

Non-interactive Mimblewimble transactions, revisited

Georg Fuchsbauer¹ and Michele Orrù²

¹ TU Wien, Austria

² UC Berkeley, USA

first.last@{tuwien.ac.at,berkeley.edu}

Abstract. Mimblewimble is a cryptocurrency protocol that promises to overcome notorious blockchain scalability issues and provides user privacy. For a long time its wider adoption has been hindered by the lack of non-interactive transactions, that is, payments for which only the sender needs to be online. Yu proposed a way of adding non-interactive transactions to stealth addresses to Mimblewimble, but we show that it is flawed. Building on Yu and integrating ideas from Burkett, we give a fixed scheme and provide a rigorous security analysis in a strengthening of the previous security model from Eurocrypt’19. Our protocol is considered for implementation by MimbleWimbleCoin and a variant is now deployed as MimbleWimble Extension Blocks (MWEB) in Litecoin.

1 Introduction

Mimblewimble (MW) is a cryptocurrency protocol that addresses the problem of ever-growing blockchain data that needs to be stored by full nodes in the system. While in all other cryptocurrencies, such as Bitcoin, the full transaction history must be kept for ever,¹ in MW, coins can be deleted after having been spent while *maintaining public verifiability of the ledger*. Instead of growing linearly (like Bitcoin [Nak08], whose blockchain is now > 400 GB)², MW-based currencies only need to store the currently existing coins (the *UTXO set*) plus some small data per transaction.

Mimblewimble achieves this by cleverly combining three ideas that were initially envisioned for Bitcoin: (1) *Confidential transactions* [Max15] hide the transacted amount by only containing *commitments* to the amounts of the inputs and outputs and giving proofs that the sum of the input values equals the sum of the output values, showing that the transaction is “balanced”. Thus, no transaction creates fresh money (apart from *coinbase transactions*, which create money explicitly). Confidential transactions are now implemented e.g. in Monero.³

(2) *CoinJoin* [Max13a] is the idea of merging (or *aggregating*) several transactions into one big transaction, in a way that makes it impossible to associate the inputs and outputs of the original transactions. In Bitcoin, this can only be done by having the creators of the transactions interact in order to merge them before being included in the blockchain. In contrast, in MW merging can be done *a posteriori* without involving the original creators. The result is that in a MW blockchain all transactions are merged into one huge transaction, and there is no information about which inputs led to which outputs.

(3) *Transaction cut-through* [Max13b] is the idea that if a transaction spends an output (which corresponds to a “coin” in the system) txo_1 and creates txo_2 , which is then spent by another transaction creating txo_3 , then this should be equivalent to a “cut-through” transaction spending txo_1 and directly creating txo_3 . While in Bitcoin this could only be done for “unconfirmed transactions”, i.e., ones not yet included in any block, MW allows cut-through to be done after confirmation, which is what enables MW’s space-efficiency improvements. As every spent coin is

¹ An exception are recent proposals building on more speculative technology such as recursive zk-SNARKs; cf. <https://minaprotocol.com/lightweight-blockchain>

² <https://www.blockchain.com/charts/blocks-size>

³ <https://www.getmonero.org/resources/moneropedia/stealthaddress.html>

removed, the result is that the huge transaction representing a MW ledger only has inputs that are the coinbase transactions and outputs that are the unspent coins. In addition, this greatly improves user privacy, as the blockchain reveals neither the transacted amounts nor the transaction graph defining how coins are being transferred (in Bitcoin all this can be inferred from the blockchain).

The main shortcoming of Mimblewimble is that the sender and the receiver(s) of a transaction need to compute the transaction in an interactive protocol. It is thus not possible for a sender to simply transfer money to a destination address without any involvement of the owner of that address, which is the standard setting in all major cryptocurrencies.

Mimblewimble (MW) was first proposed by an anonymous author in 2016 [Jed16]. After being initially investigated by Poelstra [Poe16], a formal model and an analysis of MW were provided by Fuchsbauer, Orrù and Seurin (FOS) [FOS19] in 2019. In 2020, Burkett [Bur20] proposed an extension of Mimblewimble supporting non-interactive transactions, later refined by Yu [Yu20]. We will refer to this extension as MW-YU. In this work, we first assess the security of MW-YU [Yu20, §2.1] and describe discovered vulnerabilities. We then fix the scheme, also integrating an idea from a more recent proposal by Burkett [Bur21] and give security proofs that our scheme satisfies (an appropriate adaptation of) the rigorous FOS [FOS19] security model for *aggregate cash systems*.

MimbleWimbleCoin plans to implement the protocol by year-end 2022.⁴ A variant of our protocol is used in the *MimbleWimble Extension Blocks* (MWEB), which are now supported by Litecoin (one of the top cryptocurrencies with a market capitalization of > 4 billion USD).⁵

The Mimblewimble protocol. MW uses a group \mathbb{G} (which we denote additively) of prime order p with two generators G and H . As with confidential transactions [Max15], a *coin* is a Pedersen commitment $C = \text{Cmt}(v, q) := vH + qG$ to its value v using some randomness $q \in \mathbb{Z}_p$, together with a so-called *range proof* π guaranteeing that v is contained in some interval of admissible values. In MW, knowledge of the opening q of the commitment enables spending the coin. Similarly to Bitcoin, a *transaction* in MW contains a list of input coins $\mathbf{C} \in \mathbb{G}^n$ and output coins $\hat{\mathbf{C}} \in \mathbb{G}^{\hat{n}}$, where

$$C_i = v_i H + q_i G \text{ for } i \in [n] \text{ and } \hat{C}_i = \hat{v}_i H + \hat{q}_i G \text{ for } i \in [\hat{n}] .$$

Leaving fees and coinbase (a.k.a. *minting*) transactions aside, a transaction is *balanced* if and only if $\sum \hat{\mathbf{v}} - \sum \mathbf{v} = 0$ (where for a vector $\mathbf{v} = (v_1, \dots, v_n)$, we let $\sum \mathbf{v} := \sum_{i=1}^n v_i$). For coins as defined above, this is equivalent to

$$\sum \hat{\mathbf{C}} - \sum \mathbf{C} = (\sum \hat{\mathbf{q}} - \sum \mathbf{q}) G ,$$

a quantity called the *kernel excess* $E \in \mathbb{G}$ in MW. If the transaction is balanced, then knowledge of the openings $\hat{\mathbf{q}}$ and \mathbf{q} of all involved coins implies knowledge of the discrete logarithm $\log E$ (to base G) of the excess. Intuitively, if the producer of the transaction proves knowledge of $\log E$ then, together with the binding property of Pedersen commitments, this should guarantee that the transaction is balanced. In MW this is done by generating a signature σ under public key E , using its discrete logarithm $\sum \hat{\mathbf{q}} - \sum \mathbf{q}$ as the signing key.

