}^o, \dots, k_{\pi,m}^o$ corresponding to the one-time addresses (Section 2.5.2). The sender also knows the blinding factors x_j used in commitments $C_{\pi,j}^a$ (Section 2.5.11).

A transaction consists of inputs a_1, \dots, a_m and outputs b_1, \dots, b_p such that $\sum_{j=1}^m a_j - \sum_{t=1}^p b_t = 0$. The sender re-uses the commitments $C_{\pi,1}^a, \dots, C_{\pi,m}^a$ from the previous outputs, and create commitments for b_1, \dots, b_p . Let these new commitments be $C_{\pi,1}^b, \dots, C_{\pi,p}^b$.

As hinted in Section 2.4.2, the sum of the commitments will not be 0, but a curve point zG :

$$\sum_j C_{\pi,j}^a - \sum_t C_{\pi,t}^b = zG$$

The sender will know z , allowing him to create a signature on this commitment to zero. Indeed, z follows from the blinding factors if and only if input amounts equal output amounts (recalling Section 2.4.1, we don't know γ in $H = \gamma G$).

$$\begin{aligned} & \sum_{j=1}^m C_{\pi,j}^a - \sum_{t=1}^p C_{\pi,t}^b \\ &= \sum_j x_j G - \sum_t y_t G + (\sum a_j - \sum b_t) H \\ &= \sum_j x_j G - \sum_t y_t G \\ &= zG \end{aligned}$$

2.5.13 Signature

The sender selects v sets of size m , of additional unrelated addresses and their commitments from the blockchain, corresponding to apparently unspent outputs. She mixes the addresses in a ring with her own m unspent outputs' addresses, adding false commitments to zero, as follows:

$$\mathcal{R} = \left\{ \left\{ K_{1,1}^o, \dots, K_{1,m}^o, \left(\sum_j C_{1,j} - \sum_t C_t^b \right) \right\} \right. \\ \dots, \\ \left. \left\{ K_{\pi,1}^o, \dots, K_{\pi,m}^o, \left(\sum_j C_{\pi,j}^a - \sum_t C_t^b \right) \right\} \right. \\ \dots, \\ \left. \left\{ K_{v+1,1}^o, \dots, K_{v+1,m}^o, \left(\sum_j C_{v+1,j} - \sum_t C_t^b \right) \right\} \right\}$$

Looking at the structure of the key ring, we see that if

$$\sum_j C_{\pi,j}^a - \sum_t C_t^b = 0$$

Then, any observer would recognize the set of addresses $\{K_{\pi,1}^o, \dots, K_{\pi,m}^o\}$ as the ones in use as inputs, and therefore currency flows would be traceable. With this observation made we can see the utility of zG . All commitment terms in \mathcal{R} return some EC point, and the π^{th} such term is zG . This allows us to create an MLSAG signature (Section 2.2.4) on \mathcal{R} .

MLSAG signature for inputs:

The private keys for $\left\{ K_{\pi,1}^o, \dots, K_{\pi,m}^o, \left(\sum_j C_{\pi,j}^a - \sum_t C_t^b \right) \right\}$ are $k_{\pi,1}^o, \dots, k_{\pi,m}^o, z$ which are known to the sender. MLSAG in this scenario does not use a key image for the commitment to zero zG . This means building and verifying the signature excludes the term $r_{i,m+1} \mathcal{H}_p(K_{i,m+1}) + c_i \tilde{K}_z$. The message \mathbf{m} signed in the input MLSAG is essentially a hash of all transaction information except for the MLSAG signature itself.

This ensures transactions are tamper-proof from the perspective of both transaction authors and verifiers. k^o is the essence of Beldex's transaction model. Signing \mathbf{m} with k^o proves you are the a owner/recipient of the amount committed to in C^a . Verifiers can be confident that transaction authors are spending their own funds.

Range proofs for outputs: To avoid the amount ambiguity of outputs described in Section 2.4.4, the sender must also employ the Borromean signature scheme to sign amount ranges for each output $t \in \{1, \dots, p\}$. No message \mathbf{m} is signed by the Borromean signatures.

Range proofs are not needed for input amounts because they are either expressed clearly (as with transaction fees and block rewards), or were proven in range when first created as outputs. In the current version, each amount is expressed as a fixed point number of 64 bits. This means the data for each range proof will contain 64 bit commitments and $2 * 64 + 1$ signature terms.

2.5.14 Transaction Fee

Typically transaction outputs are lower in total than transaction inputs, in order to provide a fee that will incentivize miners to include the transaction in the blockchain. Transaction fee amounts are stored in clear text in the transaction data transmitted to the network. Miners can create an additional output for themselves with the fee. This fee amount must be converted into a commitment so verifiers can confirm transactions sum to zero.

The solution is to calculate the commitment of the fee f without the masking effect of any blinding factor. That is, $C(f) = fH$, where f is communicated in clear text.

