

Abstract

This article presents Lelantus-MW, an adaptation of Lelantus [Jiv] for use with the Mimblewimble protocol. Whereas Mimblewimble achieves a limited amount of privacy by merging transactions that occur in the same block, Lelantus uses a logarithmic-size proof of membership to effectively enable merging across different blocks. At a high level, this allows value to be added to a common pool and then spent privately, but only once. We explain how to adapt this construction to Mimblewimble, while at the same time simplifying the protocol where possible.

At the same time, we also describe how we can use Bulletproof “coloring” to enable offline payments, thus addressing one of the original shortcomings of Mimblewimble.

Lelantus-MW

Pyrros Chaidos Vladislav Gelfer
University of Athens Beam Foundation

September 9, 2020

1 Introduction

In this article we describe Lelantus-MW, an adaptation of Lelantus [Jiv] for use with the Mimblewimble protocol.

Mimblewimble achieves a limited amount of privacy by merging transactions that occur in the same block, effectively mixing together all transactions that were submitted to the same block. This has significant privacy improvements when compared with e.g. Bitcoin in which transactions are individually traceable and mixing (or tumbling) requires the use of a separate service. However, the degree of privacy that is offered is limited, and is not under the control of the user. While large blocks offer better mixing, the user has little way of ensuring that.

Lelantus uses a logarithmic-size proof of membership to effectively enable merging across different blocks. At a high level, this may be thought of as a large common pool of value to which individuals can make deposits and withdrawals in private. The key feature is of course that any particular deposit can only be withdrawn once. We explain how to adapt this construction to Mimblewimble, while at the same time simplifying the protocol where possible.

At the same time, we also describe how we can use Bulletproof “coloring” to enable offline payments, thus addressing one of the original shortcomings of Mimblewimble. In our context coloring refers to hiding values inside the randomness used in the construction of zero knowledge proofs. This enables those with knowledge of the appropriate seed to detect and extract these values. With this functionality in place, senders can send funds without needing to interact with recipients, as the recipients are able to extract the necessary information on their own.

1.1 Outline

In Section 2 we give the necessary cryptographic definitions for the core concepts used in this work. Next, in Section 3 we describe the constructions which use when designing and implementing Lelantus-MW. Our choices are well-studied constructions, requiring limited assumptions.

Before delving into Mimblewimble-LW (Section 5), Section 4 revisits Mimblewimble to establish our baseline and better highlight the improvements in our construction.

Finally, in 6, we describe the technical details required to conduct transactions in a meaningful way (i.e ensure privacy and integrity are upheld). This includes the concept of bulletproof coloring, as well as the notions of tickets and addresses.

2 Preliminaries

2.1 The Discrete Log Setting

Let GGen be an algorithm that on input 1^λ returns $gk = (\mathbb{G}, p, g)$ such that \mathbb{G} is the description of a finite cyclic group of prime order p , where $|p| = \lambda$, and g is a generator of \mathbb{G} .

Definition 1 (Discrete Logarithm Assumption). *The discrete logarithm assumption holds relative to GGen if for all non-uniform polynomial time adversaries \mathcal{A}*

$$\Pr \left[(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda); h \leftarrow \mathbb{G}; a \leftarrow \mathcal{A}(\mathbb{G}, p, g, h) : g^a = h \right] \approx 0$$

2.2 Commitment Schemes

Commitment schemes allows one to commit to a value v without revealing it. We use $\text{Com}_{ck}(v; r)$ to refer to a commitment to a value v using randomness r under a previously determined key ck . A commitment can later be opened to reveal the value contained. We require that commitments are hiding i.e it must be difficult to correctly guess the value contained in a commitment given the key and the commitment itself. We also require that commitments are binding i.e. it must be difficult for a malicious opener to create a commitment and open it to two different values.

Definition 2 (Perfectly hiding). *We say a non-interactive commitment scheme $(\text{CGen}, \text{Com})$ is perfectly hiding if a commitment does not reveal the committed value. For all non-uniform polynomial time stateful interactive adversaries \mathcal{A}*

$$\Pr \left[\begin{array}{l} ck \leftarrow \text{CGen}(gk); (m_0, m_1) \leftarrow \mathcal{A}(ck); \\ b \leftarrow \{0, 1\}; c \leftarrow \text{Com}(m_b) \end{array} : \mathcal{A}(c) = b \right] = \frac{1}{2}$$

where \mathcal{A} outputs $m_0, m_1 \in \mathbb{Z}_p$.