FOS [FOS19] prove that when using Schnorr signatures (and assuming the range proofs are *simulation-extractable*; cf. Section 5), balancedness follows from the hardness of computing discrete logarithms in \mathbb{G} in the random-oracle model. FOS also show that as long as a user owning a coin C in the ledger keeps the opening private, no one can steal C (i.e., create a transaction that spends C).

Transactions in Mimblewimble can easily be merged non-interactively, in a similar way to CoinJoin [Max13a]. Consider two transactions $\text{tx}_1 = (\hat{\mathbf{C}}_1, \mathbf{C}_1, \boldsymbol{\pi}_1, E_1, \sigma_1)$ and $\text{tx}_2 = (\hat{\mathbf{C}}_2, \mathbf{C}_2, \boldsymbol{\pi}_2, E_2, \sigma_2)$. The *aggregate transaction* tx is defined as the concatenation (denoted by “ \parallel ”) of inputs and outputs, that is,

$$\text{tx} = (\hat{\mathbf{C}}_1 \parallel \hat{\mathbf{C}}_2, \mathbf{C}_1 \parallel \mathbf{C}_2, \boldsymbol{\pi}_1 \parallel \boldsymbol{\pi}_2, E_1 \parallel E_2, \sigma_1 \parallel \sigma_2) .$$

⁴ <https://www.mwc.mw/mimble-wimble-coin-articles/mimblewimble-non-interactive-transactions-review>

⁵ <https://blog.litecoin.org/litecoin-core-v0-212-release-282f5405aa11> and <https://twitter.com/DavidBurkett38/status/1555100039822954496>

A transaction of the form $\text{tx} = (\mathbf{C}, \hat{\mathbf{C}}, \boldsymbol{\pi}, \mathbf{E}, \boldsymbol{\sigma})$ is valid if all π 's and σ 's verify and if

$$\sum \hat{\mathbf{C}} - \sum \mathbf{C} = \sum \mathbf{E}. \quad (0)$$

Outputs in one transaction that also appear as inputs in the other cancel out in Equation (0) for their aggregation tx . They can therefore simply be removed from the input and output list (together with their range proofs), while validity of tx will be maintained. This has been called *transaction cut-through* in the literature [Max13b]. In MW, the ledger is defined as the cut-through of the aggregation of all transactions. Since every spent coin (a.k.a. “transaction output”, TXO) is removed by cut-through, the outputs in the ledger are precisely the *unspent TXOs* (UTXO), representing the current state of the ledger. FOS [FOS19] also remark that if the used signature scheme supports aggregation, then $\sigma_1 \parallel \sigma_2$ can be replaced by their aggregation to further save space. Then the only trace of a transaction whose outputs have been spent in the ledger is the value E .

Even if the lists of inputs and outputs in an aggregate transaction tx are shuffled, one can still link inputs and outputs that come from the same component transaction, because tx contains an excess value E_i that equals the difference between the sum of the outputs and inputs of the original transaction. This can be prevented by using *kernel offsets* [Dev20b], where E is replaced by $E + tG$ for a random $t \leftarrow \mathbb{Z}_p$ and t is included in the transaction; the aggregate of two transactions with (E_1, t_1) and (E_2, t_2) then contains $(E_1 \parallel E_2, t_1 + t_2)$. A consequence of kernel offsets is that, given an aggregated transaction, nothing can be deduced about which inputs and outputs belonged to the same original transaction. This is implied by our notion of *transaction privacy* (Section 6.6), which we prove our scheme to satisfy.

We stress that our analysis concerns the application layer, and we do not provide network-level privacy guarantees. Network adversaries that observe transactions being broadcast, or traffic analysis in general, constitute an entirely different problem that we consider outside the security of the Mimblewimble protocol per se. In practice, protocols like Dandelion [VfV17] and Tor⁶ can help mitigating attacks at the network level. See for instance their adoption in Grin⁷.

Non-interactive transactions. The key $\sum \hat{\mathbf{q}} - \sum \mathbf{q}$ for the signature σ contained in a transaction depends on the openings of the receiver’s and the sender’s coins. Most implementations of MW therefore create new transactions via an *interactive* protocol between sender and receiver to produce σ .⁸ FOS [FOS19] proposed a transaction protocol, where the sender creates all output coins, and so she can compute σ on her own. She then sends the receiver (through a separate private channel) the transaction along with the secret key associated to one of the outputs. The receiver creates a transaction spending this coin, merges it with the received transaction and broadcasts the aggregate transaction to the miners. The downside of this approach is that there is a window of time in which both sender and receiver can spend a coin, which can lead to deniability issues for payments.

In 2020, Yu [Yu20] posted on ePrint an extension of MW for achieving non-interactive transactions by adding *stealth addresses* [vS13, Tod14]. Each user has a *wallet* (or stealth address) $(A, B) \in \mathbb{G}^2$. Given a destination wallet, a sender can derive a one-time address, unique for every transaction, to which she sends the money. These one-time addresses are unlinkable to the wallet they correspond to, yet the owner of the wallet is (the only one) able to derive the secret key for it. In detail, the sender chooses a uniform element $r \leftarrow \mathbb{Z}_p$ and defines the one-time key for stealth address $(A = aG, B = bG)$ as $P = H(rA) \cdot G + B$, where H is a cryptographic hash function. Being provided $R := rG$, the owner of (A, B) can derive the secret key $\log P$ as $H(aR) + b$.

⁶ <https://www.torproject.org/>

⁷ <https://docs.grin.mw/wiki/miscellaneous/dandelion/>

⁸ For instance, in GRIN this is documented in the grin-wallet documentation: <https://raw.githubusercontent.com/mimblewimble/grin-wallet/master/doc/transaction/basic-transaction-wf.png>.

In BEAM, this is documented in the developer documentation: <https://github.com/BeamMW/beam/wiki/Cryptographic-primitives>.

Integrating stealth addresses into MW is not straightforward, as the protocol itself has no concept of addresses. Yu’s proposal [Yu20] builds on top of Burkett’s [Bur20], and after receiving feedback from the community, Yu further updated it with notes describing possible attacks and countermeasures. In essence, the idea is to extend a MW transaction output (C, π) by a one-time key P chosen by the sender, as well as an ephemeral key R that allows the receiver to compute the secret key for P (and when spending the output, a signature under signature-verification key P is required). To prevent the value P in a transaction from being modified (which would steal the output from the receiver), a signature ρ on P under key R (of which the sender knows the logarithm) is added. An output is thus of the form (C, π, R, ρ, P) .

The mechanism for letting the receiver derive the secret key for P can also be used to let the receiver derive the opening q of the commitment C (and Yu [Yu20, §2.1.1] does this by setting $q = H(H(rA)G + B)$). Note that knowing the so-called “view key” (a, B) of stealth address (aG, B) , one can derive from R both P (and thus check if the payment is for that address) and q (and thus check if C commits to a given amount).