The network verifies the MLSAG signature on \mathcal{R} by including fH as follows:

$$\left(\sum_j C_{i,j} - \sum_t C_t^b \right) - fH$$

This works because this is a commitment to zero:

$$\left(\sum_j C_{i,j} - \sum_t C_t^b \right) - fH = zG$$

2.5.15 Avoiding double-spending

An MLSAG signature (Section 2.2.4) contains images \tilde{K}_j of private keys $k_{\pi,j}$. An important property in any cryptographic signature scheme is that it should be unforgeable with non-negligible probability. Therefore, to all practical effects, we can assume a signature's key images must have been deterministically produced from legitimate private keys.

The network need only verify that key images included in MLSAG signatures (corresponding to inputs) and calculated as $\tilde{K}_j^o = k_{\pi,j}^o \mathcal{H}_n(K_{\pi,j}^o)$ have not appeared before in other transactions. If they have, then we can be sure we are witnessing an attempt to re-spend an output $C_{\pi,j}^a$ addressed to $K_{\pi,j}^o$.

If someone tries to spend $C_{\pi,j}^a$ twice, they will reveal the index π for both transactions where it appears. This has two effects: (1) all outputs at index π in the first transaction are revealed as its real inputs, and (2) all outputs at index π in the second transaction are revealed as not having been spent before. The second is a problem even considering miners would reject the double-spend transaction.

These effects could weaken the network benefits of ring signatures, and are part of the reason **RCTTypeFull** is only used for single-input transactions. The other main reason is that a cryptanalyst would know that, in general, all real inputs share an index.

2.5.16 Space and verification requirements

MLSAG signature (inputs)

From Section 2.2.4 we recall that an MLSAG signature in this context would be expressed as

$$\sigma(\mathbf{m}) = (c_1, r_{1,1}, \dots, r_{1,m+1}, \dots, r_{v+1,1}, \dots, r_{v+1,m+1}) \text{ with } (\tilde{K}_1^o, \dots, \tilde{K}_m^o)$$

As a legacy of CryptoNote, the values \tilde{K}_j^o are not referred to as part of the signature, but rather as images of the private keys $k_{\pi,j}^o$. These key images are normally stored separately in the transaction structure, as they are used to detect double-spending attacks.

With this in mind and assuming point compression, an MLSAG signature will require $((v + 1) \cdot (m + 1) + 1) \cdot 32$ bytes of storage, where v is the mixin level and m is the number of inputs. In other words, a transaction with 1 input and a total ring size of 32 would consume $(32 \cdot 2 + 1) \cdot 32 = 2080$ bytes. To this value we would add 32 bytes to store the key image of each input, for $m \cdot 32$ bytes of storage, and additional space to store the ring member offsets in the blockchain. These offsets are used by verifiers to find each MLSAG signature's ring members' output keys and commitments in the blockchain, and are stored as variable length integers, hence we cannot exactly quantify the space needed.

Including the computation of $\left(\sum_j C_{i,j} - \sum_t C_t^b\right) - fH$, and verifying key images are in G 's subgroup with $l\tilde{K} \stackrel{?}{=} 0$, we verify a `RCTTypeFull` transaction's MLSAG (Section 2.2.4) with:

PA Point addition for some points $A, B : A + B[m * (v + 1)]$

PS Point subtraction for some points $A, B : A - B[(v + 1) * (p + 1)]$

VBSM Variable-base scalar multiplications for some integer a , and point $P : aP[m]$

KBSM Known-base scalar multiplications for some integer $a : aG[1]$

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[(m + 1) * (v + 1)]$

VBA Variable-base additions for some integers a, b , and points $A, B : aA + bB[m * (v + 1)]$

Range proofs (outputs) We know that, the Borromean signature for range proofs takes the form of an n -tuple

$$\sigma = (c_1, r_{0,1}, r_{0,2}, \dots, r_{63,1}, r_{63,2})$$

Ring keys are considered part of ring signatures. However, in this case it is only necessary to store the commitments C_i , as the ring key counterparts $C_i - 2^i H$ can be easily derived (for verification purposes). Respecting this convention, a range proof will require $(1 + 64 \cdot 2 + 64) * 32 = 6176$ bytes per output.

Including $\sum_i C_i - 2^i H$ and checking $\sum_i C_i = C$, range proofs for p outputs will require this to verify.

VRSF Verify range proof special function $[p * 64]$

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[p * 2 * 64]$.

2.5.17 Ring Confidential Transactions of type `RCTTypeSimple`

In the current Beldex code base, transactions with more than one input are signed using a scheme known as **RCTTypeSimple**. The main characteristic of this approach is that instead of signing the entire set of inputs at once, the sender signs each input separately.

Among other things, this means we can't use commitments to zero in the same way as for `RCTTypeFull` transactions. It's really unlikely that we could match an input amount to an output amount, and in most cases each individual input will be less than the sum of outputs. So, we cannot directly commit $inputs - outputs = 0$.

In more detail, assume that Alice wants to sign input j . Imagine for a moment we could sign an expression like this

$$C_j^a - \sum_t C_t^b = \left(x_j - \sum_t y_t\right)G + \left(a_j - \sum_t b_t\right)H$$

Since $a_j - \sum_t b_t \neq 0$, Alice would have to solve the DLP for $H = \gamma G$ in order to obtain the private key of the expression, something we have assumed to be computationally difficult.