Definition 3 (Binding). *A non-interactive commitment scheme $(\text{CGen}, \text{Com})$ is computationally binding if a commitment can only be opened to one value. For all non-uniform polynomial time adversaries \mathcal{A}*

$$\Pr \left[\begin{array}{l} ck \leftarrow \text{CGen}(gk); \\ (m_0, r_0, m_1, r_1) \leftarrow \mathcal{A}(ck) \end{array} : \begin{array}{l} \text{Com}(m_0; r_0) = \text{Com}(m_1; r_1) \\ \text{and } m_0 \neq m_1 \end{array} \right] \approx 0$$

where \mathcal{A} outputs $m_0, m_1, r_0, r_1 \in \mathbb{Z}_p$.

2.2.1 Homomorphic Commitments

An important requirement for our application is that they are additively homomorphic, that is the sum of two commitments is a commitment to the sum of the values contained within. We require that the property holds for values and randomisers alike.

$$\text{Com}(v; r) + \text{Com}(v'; r') = \text{Com}(v + v'; r + r')$$

2.3 Zero Knowledge Arguments of Knowledge

Let R be a polynomial time decidable binary relation, i.e., a relation that defines a language in NP. We call w a witness for a statement u if $(u, w) \in R$. We define the language

$$L_R = \{x \mid \exists w : (x, w) \in R\}$$

as the set of statements x that have a witness w in the relation R .

A zero knowledge protocol for an NP-relation R enables a prover to demonstrate to a verifier that a statement u satisfies $u \in L_R$, i.e. that there exists a witness w such that $(u, w) \in R$ without disclosing anything else, in particular not disclosing the value of w that the prover has in mind.

More formally, a protocol for a relation R w.r.t. a setup $gk \leftarrow \text{GGen}(1^\lambda)$ and a common reference string crs is a tuple $(\text{CRSGen}, \mathcal{P}, \mathcal{V})$. If a common reference string, crs is used, we assume that the generator, $\text{CRSGen}(gk)$ generates it before the execution of the protocol. As we are interested in statements about commitments, we additionally consider the commitment key ck , part of the crs . Both gk and crs are provided as implicit inputs to all algorithms. We consider that the prover \mathcal{P} and the verifier \mathcal{V} , are both probabilistic polynomial time interactive algorithms. The transcript produced by \mathcal{P} and \mathcal{V} when interacting on inputs s and t is denoted by $tr \leftarrow \langle \mathcal{P}(s), \mathcal{V}(t) \rangle$.

We require that a protocol works correctly when used by honest participants (*completeness*), that it prevents provers from convincing verifiers about a statement unless they know a witness (*knowledge soundness*), and that it does not leak information about w (*zero-knowledge*).

Definition 4 (Public coin). *A protocol $(\mathcal{P}, \mathcal{V})$ is called public coin if the verifier chooses his messages uniformly at random and independently of the messages sent by the prover, i.e., the challenges correspond to the verifier's randomness ρ .*

Definition 5 (Perfect completeness). *$(\mathcal{P}, \mathcal{V})$ has perfect completeness if for all non-uniform polynomial time adversaries \mathcal{A}*

$$\Pr \left[(u, w) \leftarrow \mathcal{A}(1^\lambda) : (u, w) \notin R \text{ or } \langle \mathcal{P}(x, w), \mathcal{V}(x) \rangle = 1 \right] = 1$$

Definition 6 (n -Special Soundness). *We say that $(\mathcal{P}, \mathcal{V})$ is n -special sound [GK15] if there exists an efficient extractor χ such that for any statement u , given n accepting transcripts $\{(a, x_i, z_i)\}_{i=1}^n$ where the challenges x_i are distinct, $\chi(u, a, x_i, z_i)$ outputs w s.t. $(u, w) \in R$.*

For more complex protocols, involving a large number of rounds n -Special Soundness can be hard to obtain so we make use of the related notion of witness-extended emulation introduced by Lindell [Lin03] as it is used in [GI08, BBB⁺18]. We refer the reader to these texts for details.

A protocol is zero-knowledge if it does not leak information about the witness beyond what can be inferred from the truth of the statement. Special honest verifier zero-knowledge further specifies that if the verifier's challenges are known in advance, then it is possible to simulate the entire argument without knowing the witness.