Commonly in cryptocurrencies, when spending an output linked to a key P , the transaction is signed with the secret key of P . In Mimblewimble however, aggregation of transactions should hide which inputs and outputs come from the same component transaction, which can therefore not be signed. Yu proposes to use the logarithms of the values P in the inputs and of the values R in the outputs, denoted \hat{R} for outputs, to “authenticate” the spending by proving knowledge of the logarithm of $\sum \hat{R} - \sum P$ (similarly to how the openings of the input coins authenticate the output coins via Equation (0) in MW).

Yu proposes to simply arrange the \hat{R} values so that the above difference equals the excess E (defined as $\sum \hat{C} - \sum C$). However, this is only possible if one of the outputs \hat{C}_i goes back to the sender (who can choose the corresponding value \hat{q}_i arbitrarily); for all other coins, \hat{R}_j (together with the stealth address) defines \hat{q}_j , which defines E , for which \hat{R}_j has to be chosen, which is infeasible. We therefore modify the scheme and introduce a *stealth excess* $X := \sum \hat{R} - \sum P$, under which we add (as for E) a signature to the transaction. Our scheme then supports transactions for which all outputs are sent to destination addresses. At the time of writing, the core proposal in [Yu20, §2.1] is still affected by further issues. The ones known before our analysis are the following:

- As illustrated in [Yu20, §2.9.1], MW-YU is susceptible to a rogue-key attack [Yu20, §2.9.3]. A fix was also proposed, which requires adding one signature per transaction input, namely under the P value. The security and correctness of this proposed change are not argued further.
- Mixing NIT with non-NIT transactions, as envisaged in [Yu20], leads to correctness issues within the balance equations.⁹ No argument for why the security is preserved is given, especially w.r.t. [Yu20, Eq. ②]. (We do not consider “mixed transactions” in our scheme.)

A major drawback of Yu’s and our scheme (and MWEB) is the lack of support for transaction cut-through. Since an output is associated with a value \hat{R}_i and an input is associated with a value P_j , if an output is spent via an input in an aggregated transaction, these values do not cancel out, and removing them would thus change the stealth excess $\sum \hat{R} - \sum P$ (we discuss this in detail in Section 4.3). We note that in practice, nodes would only store and check validity of the most recent excesses and perform cut-through for coins that have been spent in the past beyond a so-called *horizon* (cf. [Bur21, §4]). As we show, cut-through enables attacks by miners, who could modify outputs of transactions (and violate our notion of *transaction-binding*, see below). This is however only relevant for *recent* transactions; once a transaction is in a block beyond the horizon, it is “protected” by the consensus mechanism (a security layer that is outside of our model).

Differences to MWEB. The variant used by Litecoin [Bur21] differs in how exactly the secrets for an output are derived from a stealth address (A, B) : In our scheme, from a Diffie-Hellman (DH)

⁹ <https://forum.mwc.mw/t/non-interactive-transaction-and-stealth-address/32>

share $R = rG$, we derive $(k, q) := H(rA)$, which defines the one-time address as $P := kG + B$ and the coin as $C := vH + qG$. Outputs in [Bur21] have an additional element K_e (in addition to K_s , which corresponds to our R) used as the Diffie-Hellman share (deriving its randomness from $\log K_s$). A symmetric key, derived from the DH-shared key, is then used to encrypt v and derive q . Our variant is arguably simpler, which facilitates our formal analysis.

Our contributions

Scheme. We propose a new protocol for non-interactive transactions, greatly inspired by Yu [Yu20] and using an idea by Burkett [Bur21] to overcome the found issues. (In Section 4 we discuss further issues that emerged after the publication of [Yu20].) A variant of our protocol is now already being used by Litecoin in its MimbleWimble Extension Blocks.

Model. We propose a security model that strengthens the one from [FOS19], which did not protect against a malleability attack by miners and had weak privacy guarantees (discussed below). We only consider non-interactive transactions, which greatly simplifies the security notions. We define security experiments that capture the following attacks:

- (i) creating money other than via coinbase transactions (*inflation resistance*)
- (ii) stealing someone else’s output in the ledger (*theft resistance*)
- (iii) changing an output of a transaction not yet included in the ledger (*transaction-binding*)
- (iv) learning anything about the transacted amounts, the destination addresses or the relations of the inputs and outputs in an aggregated transaction (*transaction privacy*)

Inflation resistance and *theft resistance* are straightforward adaptations of the notions from [FOS19, Def. 10 and 11], which become simpler for non-interactive transactions. *Transaction privacy* is stronger than FOS’s privacy notion [FOS19, Def. 12], which only guarantees that amounts are hidden (FOS’s scheme does not use kernel offsets, which means one can “disaggregate” transactions). To concisely capture all anonymity, privacy and confidentiality guarantees, we define a simulation-based notion requiring that a transaction can be simulated without knowledge of any information that transactions are supposed to hide.

We introduce *transaction-binding*, a notion that protects users against malicious miners by guaranteeing that no outputs can be removed from a transaction. In particular, after a transaction tx that spends some output txo was broadcast, no one can create a transaction tx^* that spends txo but does *not* include *all* outputs of tx . We note that *theft resistance* [FOS19, Def. 11] (which deals with interactive transactions) only guarantees the following: a user that engages in a protocol with the adversary spending coins \mathbf{C} and creating change \mathbf{C}' for herself is guaranteed that \mathbf{C}' are included in the ledger as soon as *any* of the coins from \mathbf{C} is spent.

Proof. We prove the security of our protocol by following the provable-security methodology and giving security reductions of the different security notions to standard computational hardness assumptions in idealized models.

In our security proofs, we assume that the *discrete logarithm problem* is hard in the underlying group \mathbb{G} and for transaction privacy we additionally make the *decisional Diffie-Hellman (DDH) assumption*. Our main building block is a *zero-knowledge* proof system for proving knowledge of discrete logarithms that satisfies *strong simulation-extractability* (defined in Section 5).

We show that the Schnorr signature scheme, and a variant thereof, which we use to improve efficiency of our scheme, satisfy these notions in the combination of the *algebraic group model* [FKL18] and the *random oracle model* without making any computational assumptions. Finally, we assume that the used range proofs are merely proofs of knowledge¹⁰ and do not require that they are simulation-extractable as in previous analyses [FOS19].

¹⁰ This is in some sense minimal, since for Pedersen commitments the language (see Section 2) is trivial; cf. Section 6.1.

2 Preliminaries

We let ε denote the empty string and $[a]$ the set $\{1, \dots, a\}$ (for $a \in \mathbb{N}$). We assume the existence of a group \mathbb{G} of prime order p and two “nothing-up-my-sleeve” generators $G, H \in \mathbb{G}$ (that is, the discrete logarithm of H to base G is not known to anyone). The length of the prime p is the security parameter λ . (A typical choice could be the group `Secp256k1` and hence $\lambda = 256$.) For $X \in \mathbb{G}$, we let $\log X$ denote the discrete logarithm of X to base G , that is, $\log X = x$ with $X = xG$.