2.5.18 Amount commitments

As explained, if inputs are decoupled from each other than the sender can't sign an aggregate commitment to zero. On the other hand, signing each input individually implies an intermediate approach. The sender could create new commitments to the input amounts and commit to zero with respect to each of the previous outputs being spent. In this way, the sender could prove the transaction takes as input only the outputs of previous transactions.

In other words, assume the amounts being spent are a_1, \dots, a_m . These amounts were outputs in previous transactions, in which they had commitments

$$C_j^a = x_j G + a_j H$$

The sender can create new commitments to the same amounts but using different blinding factors; that is,

$$C_j'^a = x_j' G + a_j H$$

Clearly, she would know the private key of the difference between the two commitments:

$$C_j^a - C_j'^a = (x_j - x_j')G$$

Hence, she would be able to use this value as a commitment to zero for each input. Let us say $(x_j - x_j') = z$, and call each $C_j'^a$ a pseudo output commitment.

Similarly to **RCTTypeFull** transactions, the sender can include each output's encoded blinding factor (mask) for y_t and amount for b_t in the transaction (see Section 2.4.2), which will allow each receiver t to decode y_t and b_t using the shared secret r . Before committing a transaction to the blockchain, the network will want to verify that the transaction balances. In the case of RCTTypeFull transactions, this was simple, as the MLSAG signature scheme implies each sender has signed with the private key of a commitment to zero.

For **RCTTypeSimple** transactions, blinding factors for input and output commitments are selected such that

$$\sum_j x_j' - \sum_t y_t = 0$$

This allows us to prove input amounts equal output amounts:

$$\sum_j C_j'^a - \sum_t C_t^b - fH = 0$$

Fortunately, choosing such blinding factors is simple. In the current version of Beldex, all blinding factors are random except z_j , which is simply set to

$$x'_m = \sum_t y_t - \sum_{j=1}^{m-1} x'_j$$

2.5.19 Signature

As we mentioned, in transactions of type **RCTTypeSimple** each input is signed individually. We use the same MLSAG signature scheme as for **RCTTypeFull** transactions, except with different signing keys.

Assume that Alice is signing input j . This input spends a previous output with key $\tilde{K}_{\pi,j}^o$ that had commitment C_j^a . Let $C_j'^a$ be a new commitment for the same amount but with a different blinding factor. Similar to the previous scheme, the sender selects v unrelated outputs and their respective commitments from the blockchain to mix with the real, j^{th} , input

$$\begin{aligned} &K_{1,j}^o, \dots, K_{\pi-1,j}^o, K_{\pi+1,j}^o, \dots, K_{v+1,j}^o \\ &C_{1,j}, \dots, C_{\pi-1,j}, C_{\pi+1,j}, \dots, C_{v+1,j} \end{aligned}$$

She can sign input j using the following ring:

$$\begin{aligned} \mathcal{R}_j = & \left\{ \left\{ K_{1,j}^o, \left(\sum_j C_{1,j}^a - C_{\pi,j}^a \right) \right\} \right. \\ & \dots, \\ & \left\{ K_{\pi,j}^o, \left(\sum_j C_{\pi,j}^a - C_{\pi,j}^a \right) \right\} \\ & \dots, \\ & \left. \left\{ K_{v+1,j}^o, \left(\sum_j C_{v+1,j}^a - C_{\pi,j}^a \right) \right\} \right\} \end{aligned}$$

Alice will know the private keys $k_{\pi,j}^o$ for $K_{\pi,j}^o$, and z_j for the commitment to zero ($C_{\pi,j}^a - C_{\pi,j}^a$). Recalling Section 2.5.13, there is no key image for the commitments to zero $z_j G$, and consequently no corresponding key image component in each input's signature's construction.

Each input in an **RCTTypeSimple** transaction is signed individually, applying the scheme described in Section 2.5.13, but using rings like \mathcal{R}_j as defined above. The advantage of signing inputs individually is that the set of real inputs and commitments to zero need not be placed at the same index π , as they are in the aggregated case. This means even if one input's origin became identifiable; the other inputs' origins would not. The message m signed by each input is essentially the same as for **RCTTypeFull** transactions, except it includes pseudo output commitments for the inputs. Only one message is produced, and each input MLSAG signs it.

2.5.20 Space and verification requirements

MLSAG signature (inputs) Each ring contains $(v + 1).2$ keys. Using the point compression technique, an input signature σ will require $(2(v + 1) + 1).32$ bytes. On top of this is, the key image $\tilde{K}_{\pi,j}^o$ and the pseudo output $C_{\pi,j}^a$ commitment leave a total of $(2(v + 1) + 3).32$ bytes per input. A transaction with 20 inputs using rings with 32 total members will need $((32.2 + 3).32)20 = 42880$ bytes.

For the sake of comparison, if we were to apply the RCTTypeFull scheme to the same transaction, then the MLSAG signature and key images would require $(32.21 + 1).32 + 20.32 = 22176$ bytes.