Definition 7 (Special Honest Verifier Zero-Knowledge (SHVZK)). *A public coin argument $(\mathcal{P}, \mathcal{V})$ is called a special honest verifier zero knowledge (SHVZK) argument for R if there exists a probabilistic polynomial time simulator \mathcal{S} such that for*

all interactive non-uniform polynomial time adversaries \mathcal{A}

$$\begin{aligned} & \Pr \left[(u, w, \rho) \leftarrow \mathcal{A}(1^\lambda); tr \leftarrow \langle \mathcal{P}(u, w), \mathcal{V}(u; \rho) \rangle : (u, w) \in R \text{ and } \mathcal{A}(tr) = 1 \right] \\ & \approx \Pr \left[(u, w, \rho) \leftarrow \mathcal{A}(1^\lambda); tr \leftarrow \mathcal{S}(u, \rho) : (u, w) \in R \text{ and } \mathcal{A}(tr) = 1 \right] \end{aligned}$$

where ρ is the public coin randomness used by the verifier.

If this holds also for unbounded adversaries \mathcal{A} , we say $(\mathcal{P}, \mathcal{V})$ is statistically special honest verifier zero-knowledge. If the two distributions are equal, then we say $(\mathcal{P}, \mathcal{V})$ is perfect special honest verifier zero-knowledge.

3 Constructions

Generalised Pedersen Commitments Pedersen commitments [Ped91] are a widely used homomorphic commitment scheme. Under the discrete logarithm assumption, Pedersen Commitments are perfectly hiding and computationally binding. In this paper we will also use a variant of the standard scheme that makes use of two randomisers. We will use r, s to refer to the randomisers and $\text{Com}(v, r, s)$ to refer to the commitments. The key generation and commitment functions of the pedersen scheme are as follows:

$$\begin{aligned} \text{CGen}(gk) &: h, k \leftarrow \mathbb{G}; \text{Return } ck := (g, h, k) \\ \text{Com}_{ck}(v, r, s) &: \text{Return } c = g^r h^v k^s \\ \text{Com}_{ck}(v, r) &: \text{Return } c = g^r h^v \end{aligned}$$

We point out that $\text{Com}_{ck}(v, r)$ is functionally identical to $\text{Com}_{ck}(v, r, 0)$ due to the structure of the commitment scheme. The (generalised) Pedersen Commitments are a core component of our construction, so ck is used in defining the protocol relations.

3.1 Range Proofs

The additive homomorphic property of Pedersen Commitments is a first step towards ensuring that sums of values “add up”. However, since the values themselves remain hidden, we need to enforce that users are unable to enter negative values; otherwise one might pay \$1 for a \$5 coffee and receive -\$4 as change. Such a transaction would in principle balance, but in practice the intended recipient of a negative note would simply ignore it¹. We thus need to disallow commitments to negative values. We do this by requiring that committed values lie in a certain range and use a zero knowledge protocol to enforce that without having to reveal the values.

Bulletproofs Bulletproofs are a general purpose zero knowledge proving system that can be tailored to perform efficient range proofs. Given a bound v_{\max} , the

¹Strictly speaking, a small negative value would be equivalent to a large positive value since arithmetic happens modulo q .

original construction from [BCC⁺15] can be used directly to prove knowledge of a witness (r, v) for the relation:

$$R_\rho = \{C, (r, v) : C = g^r h_1^v \wedge (0 \leq v < v_{\max})\}$$

Bulletproofs require a CRS structured as a Pedersen commitment key:

$$\begin{aligned} \text{CRSGen}_\rho(gk) : g_1, \dots, g_n, h_1, \dots, h_n &\leftarrow \mathbb{G}, \text{ for } n = \log_2(v_{\max}) \\ \text{Return } ck &:= (\mathbf{g}, \mathbf{h}) \end{aligned}$$

In section 5, we slightly modify the protocol to add a second commitment to the statement:

$$R'_\rho = \{C, \hat{C}, (r, v) : C = g^r h_1^v \wedge (0 \leq v < v_{\max})\}$$

This is immaterial for the interactive version, but when applying the Fiat-Shamir transformation, this ensures that a proof can only be verified in relation to \hat{C} as well as C .

Schnor Signatures In our applications we will often need to prove that the value of certain commitments is exactly zero. While it would be possible to do that with a range proof with a very restricted range, it is more efficient to do so directly. We use the term “signatures” as opposed to proofs of zero value due to convention. We use Schnor signatures to prove knowledge of a witness r for the relation:

$$R_\sigma = \{C, r : C = g^r\}$$

We will later use a generalised version for the relation:

$$R_{\sigma'} = \{C, (r, s) : C = g^r h_2^s\}$$

We can also use the same protocol with a different base to show that a commitment has a zero in the serial component.

$$R_{\sigma''} = \{C, (r, s) : C = g^r h_2^s\}$$

3.2 Logarithmic Membership Proofs

The protocol of Groth and Kohlweiss [GK15], later optimized in [BCC⁺15] allows one to prove they know the opening of a single commitment in a set of many without specifying which one. The size of the proofs is logarithmic in the size of the set which enables us to be efficient. Assuming a maximum set size of m_{\max} the protocol requires a common reference string as follows:

The protocol requires² a CRS with structure similar to a Pedersen commitment key:

$$\begin{aligned} \text{CRSGen}_m(gk) : g_1, \dots, g_n &\leftarrow \mathbb{G}, \text{ for } n = \log(v_{\max}) \\ \text{Return } ck &:= \mathbf{g} \end{aligned}$$

²In practice, as the structure is identical to the bulletproof *crs* we are able to reuse it if length allows.