Proofs of possession. We consider a cryptographic hash function which we model as a random oracle and denote by $H(\cdot)$. We use (key-prefixed) Schnorr signatures (see also [Figure 2](#), page 15), which are unforgeable under the DL assumption in the random oracle model (ROM), and whose security has been extensively studied in the literature [[PS00](#), [Seu12](#)]. We interpret Schnorr signatures as (zero-knowledge) proofs of knowledge of the secret key, that is, if $X = xG \in \mathbb{G}$ is the public key then a Schnorr signature is a proof of knowledge of $x = \log X$.

We generalize this to a proof of knowledge of two logarithms that has the same size as a Schnorr signature. Interpreting knowledge of $\log X$ as “possessing” X , we call the proof system PoP for “proof of possession”. More formally, PoP is a proof system for the following NP-relation (whose statements contains a part m , sometimes called a “tag”), defined w.r.t. a group description (p, \mathbb{G}, G) :

$$\{((X, Y, m), (x, y)) : X = xG \wedge Y = yG \wedge m \in \{0, 1\}^*\} .$$

A proof for a statement $(X, Y, m) \in \mathbb{G}^2 \times \{0, 1\}^*$ is computed via PoP.P using the witness (x, y) by picking a uniform $r \leftarrow \$ \mathbb{Z}_p$, defining $R := rG$, computing $(c, d) := H(X, Y, m, R)$ and returning a proof $(R, s) \in \mathbb{G} \times \mathbb{Z}_p$ with $s := r + c \cdot x + d \cdot y \pmod p$. The verification algorithm PoP.V($(X, Y, m), (R, s)$) computes $(c, d) := H(X, Y, m, R)$ and checks whether $s \cdot G = R + c \cdot X + d \cdot Y$ (see also [Figure 3](#), page 16).

We also use PoP to prove knowledge of a witness $x \in \mathbb{Z}_p$ for statements $(X = xG, m)$ by performing Schnorr signing and verification (defined like above but setting $y := 0$). A proof of possession of X with tag m is thus a Schnorr signature on m under public key X . We use PoP proofs for different *types* of values (proofs ψ for (P, D) , proofs σ for excesses (E, X) and proofs ρ for R). We assume that these are “domain-separated”, which can easily be achieved by including the type in the tag of the statement. We also assume that random oracles H used elsewhere in the scheme (e.g. to derive the value q) are domain-separated from H used for PoP.

In [Section 5](#) we show that in the *algebraic group model* [[FKL18](#)] combined with the ROM, the proof system PoP is a *strongly simulation-sound zero-knowledge proof of knowledge* of logarithms, a property that will be central in the security analysis of our protocol. *Zero-knowledge* means that there exists a simulator (which here controls the random oracle) that can simulate proofs for any statements without being given a witness that are indistinguishable from honestly generated proofs. *Proof of knowledge* (PoK) means that there exists an extractor that from any prover (which here is assumed to be algebraic; see [Section 5](#)) that produces a proof ψ for a statement (X, m) (or (X, Y, m)) can extract the witness $\log X$ (or $(\log X, \log Y)$). Proofs are *simulation-sound* (also called *simulation-extractable* (SE) for PoKs) if a witness can be extracted from a prover even if the prover can obtain simulated proofs ψ_i for statements (X_i, m_i) of its choice (except the one it is proving). For *strong* SE the only restriction is that the prover’s pair $((X, m), \psi)$ must be different from all query/response pairs $((X_i, m_i), \psi_i)$.

Pedersen commitments. We employ Pedersen commitments, which are homomorphic w.r.t. the committed values and the used randomness. A value $v \in \mathbb{Z}_p$ is committed by sampling $q \leftarrow \$ \mathbb{Z}_p$ and setting

$$C = \text{Cmt}(v, r) := vH + qG .$$

A commitment is opened by sending v and q and verified by checking $C \stackrel{?}{=} vH + qG$. Pedersen commitments are perfectly hiding (i.e., no information about the value is leaked) and computationally binding (i.e., under the discrete logarithm (DL) assumption, no adversary can change its mind about a committed value, that is, find a commitment and two openings $(v, q), (v', q')$ with $v \neq v'$ to it).

Range proofs. We require a proof system for statements on commitments, namely for the NP language defined by the following relation asserting that a committed value is contained in an admissible interval:

$$\{(C, (v, r)) : C = \text{Cmt}(v, r) \wedge v \in [0, v_{\max}]\}$$

We assume a zero-knowledge proof-of-knowledge system RaP for the above language. (Note that we do not assume RaP to be simulation-sound [FOS19, Def. 8], whereas FOS required this in their proof of theft resistance [FOS19, Theorem 14]; our scheme could thus be potentially instantiated with a more efficient range proof system.) We denote the prover algorithm by $\pi \leftarrow \text{RaP.P}(C, (v, r))$ and the verifier by $b \leftarrow \text{RaP.V}(C, \pi)$. A typical choice of proof system for RaP are Bulletproofs [BBB⁺18], which do not introduce any new trust assumption (as its parameters are random group elements, as for Pedersen commitments, and it is secure under DL in the ROM).

3 Proposal for MW with non-interactive transactions

We start with presenting the scheme and then discuss the rationale for our design choices, such as adding *stealth excesses* (Section 4.1), *doubling keys* (which prevent a sub-exponential-time attack; Section 4.2) and *epochs* (Section 4.4).

3.1 Data structures

The extension of Mumblewimble to non-interactive transactions introduces the notion of addresses.

A **stealth address** is a pair $(A = aG, B = bG) \in \mathbb{G}^2$, for which we call $(a, B) \in \mathbb{Z}_p \times \mathbb{G}$ the **view key** and $(a, b) \in \mathbb{Z}_p^2$ the **spend key**.

A **transaction** in MW-NIT is composed of (see also Figure 1):

- A list of **outputs**: tuples of the form $txo = (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi})$, each implicitly associated to an **output address** (A, B) , composed of:
 - an **ephemeral key** $\hat{R} = \hat{r}G \in \mathbb{G}$, chosen by the sender, which defines two keys from \mathbb{Z}_p as

$$(\hat{k}, \hat{q}) := H(\hat{r}A)$$

(note that (\hat{k}, \hat{q}) can be computed from \hat{R} and the view key, since $\hat{r}A = a\hat{R}$)