Including the computation of $(C_{i,j} - C'_{\pi,j})$ and $(\sum_j C_j^a \stackrel{?}{=} \sum_t C_t^b + fH)$, and verifying key images are in G 's subgroup with $l\tilde{K}$, we verify all of a RCTTypeSimple transaction's MLSAGs with:

PA Point addition for some points $A, B : A + B[m + p + 1]$

PS Point subtraction for some points $A, B : A - B[m * (v + 1)]$

VBSM Variable-base scalar multiplications for some integer a , and point $P : aP[m]$

KBSM Known-base scalar multiplications for some integer $a : aG$

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[m * 2 * (v + 1)]$

VBA Variable-base additions for some integers a, b , and points $A, B : aA + bB[m * (v + 1)]$.

Range proofs (outputs)

The size of range proofs remains the same for RCTTypeSimple transactions. As we calculated for **RCTTypeFull** transactions, each output will require 6176 bytes of storage. As before, each range proof will require this to verify (Section 2.4.4):

VRSF Verify range proof special function $[p * 64]$

DVBA Double-variable-base addition for some integers a, b , and point $B : aG + bB[p * 2 * 64]$.

3 Beldex Blockchain

3.1 Masternodes

Beldex improvises on the CryptoNote protocol to provide greater confidential computation; however, a group of incentivized masternodes determine most of the functionality and scalability in Beldex. These masternodes are full nodes that store the entire copy of the Beldex blockchain. To set up a Beldex masternode, a node operator must lock-in a minimum collateral and possess the required minimum bandwidth and storage capabilities. The collateral ensures that the masternode operators refrain from acting dishonestly, as they have a significant stake in the network. As compensation for their services, masternode operators receive a portion of the block reward. A masternode is continually reviewed and if it is found to act dishonestly or doesn't provide the required bandwidth at any time, then it is penalized and moved to the end of the reward queue. A maximum of three warnings is given to a masternode operator before they are removed from the network.

The network thus formed is resistant to Sybil attacks, in which a single or a group of

bad actors work in coalition to undermine the network. The initial setup collateral and the periodical review of masternodes ensure that such attacks are prevented. The maximum amount of stake a single validator can have is limited by supply and demand, preventing one validator to have a disproportionately large stake so as to sabotage the second-layer ecosystem services. Resistance to Sybil attacks can also be acquired through Cryptoeconomics [QC19]. When a bad actor begins to accumulate Beldex (BDX) in order to obtain a majority stake within the network, the supply in circulation decreases and in turn drives the price up. Acquiring additional Beldex becomes increasingly expensive, thus making the attack economically unfeasible. This disincentivizes the attacker. To ensure that masternode validators obtain sustainable returns for their services, the circulating supply is actively reduced and the emission curves and minimum collateral requirements are designed such that enough circulating supply is locked within the masternodes.

3.2 Block Reward

Block rewards in the Beldex network are currently distributed through proof-of-Stake consensus. A node or a masternode creates a new block every 30 seconds by collecting the transactions, validating them, including them within the block, and adding the block to the end of the network. To perform this action, they maintain a certain bandwidth and propose a new block periodically. In return, they collect block rewards as new BDX coins are emitted. The Beldex block reward is set to a constant value of 6.5 BDX.

Node Reward: In addition to the transaction fees, 10% of the block reward is awarded to the node that constructs the block.

Master Node Reward: 90% of total Beldex rewards go to a Masternode, or two Masternodes if both the masternodes construct the block at the same time. Once the Masternode receives a reward, it assumes the last position in the reward queue. Thus, nodes that recently received rewards are moved to the end of the reward queue while nodes that have been waiting for a longer period are moved to the start of the queue. The very first time a Masternode validator is indexed on the network's distributed decentralized directory, it assumes the last position in the reward queue. The masternodes that are found to provide optimum services to the network continually move up the queue, if they are not otherwise penalized.

3.3 Authenticated Collateralization

A masternode operator may stake this amount entirely by themselves or stake part of the amount and let other users stake the remaining. In this case, they share the rewards based on their stake within the masternode. To be added to the Beldex distributed directory, each Masternode must prove to the network that they hold the necessary stake, which is 10000 BDX. The innate confidentiality features of Beldex prevent the masternodes from doing so as public addresses do not reveal the amount of transaction. They also cannot use viewkeys to see transactions going out of the network.

To solve this dilemma of Masternodes, Beldex uses a novel method of time-locking the collateral. The initial collateral by a masternode is locked for a specific block-height of 86,400 blocks. Thus, from the time of registration until 86,400 blocks are created, the masternodes will be unable to spend this collateral. Beldex leverages this method to

ensure that a particular masternode holds the required collateral amount.

In the additional field of transaction, the Masternode operator provides the Beldex address to receive rewards. This address is the public key of the Masternode operator. The transactions for the purpose of registration are open transactions, thus they may not be included in a ring transaction as they do not serve to provide anonymity. The existing Masternodes in the network must verify that the first time Masternode registrants possess the required collateral before they are added to the distributed decentralized directory.