The protocol enables us to reveal a serial number s and prove that a commitment C contains that serial number, and is contained in set \mathbf{C} , as in the following relation:

$$R_m = \{(\mathbf{C}, s), r : C = g^r k^s \wedge (C \in \mathbf{C})\}$$

3.3 The Fiat-Shamir Heuristic

The protocols we describe above operate in multiple rounds which is not practical for our application. The Fiat-Shamir transformation [FS87] produces a non-interactive version of a Σ -protocol by substituting the verifier's challenge with the output of a cryptographic hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ on the protocol's statement and current transcript. The transformation can be proven secure in the random oracle model [BR93]. The transformation crucially relies in the public coin and HVZK property of the protocols.

4 Mimblewimble

In this section we briefly describe the original version of Mimblewimble as adapted from [Jed]. For simplicity, we omit the discussion of fees, minting and block structure.

4.1 Transactions

Intuitively, a transaction in Mimblewimble consists of gathering an amount of value stored in existing notes and re-distributing it in a set of new notes, so that the sums of value in the old (input) notes and the new (output) notes are equal. This is augmented by proofs certifying all the new notes are within the allowed range of values, and by a number of special zero-value notes (kernels) that exist to maintain privacy.

Formally, a transaction T consists of three lists: $\mathsf{T.I}$, $\mathsf{T.O}$, $\mathsf{T.K}$. The list $\mathsf{T.I}$ consists of commitments C_{I1}, C_{I2}, \dots . The list $\mathsf{T.O}$ consists of tuples $(C_{O1}, \pi_{O1}), (C_{O2}, \pi_{O2})$, such that C_{Oj} is a commitment and π_{Oj} is a range proof on C_{Oj} . The list of kernels $\mathsf{T.K}$ consists of tuples $(C_{K1}, \sigma_{K1}), (C_{K2}, \sigma_{K2})$, such that C_{Kl} is a commitment and σ_{Kl} is a signature demonstrating C_{Kl} is zero-valued.

Definition 8. *A transaction $\mathsf{T} = (\mathsf{T.I}, \mathsf{T.O}, \mathsf{T.K})$ is valid if:*

- *The list of commitments $\mathsf{T.I}$ is a subset of the current UTXO set.*
- *For each commitment C_{Oj} , the corresponding proof π_{Oj} is valid.*
- *For each commitment C_{Kl} the corresponding signature σ_{Kl} is valid.*
- *The the sum of inputs minus the sum of outputs is equal to the sum of zero-valued commitments: $\sum_i C_{Ii} - \sum_j C_{Oj} = \sum_l C_{Kl}$.*

4.2 Merging Transactions

The way transactions are structured is highly conducive to merging: given two transactions $\mathsf{T}_1 = (\mathsf{T}_1.I, \mathsf{T}_1.O, \mathsf{T}_1.K)$, $\mathsf{T}_2 = (\mathsf{T}_2.I, \mathsf{T}_2.O, \mathsf{T}_2.K)$ such that $\mathsf{T}_1.I$ has no common elements with $\mathsf{T}_2.I$ and $\mathsf{T}_1.O$ has no common elements with $\mathsf{T}_2.O$, we can

merge³ them as: $T_M = (T_1.I \cup T_2.I, T_1.O \cup T_2.O, T_1.K \cup T_2.K)$. Furthermore, if a commitment appears in both the input and output lists, it can be deleted (cut) from both without affecting the sum. With this in mind we arrive at:

$$\begin{aligned} T_{M.I} &:= (T_1.I \cup T_2.I) \setminus (T_1.O \cup T_2.O) \\ T_{M.O} &:= (T_1.O \cup T_2.O) \setminus (T_1.I \cup T_2.I) \\ T_{M.K} &:= T_1.K \cup T_2.K \end{aligned}$$

We consider the ledger L as a special transaction with an empty input set. The output set of the ledger is termed the UTXO set.

$$\text{UTXO} := L.O$$

The ledger is updated by merging valid transactions onto it. The issue of deciding which transactions are to be merged and in which order is outside the scope of this text. When a transaction is accepted by the network, its outputs are added to the UTXO set, and its inputs deleted from it.