- a **commitment** $\hat{C} := \hat{v}H + \hat{q}G$ to the output value \hat{v} , using randomness \hat{q}
 - a **range proof** $\hat{\pi}$ proving knowledge of an opening (v, q) of \hat{C} , with $v \in [0, v_{\max}]$
 - a **one-time output public key** $\hat{P} \in \mathbb{G}$, computed from \hat{k} as $\hat{P} := \hat{B} + \hat{k}G$ (note that the *spend* key is required to compute $\log \hat{P}$)
 - a **proof of possession** $\hat{\rho}$ of \hat{R} with tag $\hat{C} \parallel \hat{\pi} \parallel \hat{P} \parallel \hat{\chi}$ (i.e., a proof for statement $(\hat{R}, \hat{C} \parallel \hat{\pi} \parallel \hat{P} \parallel \hat{\chi})$)
 - an **epoch** $\hat{\chi}$ in which the output was created
- A list of **inputs** of the form (P, D, ψ) where
 - $P \in \mathbb{G}$ is the one-time **public key** of the transaction output being spent (each value P is only allowed once in the ledger)
 - $D \in \mathbb{G}$ is the one-time **doubling key**, chosen by the sender, that “doubles” P
 - ψ is a **proof of possession** of P and D with tag the transaction output being spent

inputs	outputs	inputs	outputs
$txo_1 \leftarrow P_1, D_1, \psi_1$	$\hat{C}_1, \hat{\pi}_1, \hat{R}_1, \hat{\rho}_1, \hat{P}_1, \chi_1$	$txo_1 \leftarrow P_1, D_1, \psi_1$	$\hat{C}_1, \hat{\pi}_1, \hat{R}_1, \hat{\rho}_1, \hat{P}_1, \chi_1$
$txo_2 \leftarrow P_2, D_2, \psi_2$	$\hat{C}_2, \hat{\pi}_2, \hat{R}_2, \hat{\rho}_2, \hat{P}_2, \chi_2$	$txo_2 \leftarrow P_2, D_2, \psi_2$	$\hat{C}_2, \hat{\pi}_2, \hat{R}_2, \hat{\rho}_2, \hat{P}_2, \chi_2$
\vdots	\vdots	\vdots	\vdots
s, f, t, y, E, X, σ		$\mathbf{ct}, s, f, t, y, \mathbf{E}, \mathbf{X}, \sigma$	

Fig. 1. Left: Visualization of a (simple) MW-NIT transaction. Inputs consist of a value P_i contained in a previous transaction output, a “doubling key” D_i and a proof of possession (PoP) ψ_i of P_i and D_i . Outputs consist of a commitment \hat{C}_i to their value, an associated range proof $\hat{\pi}_i$, an ephemeral key \hat{R}_i , a signature (or PoP) $\hat{\rho}_i$ under \hat{R}_i , the one-time address \hat{P}_i and an epoch χ_i . The kernel consists of the supply s , the fee f , the offsets t and y and the PoP σ of E and X . The excess E and the stealth excess X can be computed as in (1) and (2). **Right:** Visualization of an aggregated MW-NIT transaction. It additionally contains a cut-through list \mathbf{ct} , and lists of excesses, stealth excesses and corresponding PoPs.

– The **kernel**, which is composed of:

- the **supply** $s \in [0, v_{\max}]$, indicating the amount of money created in the transaction
- the **fee** $f \in [0, v_{\max}]$, indicating the fee paid for the current transaction
- the **offset** $t \in \mathbb{Z}_p$
- the **excess** $E \in \mathbb{G}$, defined as the difference between the commitments in the outputs (including the fee) and the inputs (including the supply), shifted by the offset. If C_i is the i -th input commitment, that is, the value contained in the output in which P_i appears, then

$$E := \sum \hat{C} + fH - \sum C - sH - tG, \quad (1)$$

which can be seen as $E := E' - tG$ in terms of the *true excess* $E' := \sum \hat{C} + fH - \sum C - sH$

- the **stealth offset** $y \in \mathbb{Z}_p$
- the **stealth excess** $X \in \mathbb{G}$, defined as the difference between the ephemeral keys \hat{R}_i from the outputs and the doubling one-time keys D_i from the inputs, shifted by the stealth offset y

$$X := \sum \hat{R} - \sum D - yG \quad (2)$$

- a **proof of possession** σ of E and X (with empty tag ε)

A (simple, i.e., non-aggregated; see below) transaction is thus of the form:

$$\mathbf{tx} = ((P, D, \psi), (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi}), (s, f, t, y, \sigma))$$

3.2 Transaction creation

Consider a transaction output $txo = (C, \pi, R, \rho, P, \chi)$ spent to an address (A', B') .

- Given the corresponding view key (a', B') , one can compute the shared keys k and q (the opening for the commitment C) as

$$(k, q) := H(a'R).$$

- Given the corresponding spend key (a', b') , one can compute the secret key for P as $p := b' + k$.

To create a transaction that in epoch $\hat{\chi}$ spends transaction outputs txo_i of values v_i with one-time keys P_i , for $i \in [n]$, and creates outputs of values $\{\hat{v}_i\}_{i \in [\hat{n}]}$ for destination addresses $\{(A_i, B_i)\}_{i \in [\hat{n}]}$, creating an amount s of new money and paying f in fees so that $\hat{v}_1, \dots, \hat{v}_{\hat{n}}, s, f \in [0, v_{\max}]$ and $\sum \hat{\mathbf{v}} + f = \sum \mathbf{v} + s$ (if this is not the case, then abort), do the following:

- for each input index $i \in [n]$:
 - compute all values q_i and $p_i := \log P_i$, for $i \in [n]$, as described above

- if $C_i \neq v_i H + q_i G$ (with C_i from txo_i) or if $P_i \neq p_i G$ then abort
- sample a random $d_i \leftarrow \mathbb{Z}_p$ and define $D_i := d_i G$
- compute a proof of possession

$$\psi_i \leftarrow \text{PoP.P}((P_i, D_i, txo_i), (p_i, d_i))$$

– for each output index $i \in [\hat{n}]$:

- sample a random $\hat{r}_i \leftarrow \mathbb{Z}_p$ and define ephemeral key $\hat{R}_i := \hat{r}_i G$
- compute the shared secrets for the destination address (A_i, B_i)

$$(\hat{k}_i, \hat{q}_i) := \text{H}(\hat{r}_i A_i) \quad (3)$$

and from them compute the output commitment and the one-time key

$$\hat{C}_i := \hat{v}_i H + \hat{q}_i G \quad (4)$$

$$\hat{P}_i := \hat{B}_i + \hat{k}_i G \quad (5)$$

- compute a range proof $\hat{\pi}_i \leftarrow \text{RaP.P}(\hat{C}_i, (\hat{v}_i, \hat{q}_i))$
- compute a proof of possession of the ephemeral key

$$\hat{\rho}_i \leftarrow \text{PoP.P}((\hat{R}_i, \hat{C}_i \parallel \hat{\pi}_i \parallel \hat{P}_i \parallel \hat{\chi}), \hat{r}_i) \quad (6)$$

– sample a random $t \leftarrow \mathbb{Z}_p$ and compute

$$e := \sum \hat{q} - \sum \mathbf{q} - t = \log E$$

with E as in (1)

– sample a random $y \leftarrow \mathbb{Z}_p$ and compute

$$x := \sum \hat{r} - \sum \mathbf{d} - y = \log X$$

with X as in (2).