3.4 Flash

Flash instant transactions are a second layer atop the Beldex blockchain that allows for confirmation of transactions before they are included in a block. Blockchain transactions typically take the time required for the transaction to be included in a block and confirmed by the network. A Beldex transaction requires 10 confirmations and the block creation time is 30 seconds. Thus, it may take anywhere between 5 – 10 minutes for a transaction to be confirmed. But with flash, it only takes a few seconds.

The Beldex second layer architecture (Flash) is similar to Bitcoin’s lightning network in that a shortest possible route is selected between the Flash channels [PD16]. A channel path reducing algorithm is used to send the outputs to the receiver, where channels from users who have made transactions with each other before become relays. If B has sent or received inputs from A and C , then B can act as a relay in transactions between A and C .

Flash instant transactions are feasible through the use of Masternodes that confirm a transaction’s authenticity by securing the key images associated with the UTXOs. In a ring signature transaction, the UTXOs are linked to a distinctive key image. Select Masternodes store this key image and hold onto it until the transaction is added to a block on the Beldex network. If the said key image is produced more than once, then that signals a double spending attempt. The corresponding duplicate transaction is rejected.

Similar to Ethereum, Flash will enable a mechanism for competing transactions to pay a higher fee to be completed first. But this is several times faster and will be in the order of a few seconds since it takes place outside the blockchain. Confirmation time on the blockchain is also significantly reduced due to the network transition to a Proof-of-Stake (PoS) consensus.

3.5 Coin Burning Mechanism

The Beldex coin burning mechanism is a way to curb inflation and deliver sustainable price discovery. To achieve this, the network fee obtained through Flash transactions will be burned. The Flash second layer serves to provide instant transactions and thus will have a higher fee. This fee is burned as soon as the transaction is completed.

3.6 Proof-of-Stake (PoS) in Beldex Chain

POS is an alternative to POW that is designed to be economical and energy-efficient. At present, most cryptocurrencies run on a POW algorithm that consumes a lot of energy through mining hardware. Newer and powerful hardware is developed to plow through

the hashrate. This hardware is often not affordable to everyone, making them less decentralized and they consume several KW of power every day. Bitcoin's annual energy consumption index is an estimated 169.59 TWh, according to Digiconomist³.

A full POS Beldex network will enable users to stake BDX on nodes. Those with a stake can participate in validating blocks on the network. The higher the stake, the higher the chances for the node to be selected to validate a block, and the lower is the difficulty. POS on Beldex is based on delimited competition of nodes.

Anyone is free to contribute to the Beldex network, given the required collateral staking amount is locked in the node. Dishonest nodes are disincentivized and their collateral is devalued. This prevents the nodes from acting against the network whereas honest nodes are rewarded and appraised. Beldex also provides pool staking via masternodes where the hardware is shared between a few individuals who share the stake amount (10,000 BDX).

3.7 How Does Proof of Stake Work?

Proof of Stake works on the mechanism that for a node to participate in the network it must possess the minimum required stake. The minimum required stake to run a POS node is 10000 BDX. Bucephalus is Beldex's Proof of Stake implementation. POS makes the Beldex network fast, sustainable, secure, and scalable.

In POS consensus, every 30 seconds, a node is selected at random to produce a block which is then sent to a quorum vote consisting of 11 other participating validator nodes. The validators in a quorum verify the authenticity of the block and if any 7 validator nodes of the 11 attest to the block's authenticity, it is broadcast to the network as a valid block. For the creation of the next block, a new quorum or consensus group is formed and the process repeats.

3.8 Power Consumption In A POS Blockchain

One area of concern for blockchains is sustainability. 50 or 100 years into the future, they need to be able to validate transactions the same way they do now without consuming 50X more power. The Bitcoin blockchain consumes over 200 TWh per year, that's as much power consumed by entire nation-states.

POS solves the blockchain energy problem by decoupling the hash power from the computing power of the nodes. The minimum stake provides one the hash power required to participate in the network. By transferring authority from computing power to one's stake within the network, POS essentially reduces the entry barrier to participation. POS networks could therefore have a greater degree of decentralization, scalability, and flexibility.

3.9 Security in Proof of Stake Consensus Model

In Proof of Work, an attacker would have to gain 51% of the hash power within the network. To do this, they set up a majority of the nodes, gain computational power which translates to hash power. For this reason, blockchain networks such as Bitcoin

³<https://digiconomist.net/bitcoin-energy-consumption/>

increase the cost of hardware required to attack a network. They use ASICs, FPGAs, and GPUs which are cost-intensive. However, this leads to additional complications of centralization. If the network is minable only with a specific kind of ASIC, then the manufacturers of this equipment can play a role in who gets to participate in the network.

On the other hand, Proof of Stake provides node operators with the authority to validate a block based on their stake within the network. An attacker, in this case, would have to possess more than 50% of the staked $\$BDX$ to take over the network, which is cost-intensive and economically not feasible. The threat of a 51% attack can be negated by sufficient decentralization of the network. If more than 50% of the network's tokens are locked, then it becomes virtually impossible to take over the network. Moreover, any malicious node risks slashing, in which they stand to lose their collateral and rewards. Thus, node operators are disincentivized from acting against the network.