In practice, this can be used to merge all submitted transactions into one transaction per block, and subsequently merge all block transactions into a single transaction. This practice of merging and cutting can decrease the storage requirements while still allowing one to fully verify transactions. It also provides a small level of confidentiality.

4.3 Mimblewimble Anonymity

Due to merging, Mimblewimble blocks contain a single transaction, produced as a merging of all user-submitted transactions during that time. This can obfuscate the transaction graph in the view of adversaries with access only to the block-level, but individual transactions will still be distinguishable with regard to adversaries with access to the network layer: even with the values hidden, “who paid whom” will be in the clear.

5 Lelantus-MW

While the standard Mimblewimble protocol offers some privacy by obfuscating the transaction graph in the block level, it has two shortcomings: first, small blocks offer less obfuscation and second, the layer of obfuscation collapses if the adversary is able to closely monitor the network. Lelantus offers a design that enables *shielded* transactions, i.e. transactions that store or withdraw value from a universal set without revealing which element they are withdrawing from. This raises the issue of double-spending: if spent elements are unknown they cannot be removed from the set, and thus users might attempt to spend them again. This is solved by adding an extra “serial number” field to the commitments of shielded values. When a shielded value is being spent the serial number is revealed and checked against a list of previously-revealed serial numbers. This prevents double spending (due to the binding property of the commitment scheme), without compromising privacy as serial numbers are not revealed when shielded values are created.

³We note here that we allow repetitions in the kernel lists.

Given the above discussion, we extend our definition of transactions in two ways:

- We allow shielded inputs. Whereas a transparent input “justifies” its existence by appearing in the UTXO set, a shielded one is supported by the combination of a membership proof and a single-use serial number.
- We allow shielded outputs. Whereas transparent outputs simply enter the UTXO set, shielded ones are combined with an additional masking factor and added to the STXO set.

Shielded outputs are in part similar to transparent ones: they have a commitment containing their value and corresponding range proof. To disallow double-spending, they also need to have a serial number, which is hidden in a second commitment, to be combined with the first one when added to the STXO set,. Obviously, the second commitment must carry zero value. Because we wish to exclude serial number components from entering the balance sum, we need to tie the two commitments together in some fashion, which we accomplish by adding the second one to the statement of the range proof.

A shielded output is a tuple $(C_{SO}, \pi_{SO}, \hat{C}_{SO}, \sigma_{SO})$ such that:

- C_{SO} is a commitment (i.e a group element)
- π_{SO} is a proof for the relation R'_ρ with statement C_{SO}, \hat{C}_{SO} (i.e a range proof counter-signing \hat{C}_{SO}).
- \hat{C}_{SO} is a generalised commitment (i.e a group element)
- σ_{SO} is a proof for relation $R_{\sigma'}$ for statement \hat{C}_{SOj} (i.e a proof that \hat{C}_{SO} has zero value).

The tuple $\hat{C}_{SO}, \sigma_{SO}$ is called a *ticket* because it can be pre-generated ahead of time and also because it contains the serial number. To prevent (potentially accidental) misuse leading to loss of funds we require that tickets are unique. Concretely, we check that \hat{C}_{SO} has not been used previously.

Shielded inputs are different. Like transparent ones, they contain a commitment carrying value. However, that commitment does not appear (directly) in the STXO set as that would provide no privacy. The input also provides a group element which is interpreted as a hash pre-image of the serial number, accompanied with a signature proving knowledge of the discrete log. The last part of the input is a membership proof of the combined commitment and serial number in STXO.

A shielded input is a tuple $(C_{SI}, \sigma_{SI}, Q_{SI}, \sigma'_{SI}, \pi_{SI})$ such that:

- C_{SI} is a commitment (i.e a group element)
- σ_{SI} is a proof for relation $R_{\sigma''}$ for statement C_{SI} (i.e a proof that C_{SI} has zero serial number).
- Q_{SI} is a group element.
- σ'_{SI} is a proof for relation R_σ for statement Q_{SI} (i.e a proof of knowledge for the log of Q_{SI}).
- π_{SI} is a proof for the relation R_m on statement (C, s) where $C = \frac{\text{STXO}}{C_{SI}}$ and $s = H(Q_{SI})$

With the above changes in mind we are able to adapt Definition 8 as follows:

Definition 9. A transaction $T = (T.I, T.SI, T.O, T.SO, T.K)$ is valid if:

- The list of commitments $\mathsf{T.I}$ is a subset of the current UTXO set.
- For each commitment $C_O \in \mathsf{T.O}$, the corresponding proof π_O is valid.
- For each commitment C_{KI} the corresponding signature σ_{KI} is valid.
- For each shielded input, $(C_{SI}, \sigma_{SI}, Q_{SI}, \sigma'_{SI}, \pi_{SI})$ the corresponding proofs, $\sigma_{SI}, \sigma'_{SI}, \pi_{SI}$, are valid, and $H(Q_{SI}) \notin \mathsf{USN}$.
- For each shielded output, $(C_{SO}, \pi_{SO}, \hat{C}_{SO}, \sigma_{SO})$ the corresponding proofs π_{SO}, σ_{SO} are valid.
- The the sum of inputs minus the sum of outputs is equal to the sum of zero-valued commitments: $\sum_i C_{Ii} + \sum_{i'} C_{SIi'} - \sum_j C_{Oj} - \sum_{j'} C_{SOj'} = \sum_l C_{KI}$.

5.1 The Shielded Set

As with standard Mimblewimble, the non-shielded outputs of the ledger form the UTXO set. We define the shielded (STXO) set as the list of combined shielded output commitments in the ledger:

$$\mathsf{STXO} := \{C_{SOi} \cdot \hat{C}_{SOi} \mid (C_{SO}, \pi_{SO}, \hat{C}_{SO}, \sigma_{SO}) \in \mathsf{L.SO}\}$$

Items cannot be deleted from the shielded set. In order to disallow double-spending, we keep a list of used serial numbers:

$$\mathsf{USN} := \{H(Q) \mid (C_{SI}, \sigma_{SI}, Q_{SI}, \sigma'_{SI}, \pi_{SI}) \in \mathsf{L.SI}\}$$

A note on the Membership Proof Even though the list of shielded outputs consists of generalised commitments, we observe that given a serial number s and a commitment $C = \text{Com}(v; r)$, we can divide the list by C . Assuming one of the original list elements has serial number s and value v with some randomness r' , the corresponding quotient will be $\frac{\text{Com}(v, r', s)}{\text{Com}(v, r)} = \text{Com}(0, r' - r, s) = \text{Com}(0, r'', s)$. This allows us to directly use the membership proof of [GK15, BCC⁺15] without the need to adapt them for the generalised commitments.

6 Transacting using Lelantus-MW

At a high level, sending funds consists of creating an output that the receiver can subsequently open. The receiver can then use the commitment as an input to one of their transactions. The data structures and proofs used for shielded transactions imply that “owners” of value can spend it as they wish, with strong privacy guarantees. However, the “owner” is defined as the person who knows the relevant witnesses to construct a proof. This raises two issues:

- How is it possible for the sender to create commitments they cannot open?
- How will the intended recipient be able to spend them?

The first issue is easy to handle, by means of a Diffie-Hellman like protocol, as with standard mimblewimble. The second issue is the main focus of this chapter. While a trivial solution would be to have the sender simply send the necessary data out of band (e.g. via email), it is obviously preferable if this can happen in-band and at the same time without introducing any additional overheads.

We rely on the concept of “Coloring”, which is originally related to protocols used to ascribe and recognise additional properties to otherwise ordinary bitcoins without breaking compatibility with existing software. In our case, coloring has to do with establishing a covert channel [LMS05, BDI⁺96] over the randomizers used in the zero-knowledge proofs used when creating a commitment (i.e schnorr signatures and range proof bulletproofs). This will enable the intended recipient to both recognise commitments addressed to her, as well as extract the necessary witnesses to actually spend them at a later time.

6.1 Bulletproof Coloring

A requirement for direct anonymous payments is for an in-band method of privately notifying the recipient of a payment. Coloring offers a way of doing so by modifying the prover’s randomness generation inside a bulletproof. For some of the blinders in the proof the corresponding randomness will be generated as a deterministic function of the payee’s address. This is done in a way that is computationally indistinguishable from the standard version for any outsiders, but enables the intended recipient to both detect the payment as one directed to them, and also extract a short message embedded within.

We begin, by reproducing the standard range proof for commitments as described in [Jiv]. For brevity, we omit the logarithmic inner product argument as it is unchanged from the original.