– compute a proof of possession of E and X with empty tag: $\sigma \leftarrow \text{PoP.P}((E, X, \varepsilon), (e, x))$

The final transaction is

$$\text{tx} := ((P_i, D_i, \psi_i)_{i \in [n]}, (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi})_{i \in [\hat{n}]}, (s, f, t, y, \sigma)) \quad (7)$$

3.3 Transaction aggregation

Aggregate transactions are essentially concatenations of the composing transactions. In contrast to Jedusor's [Jed16] and FOS' [FOS19] protocols, *MW-NIT does not perform any cut-through*, as this is insecure, as we show in Section 4.3. Outputs of one transaction that are spent as inputs of another one in the aggregation are therefore kept in a *cut-through list* \mathbf{ct} , which stores the concatenation of the output and the input spending the latter.

While simple transactions do not (need to) contain the (stealth) excesses (displayed in light gray in Figure 1), aggregate transactions (also displayed in Figure 1) contain lists of excesses \mathbf{E} , stealth excesses \mathbf{X} and associated proofs $\boldsymbol{\sigma}$. An aggregated transaction is thus of the form

$$\text{tx} = ((P_i, D_i, \psi_i)_{i \in [n]}, (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi})_{i \in [\hat{n}]}, (\mathbf{ct}, s, f, t, y, (E_i, X_i, \sigma_i)_{i \in [\bar{n}]}) \quad (8)$$

where $\mathbf{ct} := (C'_i, \pi'_i, R'_i, \rho'_i, P'_i, \chi'_i, P'_i, D'_i, \psi'_i)_{i \in [n']}$. A simple transaction (7) can be cast as (8) by setting $\mathbf{ct} := ()$, $\bar{n} = 1$ and computing E and X as in Equations (1) and (2). Given transactions tx_1 and tx_2 , assuming w.l.o.g. that they are of the form (8), their aggregation tx is computed as follows:

- define \mathbf{txi} as the concatenation of the inputs of tx_1 and tx_2 , and \mathbf{txo} as the concatenation of their outputs, \mathbf{ct} as the concatenation of their cut-through lists and \mathbf{ker} as the concatenation of their lists of excesses, stealth excesses and associated signatures σ .

- if the same value P appears in two entries of $\mathbf{txi} \parallel \mathbf{ct}$, or if the same value P appears in two entries of $\mathbf{txo} \parallel \mathbf{ct}$, then abort
- if a P value appears in an entry (P_i, D_i, ψ_i) of \mathbf{txi} and in an entry txo_j of \mathbf{txo} , do the following: if $\text{PoP.V}((P_i, D_i, txo_j), \psi_i) = \mathbf{true}$ then remove (P_i, D_i, ψ_i) from \mathbf{txi} , remove txo_j from \mathbf{txo} and include $txo_j \parallel (P_i, D_i, \psi_i)$ in \mathbf{ct} ; else abort.
- sort each list \mathbf{txi} , \mathbf{txo} , \mathbf{ct} and \mathbf{ker} lexicographically (required to hide which inputs and outputs come from which transaction)
- compute the aggregated supply, fee, offset, and stealth offset (from the supplies s_i , etc., of \mathbf{tx}_i)

$$s := s_1 + s_2 \quad f := f_1 + f_2 \quad t := t_1 + t_2 \quad y := y_1 + y_2$$

- return $\mathbf{tx} := (\mathbf{txi}, \mathbf{txo}, (\mathbf{ct}, s, f, t, y, \mathbf{ker}))$.

3.4 Output verification

We define when a view key (a, B) accepts a transaction output. Given an amount v and an output $txo = (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi})$, compute $(k, q) := \text{H}(a\hat{R})$, and accept txo if it has not been previously received in epoch $\hat{\chi}$ and if

$$\hat{C} = vH + qG \quad \text{and} \quad \hat{P} = B + kG .$$

3.5 Transaction verification

Simple transactions. A transaction

$$\mathbf{tx} = ((P_i, D_i, \psi_i)_{i \in [n]}, (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi}_i)_{i \in [\hat{n}]}, (s, f, t, y, \sigma)) ,$$

is valid w.r.t. a list of (previous) outputs \mathbf{txo} with $txo_i = (C_i, \pi_i, R_i, \rho_i, P_i, \chi_i)$, for $i \in [n]$, if

- for all $i \in [n]$, the value P_i in \mathbf{tx} is unique and the value P_i in txo_i are the same
- all input proofs are valid for all $i \in [n]$: $\text{PoP.V}((P_i, D_i, txo_i), \psi_i) = \mathbf{true}$
- all range proofs are valid for all $i \in [\hat{n}]$: $\text{RaP.V}(\hat{C}_i, \hat{\pi}) = \mathbf{true}$
- all proofs of possession of \hat{R} are valid for all $i \in [\hat{n}]$: $\text{PoP.V}(\hat{R}_i, \hat{C}_i \parallel \hat{\pi}_i \parallel \hat{P}_i \parallel \hat{\chi}_i, \hat{\rho}_i) = \mathbf{true}$
- the excess proof of possession is valid $\text{PoP.V}((E, X, \varepsilon), \sigma) = \mathbf{true}$, for

$$E := \sum \hat{C} - \sum C + (f - s)H - tG \tag{9}$$

$$X := \sum \hat{R} - \sum D - yG \tag{10}$$

Aggregate transactions. An aggregate transaction

$$\mathbf{tx} = (\mathbf{txi}, \widehat{\mathbf{txo}}, (\mathbf{ct}, s, f, t, y, \mathbf{E}, \mathbf{X}, \sigma))$$

as in (8), with $\mathbf{txi} = (P_i, D_i, \psi_i)_{i \in [n]}$, $\widehat{\mathbf{txo}} = (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi}_i)_{i \in [\hat{n}]}$, $\mathbf{ct} = (txo'_i \parallel txi'_i)_{i \in [n']}$ is verified w.r.t. previous outputs \mathbf{txo} as follows:

Check that \mathbf{txi} , $\widehat{\mathbf{txo}}$ and \mathbf{ct} are sorted lexicographically, then re-arrange the cut-through terms: set the overall outputs being spent and the inputs spending them, as well as the freshly created outputs, as

$$\mathbf{txo}^* := \mathbf{txo} \parallel \mathbf{txo}' = (C_i, \pi_i, R_i, \rho_i, P_i, \chi_i)_{i \in [n+n']}$$

$$\mathbf{txi}^* := \mathbf{txi} \parallel \mathbf{txi}' = (P_i, D_i, \psi_i)_{i \in [n+n']}$$

$$\widehat{\mathbf{txo}}^* := \widehat{\mathbf{txo}} \parallel \mathbf{txo}' = (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi}_i)_{i \in [\hat{n}+n']}$$

(Note that the last n' elements of \mathbf{txo}^* and $\widehat{\mathbf{txo}}^*$ are equal but denoted differently). Verify the transaction $(\mathbf{txi}^*, \widehat{\mathbf{txo}}^*, (s, f, t, y, \mathbf{E}, \mathbf{X}, \sigma))$ w.r.t. \mathbf{txo}^* : check (i)–(iv) as for simple transactions, and the following instead of (v):

(v') for all $i \in [\bar{n}] : \text{PoP.V}((E_i, X_i, \varepsilon), \sigma_i) = \mathbf{true}$
(vi') additionally, the following “balance equations” are checked:

$$\sum E = \sum \hat{C} - \sum C + (f - s)H - tG \quad (11)$$

$$\sum X = \sum \hat{R} - \sum D - yG \quad (12)$$

(Note that in (11) it suffices to sum the \hat{C} only up to \hat{n} and the C only up to n , as the remaining terms from \mathbf{ct} cancel out.)