3.10 Checkpoints

Checkpoint is a mechanism employed as an additional layer of security within the network. Checkpoint works by taking a snapshot of the blockchain history every 4 blocks. Thus, even if the Beldex network were to be compromised by a group of bad actors banding together, they will not be able to alter more than 4 blocks, since everything prior to it is recorded.

3.11 Bucephalus Hardfork

The Beldex network transitioned to Proof of Stake on Dec 10, 2021, via the Bucephalus Hardfork. The proof of stake consensus introduced several changes to the network, a few of them being reduced block production time. The previous 2 minute block time is reduced four times and now sits at an average of 30 seconds. The transaction size is significantly reduced while block rewards have increased $5\times$ from 2 BDX to 10 BDX. 62.5% of the new block rewards are allocated to masternode block producers while 37.5% is allocated to governance. Flash instant transactions were launched with a network confirmation of 2 nodes.

3.12 Implementation of Bucephalus PoS Consensus

The Bucephalus consensus works on a number of quorum states that run simultaneously and have different functionalities.

For example, the PoS quorum determines the block producer and the 11 validator nodes. The Flash quorum is based on two separate consensus groups that verify the transaction and then include them in a block. The Masternode quorum checks the nodes' uptime proof before being selected to participate in block validation while the checkpoint quorum takes a snapshot of the chain history every 4 blocks.

3.12.1 PoS Quorum

A POS quorum is made up of 12 nodes in which 1 node produces the block and the remaining 11 nodes validate it. The producer selects valid transactions from the mempool and creates the nominee block. The producer then signs the nominee block and broadcasts

it to the validator nodes in the consensus group. Each validator node broadcasts it to 3 other nodes in the network to ensure that all the nodes receive the block.

Each validator now creates a 128 bit integer and sends it to other validators. Validators reveal the pre-image of the hash, so that each validator now has a subgroup values. Validator combine the pre-images to produce a new random value. The new random value is added to the nominee block and then signed. The signature is then sent to all other validators.

The validators then check each other's signatures, ensuring they are all signing both the same nominee block and the same combined random integer. Once signatures have been validated, any validator in the quorum can submit the block to the network. For the block to be confirmed, 7/11 validator nodes should approve it.

Quorum state is checked every 1 minute. This could be set to less than 1 minute, however, a handshake between 12 nodes in a quorum requires at least 30 seconds. Thus, to ensure a quorum is formed, the maximum time is set to 1 minute, while in reality, it could be lesser. On the other hand, the blocktime is adjusted such that the average block creation time is 30 seconds.

A masternode that produces the block is rewarded 6.25 BDX and assumes the last position in the reward queue. The successive block is produced once the next set of 12 quorum nodes is selected from the 50 nodes that were set aside in Masternode Quorum.

3.12.2 Flash Quorum

Flash transactions are instant and have a slightly higher fee than regular transactions (which is still orders of magnitude lesser than most other blockchain networks). Flash transactions differ from other transactions in the following way. Regular $\$BDX$ transactions are sent to the mempool where they wait to be included in a block. However, flash transactions are sent to two different quorums that instantly verifies them. The sender and receiver see that the locked balances have been updated on their wallets. Transactions are then sent to the mempool to be included in the block.

3.12.3 Masternode Quorum

The masternode quorum checks if any given masternode is active and eligible for the next round of PoS quorum. 10 validator nodes in POS quorum check other masternodes for proof of uptime. 50 nodes are set aside for this process. The nodes that do not submit uptime proof are decommissioned in order to prevent them from obstructing the network.

3.12.4 Checkpoint Quorum

As mentioned in section 3.10, checkpoints are an additional layer of security. Checkpoint quorums take a snapshot of the blockchain and record the hash of every 4th block. Thus, an attacker can virtually only be able to change 4 blocks. Beyond that, the network will self-correct using the snapshot history from the Checkpoints.

3.13 Byzantine Fault Tolerance

When a valid block is not produced for more than 60 seconds, it results in a pause in block production. When this happens, the next block producer can produce the next nominee block and send it to the validators by including a switch tag. This tag indicates to new handshaking nodes that there has been a failure in creating a quorum.

Fault cases: (1) the block producer doesn't create the block, (2) less than 11 validators reveal the pre-image of the hash, (3) validators signing two different blocks (4) validators reveal different values, thus they're unable to submit the block.

3.13.1 Block producer doesn't create the block

If a block producer fails to construct and send a nominee block, the validators will be triggered to submit intent to demerit message. If this message is signed by a simple majority of the quorum, a demerit point is to be applied to the producer and the next producer creates a switch block.