Figure 1: Range proof for R_ρ

In the protocol, the t_1, t_2 values are derived from the following polynomials, where $l(X), r(X)$ are vector polynomials linear in X , and $t(X)$ is a scalar polynomial

quadratic in X :

$$\begin{aligned}
l(X) &= (\mathbf{a}_L - z \cdot \mathbf{1}^n) + \mathbf{s}_L \cdot X && \in \mathbb{Z}_p^n[X] \\
r(X) &= \mathbf{y}^n \circ (\mathbf{a}_R + z \cdot \mathbf{1}^n + \mathbf{s}_R \cdot X) + z^2 \cdot X^2 && \in \mathbb{Z}_p^n[X] \\
t(X) &= \langle l(X), r(X) \rangle = t_0 + t_1 \cdot X + t_2 \cdot X^2 && \in \mathbb{Z}_p[X] \quad (1)
\end{aligned}$$

6.1.1 Recipient verification

In the protocol of Figure 1, we observe that α and ρ are random elements of \mathbb{Z}_p to be used as blinders. As a warmup, let us consider a modified protocol where they are instead generated via a keyed PRF. For example, we replace $\boxed{\alpha \leftarrow \mathbb{Z}_p, \rho \leftarrow \mathbb{Z}_p}$ by $\boxed{\alpha \leftarrow H(S || \text{“alpha”} || \text{seed}), \rho \leftarrow H(S || \text{“rho”} || \text{seed})}$, where S represents the concatenation of reference strings gk, ck, crs with the statement C , and $\text{seed} \in \mathbb{Z}_p$ represents the coloring key.

Theorem 1 (Informal). *The modified range proof protocol is complete, has witness-extended emulation and is computationally HVZK.*

Proof. Completeness holds by inspection, as the modified values of α, ρ are used consistently with the original protocol.

We also observe that witness-extended emulation is not impacted, as the definition is not conditioned on any particular prover. Thus, we are able to use the same extractor to repeat the proof with no changes needed (in a way, we consider the modified prover to be an adversarial one and carry out the extraction regardless).

Assuming the hash outputs are computationally indistinguishable from random ensures that zero knowledge still holds, albeit computationally: if the existing simulator fails then an adversary is able to efficiently distinguish the hash function from a random one. \square

Even though security is only minimally impacted, a user with the correct value of seed will be able to generate α', ρ' independently, and use the challenge x from the transcript to produce $\mu' = \alpha' + \rho' \cdot x$. She can then “check” the value of μ from the transcript against μ' to determine if she is the intended recipient.

6.1.2 Conditional witness extraction

We now observe we can develop the same idea further. We further modify the protocol so that τ_1, τ_2 are also derived via the PRF. This enables us to revisit τ_x , that is $\tau_x = \tau_2 \cdot x^2 + \tau_1 \cdot x + z^2 \cdot r$. Assuming we know the correct value of seed , this enables us to reconstruct τ_1, τ_2 , and using our knowledge of x , this allows us to extract r .

6.1.3 In-Band Data

We are able to expand on our previous idea, by adding some additional data δ as a perturbation to α , i.e. $\alpha \leftarrow H(S || \text{“alpha”} || k) + \delta$. This enables us to recover δ from μ by the above equation. We are still able to check if the value of seed is correct by enforcing a specific range of allowed values for δ that is small compared to the group order q , and accepting the possibility of false positives.

6.1.4 Additional Data

If we additionally know the value specified via the bulletproof, we can also add (and extract) additional data as perturbations in the $\mathbf{s}_L, \mathbf{s}_R$ values. We can perform the extraction from the \mathbf{l}, \mathbf{r} vectors since they only depend on v and public data.

6.1.5 Bulletproof Coloring in practice

In the logarithmic version used in the implementation, we are able to hide one value mod q in each of $\mathbf{s}_L, \mathbf{s}_R$, as they are subsequently compressed, preventing further extraction. We are still able to extract the randomizer r , as well as a third value mod q using α . Extracting any value is conditional on knowing the correct `seed` used in the PRF, and for $\mathbf{s}_L, \mathbf{s}_R$, we also need to know the value v . For this reason, we assume v is contained in the value contained inside α . Note here, that the range of values is significantly smaller than q , so there is additional space for data.

In total, we extract 2 arbitrary values mod q as well as v, r , with room for some additional data (which we leave unspecified). The only additional information required for the receiver to be able to spend the output is knowledge of the log of the serial number pre-image. We will cover this in the next section, as it is not included in the bulletproof.

6.2 Sending and Receiving Shielded Values

We now revisit our main goal: enabling senders to create shielded outputs they cannot spend (and allowing receivers to spend them).

The bulletproof coloring technique described above gives us ample space to extract data from π_{SO} , including one of the needed witnesses (r) for C_{SO} . We will now focus our attention to enabling the receiver to obtain the PRF `seed` as well as to recognise the transaction. We do so by applying similar techniques to the $\hat{C}_{SO}, \sigma_{SO}$ part of the shielded output (i.e the ticket), and also adding a DH-like protocol so that receivers can determine `seed` as well as the serial number pre-image.