3.6 Inclusion of transactions in the ledger

A ledger Λ is simply an aggregated transaction of the form (8) that has no inputs (as the inputs of any transaction added to Λ must spend existing outputs, these are moved to the cut-through list).

A transaction \mathbf{tx} of the form (7) or (8) is included in Λ by aggregating (as defined in Section 3.3) Λ and \mathbf{tx} to Λ' , checking that Λ' has no inputs (thus all inputs of \mathbf{tx} were spent/cut through) and checking validity (as defined in Section 3.5) of Λ' w.r.t. an empty list \mathbf{txo} . If any of the checks fail, return \perp , otherwise Λ' . (If Λ is known to be valid, it suffices to identify, for every input (P_j, D_j, ψ_j) of \mathbf{tx} , the ledger output \mathbf{txo}_{i_j} containing P_j , and check validity of \mathbf{tx} is w.r.t. $\mathbf{txo}_{i_1}, \dots, \mathbf{txo}_{i_n}$.)

4 Fallacies in the initial proposal

We discuss the main issues found with Yu’s scheme, which motivated the design choices for our scheme in Section 3. Originally [Yu20, §2.2.2], transactions did not include the values D_i , nor the stealth excess X with the respective offset y . Instead, a valid transaction had to satisfy

$$E + tG = \sum \hat{R} - \sum P \quad (13)$$

(in our notation) [Yu20, Eq. ②] instead of Equation (10), and ψ and σ only proved knowledge of the discrete logarithms of P and E , respectively.

4.1 Correctness

Equation (13) can be made true if one of the outputs, say the i -th, goes back to the creator of the transaction (e.g., because it is a “change output”). She can then set $\hat{R}_i := E + tG + \sum P - \sum_{j \neq i} \hat{R}_j$ (for which she knows $\log \hat{R}_i$) and sample q_i uniformly. However, it is infeasible to create a transaction whose outputs are all linked to destination addresses, e.g., a transaction with a single output: \hat{R} (together with the address) determines the coin opening q , which defines the value E ; but (re-)defining \hat{R} so that (13) holds would lead to a new value E . (In Yu’s notation [Yu20, Eq. ②], the value r_o depends on q , which in turn is computed from r_o .)

In order not to restrict the format of transactions (and because it allows us to prove the scheme secure), we introduced stealth excesses X , which along with the proof σ accounts for the “excess” in stealth addresses. We also introduce a *stealth offset* to preserve privacy of aggregated transactions.

4.2 The Feed-Me attack

It turns out that merely adding a stealth excess leads to an attack (against *transaction-binding*, see Section 6.5), which was first found by @south_lagoon77, alias kurt.¹¹ Consider the scheme in Section 3 but without any D values, which are replaced by the corresponding P values in the equations; in particular, the balance equation for one-time keys is

$$X = \sum \hat{R} - \sum P - y^*G \quad (10^*)$$

¹¹ See: <https://twitter.com/davidburkett38/status/1466460568525713413>

instead of Equation (10).

To explain the attack, it suffices to focus on non-aggregated transactions. Consider Alice, an honest user with address (aG, B) , that creates two transactions tx_1 and tx_2 , both spending one output txo_1 and txo_2 , resp., and creating one output each, transferring v_1 to Bob and $v_2 > v_1$ to Charlie (all supplies and fees are 0):

$$\text{txo}_1 = (C_1, \dots) \leftarrow \boxed{P_1, \psi_1 \mid \dots, \hat{R}_1, \hat{\rho}_1, \dots}_{t_1, y_1, E_1, X_1, \dots} \quad (\text{tx}_1)$$

$$\text{txo}_2 = (C_2, \dots) \leftarrow \boxed{P_2, \psi_2 \mid \dots, \hat{R}_2, \hat{\rho}_2, \dots}_{t_2, y_2, E_2, X_2, \dots} \quad (\text{tx}_2)$$

Both transaction are broadcast to the miners. A malicious miner can now forge a new two-output transaction tx^* , transferring the amount $v_2 - v_1$ to himself, as follows:

$$\text{txo}_2 \leftarrow \boxed{P_2, \psi_2 \mid \begin{array}{c} \dots, \hat{R}_1, \hat{\rho}_1, \dots \\ C^*, \pi^*, R^*, \rho^*, \dots \end{array}}_{t_1, y^*, E_1, X_1, \dots} \quad (\text{tx}^*)$$

It combines the input of tx_2 and the output and the excesses from tx_1 . A second output is computed “honestly” by choosing $r^* \leftarrow_{\$} \mathbb{Z}_p$, setting $R^* = r^*G$ and signing any value P^* (knowing $\log P^*$) as ρ^* . The miner also creates the corresponding coin C^* , so that tx^* satisfies Equation (11), by setting $C^* := \text{Cmt}(v_2 - v_1, q_2 - q_1)$, where q_1 and q_2 are the openings of C_1 and C_2 (which the miner can either obtain by knowing Alice’s view key, or *by having sent the outputs txo_1 and txo_2 to Alice in the first place.*) Finally, the miner needs to compute y^* so that tx^* satisfies Equation (10*), that is,

$$y^*G = R^* + \hat{R}_1 - P_2 - X_1 .$$

By validity of tx_1 , again from (10*), we get $0 = -\hat{R}_1 + P_1 + y_1G + X_1$ and by adding the two equations:

$$y^*G = R^* + P_1 - P_2 + y_1G . \quad (14)$$

As per Equation (5), P_1 and P_2 are defined as $P_i = B + k_iG$, where k_i is obtained as in Equation (3). The crucial observation now is that if the miner knows the values k_1 and k_2 (which it can either obtain from Alice’s view key or by being the creator txo_1 and txo_2), then it knows the discrete logarithm of $P_1 - P_2 = (k_1 - k_2)G$, as the value B from Alice’s address *cancels out*. The miner can thus compute y^* satisfying (14), thus completing the forged transaction tx^* that transfers the value $v_2 - v_1$ to the miner’s one-time key P^* .

Fixes. To prevent this attack, our first fix was to derive the one-time keys multiplicatively, that is, setting $P_i = k_iB$ in Equation (5). The term $P_1 - P_2$ then becomes $(k_1 - k_2)B$, which is non-zero with overwhelming probability and therefore it becomes hard to compute y^* . More generally, we showed that as long as the adversary cannot find distinct hash function outputs $k_{1,1}, \dots, k_{1,n_1}, k_{2,1}, \dots, k_{2,n_2}$ in Equation (3) so that

$$\sum \mathbf{k}_1 - \sum \mathbf{k}_2 = 0 , \quad (15)$$

this variant of the scheme (without the D values) satisfies transaction binding.