3.13.2 Less than 11 validators reveal the pre-image of the hash

Although all Beldex Master Nodes are incentivised to cooperate in order to avoid demerits, some participants may either act maliciously or simply fail. The protocol should be able to tolerate the failure of 6 of 11 validator nodes during the reveal pre-image of hash stage. As above, any validator that doesn't participate in this round, or reveals a different value from their commitment can be flagged by other validators with an intent to demerit message. If this message is signed by a majority of nodes, this message is submitted to the network as evidence of a node's non-performance in a block creation quorum.

If more than 6 of 11 nodes don't participate in the commit or reveal round, the quorum will not be able to create a block - which means the next quorum can take over block creation once a timeout occurs.

3.13.3 Validators reveal different values, thus they're unable to submit the block

There may be a number of cases where not all of the masternodes correctly receive all of the committed or revealed values. This will result in differing final random integers being calculated by the validators.

As with other errors, the scheme will progress normally as long as the revealed values for 7 of the 11 validators match. Nodes with invalid revealed values can query the other validators in their quorum again to receive any values they missed during the two rounds. If more than 7 nodes disagree on the final revealed value, then the scheme will fail once a timeout occurs and the next producer produces a switch block.

3.13.4 Validator signs two different blocks

Any masternode node, including a validator or leader, may submit a deregistration transaction if they can produce signatures from any other masternode that signs multiple

blocks which compete for the same height or signs a competing chain once more than two checkpoints have been applied.

3.14 Decommission and Deregistration

Validators can send intent to decommission message if they deem that any peer hasn't followed the rules (e.g. the block producer failing to create a nominee block) or if a validator doesn't submit proof of uptime for any of the services such as Beldex Storage Server and BelNet.

Block Credits The decommissioned nodes have a specific block credit before which they need to submit an uptime proof. Block credits are received by the nodes based on the time elapsed since the masternode joined the network. Each masternode is awarded 96 block credits in a day approx. The maximum block credit is 5760 blocks. Thus, the longer the masternode has been active in the network, the greater will be its block credits.

4 Beldex Services

4.1 Beldex Decentralized Wallets

4.1.1 Android & iOS Wallet

The Android and iOS wallets come with the same features as the Beldex electron wallet. The mobile wallets provide instant access to your Beldex (BDX). Store, send, receive, and even stake BDX (feature to be launched in the upcoming release). You can download the Beldex wallet for your Android device and iOS device [here](#).

4.1.2 Beldex Electron Desktop Wallet

The Beldex Electron desktop wallet is a DApp built for Windows, Linux, and Mac. The electron desktop wallet is fast, secure, and decentralized. It can be used to hold, send, and receive BDX. In addition to this, the electron wallet provides an extensive and easy to use GUI to setup and host your Beldex masternode. You can download it [here](#).

4.2 Decentralized Messaging App (B-Chat)

B-Chat is an anonymous messaging application that routes messages through BelNet. While modern applications like WhatsApp and Signal provide end-to-end encryption, the Signal protocol [CGCD⁺17] is particularly noted for its high degree of confidentiality. However, end-to-end encryption alone does not guarantee complete confidentiality, as data is ultimately stored on centralized servers. These centralized data centers are vulnerable to attacks and are often inadequately protected against government policies and interventions. Authorities can compel centralized messaging services to share user data, undermining the confidentiality that these applications aim to provide.

B-Chat introduces routing of messages through the BelNet, a decentralized Virtual Private Network (dVPN). The decentralized VPN, the BelNet, ensures that messages are routed to their destination through the shortest possible node distance while encrypting

them at each hop. The minimum hop length is set as 3, not counting the origin and destination nodes. This prevents a dishonest node from sabotaging the message as it only receives an encrypted form of the message to begin with.

Messages can be sent by the sender even if the receiver is online or offline, with the help of Beldex Secure Masternodes. When the receiver is online, the message packets get routed through the Masternodes. The secure Masternodes act as the routing channel when user *A* wants to send a message to user *B*. The Masternode determines the shortest possible distance over which the message can be routed on the routing channel. To do this, the origin node (node where the message originates) gets the destination nodes' information from the distributed decentralized directory stored by the Masternodes. Then, it establishes a safe route while checking for available bandwidth and traffic congestion. It then transmits the message, which gets encrypted at each node hop and a minimum of three node hops are required for each message packed. The message is received by the destination node which in turn relays it to the user *B*. When the user *B* is offline, the closest node, or the destination node, stores the message and relays it when they return online.

4.3 BelNet

Virtual Private Networks (VPNs) have allowed individuals to maintain anonymity in their online activities. However, many current VPNs are not fully private, as they are typically operated by centralized entities. Existing onion routing protocols, like Tor, provide limited decentralization, with their Distributed Hash Tables (DHTs) being overseen by a small group of nodes known as Directory Authorities⁴. As a result, these systems offer only basic levels of confidentiality.

In contrast, an onion routing protocol implemented on a blockchain network can achieve complete decentralization, providing enhanced confidential routing. BelNet is a private decentralized VPN that operates using the Beldex Routing Protocol (BRP). The BRP optimizes communication between nodes, ensuring that traffic is routed through the network with minimal load on secure Masternodes.