6.2.1 Addresses

A receiver first creates a triple keypair $P_{1,g}, P_{1,k}, P_2$ where $P_{1,g} = s_1 \cdot g, P_{1,k} = s_1 \cdot k$ and $P_2 = s_2 \cdot k$ and $s_i \leftarrow \mathbb{Z}_q$. The corresponding secret keys are (s_1, s_2) , used for recognizing and spending. As such, s_1 can be revealed to a semi-trusted service with no risk to funds. The triple $(P_{1,g}, P_{1,k}, P_2)$ is a user's address, and it enables other users to directly pay her.

6.2.2 Tickets

To make a commitment recognisable, the sender will create \hat{C}_{SO} in a specific way. First the sender samples a random value $w \leftarrow \mathbb{Z}_q$ and sets $skp \leftarrow H_1(w)$. Then $Q \leftarrow P_2^{skp}$. Finally $s \leftarrow H(Q)$ and $\hat{C}_{SO} \leftarrow g^{skw}$. Looking forward, the sender does not know the log of Q w.r.t g . The intended recipient however could calculate it as $w \cdot s_2$, if provided with w .

The commitment is used to seed the randomness of the generalised Schnorr signature σ_{SO} as follows: $\text{seed} \leftarrow (P_{1,g}^s \cdot P_{1,k}^w)^{H_1(\hat{C}_{SO})}$. This implies $\text{seed} = \hat{C}_{SO}^{s_1 \cdot H_1(\hat{C}_{SO})}$

i.e the recipient can determine the seed from s_1 and \hat{C}_{SO} . Once the key is determined, the recipient can extract the witnesses s, w from the signature, using the techniques of section 6.1.3 which constitutes a full opening of \hat{C}_{SO} . Combined with the above, this also allows the recipient to calculate the log of Q , which is required for spending, check that the **seed** value is valid, and finally extract the values from the bulletproof.

6.2.3 Ticket Distribution

Looking forward, when the receiver attempts to spend the received commitment, the value of Q will be revealed as Q_{SI} in the shielded input. This will allow the sender to detect that a particular shielded input has been spent. Conceptually, this can be mitigated by having the receiver immediately re-send herself the same value, but that would be wasteful in practice. A more efficient solution is to have the receiver pre-compute a number of tickets and make them available to known parties. This can substitute distributing addresses.

The only required change is that the receiver needs to communicate **seed** in addition to $\hat{C}_{SO}, \sigma_{SO}$, as the calculation to produce it requires either the pair (s, w) or s_1 , of which the sender knows neither. The receiver is able to independently create C_{SO}, π_{SO} at a later time. Note that whilst $\hat{C}_{SO}, \sigma_{SO}$ do not depend on C_{SO}, π_{SO} , the statement of the bulletproof π_{SO} explicitly includes \hat{C}_{SO} so that the two parts are effectively “tied”.

References

- [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. IEEE, 2018.
- [BCC⁺15] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit. Short accountable ring signatures based on ddh. In *European Symposium on Research in Computer Security*, pages 243–265. Springer, 2015.
- [BDI⁺96] Mike Burmester, Yvo G Desmedt, Toshiya Itoh, Kouichi Sakurai, Hiroki Shizuya, and Moti Yung. A progress report on subliminal-free channels. In *International Workshop on Information Hiding*, pages 157–168. Springer, 1996.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM conference on Computer and communications security – CCS 1993*, pages 62–73. ACM, 1993.
- [FS87] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology—Crypto’86*, pages 186–194. Springer, 1987.
- [GI08] Jens Groth and Yuval Ishai. Sub-linear zero-knowledge argument for correctness of a shuffle. In *Advances in Cryptology – EUROCRYPT 2008*, pages 379–396. 2008.
- [GK15] Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 253–280. Springer, 2015.
- [Jed] TE Jedusor. Miblewimble, 2016, defunct hidden service.
- [Jiv] Aram Jivanyan. Lelantus: Towards confidentiality and anonymity of blockchain transactions from standard assumptions.
- [Lin03] Yehuda Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 16(3):143–184, 2003.
- [LMS05] Matt Lepinski, Silvio Micali, and Abhi Shelat. Fair-zero knowledge. In *Theory of Cryptography Conference*, pages 245–263. Springer, 2005.
- [Ped91] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO*, volume 576 of *LNCS*, pages 129–140. Springer, 1991.