The BRP uses the packet switching routing model and supports a wide range of internet protocols. It's Distributed Hash Table that lists the nodes on the network is maintained by Masternodes, based on a node limiting and authorizing mechanism. The governing Masternodes are a set of nodes that get randomly assigned to verify an incoming new node.

A new node must meet the bandwidth and staking requirements to obtain the status of a relay or exit node on the BelNet. The governing nodes are reassigned to each node and at fixed time intervals, they keep verifying if a node has sufficient bandwidth. If any node fails to meet the bandwidth requirements, then they are penalized by the governing nodes and their rewards get cut until they provide the required bandwidth. The significant staking capital also prevents the nodes from acting dishonestly in the network.

⁴<https://www.technologyreview.com/2012/04/04/186902/how-china-blocks-the-tor-anonymity-network/>

4.4 Beldex Browser

The online advertising industry often relies on invasive tracking of user activity, disregarding user confidentiality on a daily basis. Many applications even monitor users' offline data. In the current advertising landscape, data processors and intermediaries siphon off a significant portion of ad revenue from content creators, while large corporations dominate the advertising market⁵. To address these issues, Beldex introduces the Beldex browser—a secure, confidentiality-focused browser designed to empower both advertisers and content creators while safeguarding user data.

The Beldex browser routes user traffic using BelNet and the Beldex Masternodes, which anonymizes the source and destination of the traffic. Users choose the content they want to pay their attention while content creators and advertisers can target a more specific audience. Donning the peer-to-peer agenda, content authors can publish ad-free content to their audience who may choose to pay creators for their content with BDX.

5 Marketing

5.1 Community

Beldex has had a very strong community base since inception. Beldex has various channels for engaging with the community. You can find the full list of announcement boards and social media below.

- Telegram Announcements: https://t.me/official_beldex
- Telegram Chat: <https://t.me/beldexcoin>
- Twitter: <https://twitter.com/BeldexCoin>
- Discord: <https://discord.gg/Hj4MAmA5gs>
- Facebook: <https://www.facebook.com/beldexofficial>
- Instagram: <https://www.instagram.com/beldexcoin/?hl=en>
- LinkedIn: <https://www.linkedin.com/company/beldex-coin>
- Medium: <https://beldexcoin.medium.com/>

5.2 Bounty and Airdrop

Tokens reserved for the purpose of bounty and airdrop are 80,000 BDX tokens valued at \$10,000 USD. Airdrops will be conducted at regular intervals at the discretion of Beldex.

⁵<https://fortune.com/2017/01/04/google-facebook-ad-industry/>

6 Contributions

Beldex contributors are primarily masternode validators who strengthen the network by setting up masternodes and validating blocks. To contribute to the Beldex network, you can either set up a [dedicated masternode](#) or stake on platforms that support Beldex masternode staking. Validators can also share the minimum collateral of 10000 BDX to set up a masternode and for this purpose, Beldex will partner with shared masternode pools where they can easily stake a BDX amount of their choosing and receive rewards proportional to their stake.

Independent developers who are interested in contributing may contribute to Beldex on their own volition. Developers can find the open source Beldex source code [here](#). To contribute, leave a message on our discord channel and our team will get back to you.

References

- [AHJ17] Kurt M. Alonso and Jordi Herrera-Joancomartí. Monero - privacy in the blockchain. *IACR Cryptol. ePrint Arch.*, 2018:535, 2017.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 315–334, 2018.
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo G. Desmedt, editor, *Advances in Cryptology — CRYPTO '94*, pages 174–187, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.
- [CGCD⁺17] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *2017 IEEE European Symposium on Security and Privacy (EuroSP)*, pages 451–466, 2017.
- [DH76] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [LWW04] Joseph K. Liu, Victor K. Wei, and Duncan S. Wong. Linkable spontaneous anonymous group signature for ad hoc groups. In Huaxiong Wang, Josef Pieprzyk, and Vijay Varadharajan, editors, *Information Security and Privacy*, pages 325–335, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [MSBL⁺19] Pedro Moreno-Sanchez, Arthur Blue, Duc V. Le, Sarang Noether, Brandon Goodell, and Aniket Kate. Dlsag: Non-interactive refund transactions for interoperable payment channels in monero. *Cryptology ePrint Archive*, Paper 2019/595, 2019. <https://eprint.iacr.org/2019/595>.
- [Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.

- [NMRL16] Shen Noether, Adam Mackenzie, and the Monero Research Lab. Ring confidential transactions. *Ledger*, 1:1–18, Dec. 2016.
- [PD16] Poon and Dryja. The bitcoin lightning network: Scalable off-chain instant payments. 2016.
- [QC19] Mikerah Quintyne-Collins. Short paper: Towards characterizing sybil attacks in cryptocurrency mixers. Cryptology ePrint Archive, Paper 2019/1111, 2019. <https://eprint.iacr.org/2019/1111>.
- [RST01] Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In Colin Boyd, editor, *Advances in Cryptology — ASIACRYPT 2001*, pages 552–565, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [vS13] Nicolas van Saberhagen. Cryptonote v 2.0. 2013.