However, Wagner’s *k-list tree algorithm* [Wag02] can be used to find such values in sub-exponential time for sufficiently large n_1 and n_2 . To be protected against active adversaries ready to invest substantial computing power, a user would therefore need to limit the number of her pending transactions at any point in time (which could degrade scalability of the system). To overcome this downside, we follow David Burkett’s approach [Bur21] and introduce an additional group element D in every transaction input, which replaces P in the balance equation (10*), yielding (12). Since the D_i ’s are chosen by the creator of the transaction (whereas the P_i ’s are chosen by the previous spender, who in *Feed-Me* types of attacks is malicious), the values corresponding to the $k_{i,j}$ above are random, and thus the probability that (15) holds is negligible.

Since the D values are chosen by the honest user, the adversary in the Feed-Me attack does not know $\log(D_1 - D_2)$ (after replacing P values by D values in (14)), so we just reverted to the original format $P := B + kG$ for one-time keys and prove this variant satisfies transaction-binding in Section 6.5.

4.3 On transaction cut-through

Suppose that, in an aggregate transaction, an output $(C, \pi, R, \rho, P, \chi)$ of one transaction is spent as input (P, D, ψ) of another transaction. One may wonder whether, as with original MW, cut-through can be applied, that is, remove the spent output and the input referring to it from the aggregate transaction.

While validity of the coin-balance equation (11) would be maintained, this is not the case for the “address equation” (12). One may thus consider (as Yu does [Yu20, §2.1.1] for (sufficiently old parts of) the ledger) adding a value Z defined as the sum of all removed \hat{R} values minus all removed D values to the aggregated transaction. The check in Equation (12) would be replaced by $\sum \hat{R} - \sum D + Z = \sum X + yG$, where the sums are only over the indices that have not been removed in the outputs ($\sum \hat{R}$) and the inputs ($\sum D$).

However, since Z is not bound to anything, this scheme would be insecure. Consider a miner that collects transactions and aggregates them. Then she simply replaces one of the remaining \hat{R} values by a value of which she knows the discrete logarithm, puts a new \hat{P} value, produces a corresponding proof $\hat{\rho}$ and defines Z as $\sum X + yG + \sum D$ minus the sum of the \hat{R} values including her own. The result is a valid transaction of which the miner now owns one of the outputs (assuming the miner knows the view key of the stealth address it stole the coin from; otherwise it made the coin unspendable by its owner).

We suspect that Yu assumed all R and P values remain for each (possibly aggregated) transaction when included in the blockchain, since in [Yu20, Eq. ③], it says

$$SUM(R - P')_{\text{spent at height}}$$

which suggests that all these values need to be present. In addition, we note that simply removing cut-through inputs or outputs would make Equations ③ and ④ incorrect.

Example. We illustrate the attack against a scheme that allows cut-through and includes a value Z to compensate for the removed terms. Consider two 1-input/1-output transactions tx_2 and tx_3 where tx_2 spends some output tx_{01} , creating one output, which is then spent by tx_3 (assume that all supplies and fees are 0).

$$\begin{aligned} \text{tx}_{01} = (C_1, \dots) &\leftarrow \boxed{P_1, D_1, \psi_1 \mid C_2, \pi_2, R_2, \rho_2, P_2, \chi_2}_{t_2, y_2, E_2, X_2, \sigma_2} && (\text{tx}_2) \\ &\leftarrow \boxed{P_2, D_2, \psi_2 \mid C_3, \pi_3, R_3, \rho_3, P_3, \chi_3}_{t_3, y_3, E_3, X_3, \sigma_3} && (\text{tx}_3) \end{aligned}$$

(Note that we have $\text{PoP.V}((P_2, D_2, \text{tx}_2.\text{out}), \psi_2) = \text{true}$). Suppose tx_2 and tx_3 could be merged as

$$\text{tx}_{01} \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_3, \pi_3, R_3, \rho_3, P_3, \chi_3}_{t_2+t_3, y_2+y_3, (E_2, E_3), (X_2, X_3), (\sigma_2, \sigma_3), Z}$$

which is valid if $\psi_1, \pi_3, \rho_3, \sigma_2$ and σ_3 are valid and the following holds:

$$\begin{aligned} C_3 - C_1 &= E_2 + E_3 + (t_2 + t_3)G \\ R_3 - D_1 + Z &= X_2 + X_3 + (y_2 + y_3)G \end{aligned}$$

Then a miner could simply choose r^*, p^*, χ_3^* set $R_3^* := r^*G, P_3^* := p^*G$, create ρ_3^* honestly, define $Z^* := X_2 + X_3 + (y_2 + y_3)G - R_3^* + D_1$ and create the (valid!) transaction

$$\text{tx}_{01} \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_3, \pi_3, R_3^*, \rho_3, P_3^*, \chi_3^*}_{t_2+t_3, y_2+y_3, (E_2, E_3), (X_2, X_3), (\sigma_2, \sigma_3), Z^*}$$

for which she knows the temporary spending key p^* .

On keeping ρ and ψ . One may wonder then if it is possible to keep the R and D values but elide the proofs ρ and ψ , or at least one of them. Again, each of these removals would lead to an attack. We start with considering **removing ρ but keeping ψ** , that is, an aggregated transaction in the example above looks as follows:

$$txo_1 = (C_1, \dots) \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_3, \pi_3, R_3, \rho_3, P_3, \chi_3}_{\mathbf{ct}, t_2+t_3, y_2+y_3, (E_2, E_3), (X_2, X_3), (\sigma_2, \sigma_3)}$$

where $\mathbf{ct} = (R_2, P_2, D_2, \psi_2)$ and ψ_2 is a valid proof for P_2, D_2 (note that the tag for proofs ψ must change to account for removed information).

Intuitively, not having ρ_2 means that P_2 is not bound to R_2 anymore, which the following attack leverages, where given tx_2 and tx_3 , the miner replaces tx_3 by a transaction it owns: the miner chooses p_2^*, d_2^* and r_3^* , sets $D_2^* := d_2^*G$ and $R_3^* := r_3^*G$, computes ψ_2^* and ρ_3^* honestly and sets $X_2^* := (r_3^* - d_2^* - y_2)G$ and computes σ_2^* honestly. It also chooses p_3^* , sets $P_3^* := p_3^*G$, and creates a proof ρ_3^* on it using r_3^* , and computes X_3^* and σ_3^* honestly. Then the following is a valid transaction for which the miner knows the key to spend the output

$$txo_1 = (C_1, \dots) \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_3, \pi_3, R_3^*, \rho_3^*, P_3^*, \chi_3^*}_{\mathbf{ct}^*, t_2+t_3, y_2+y_3, (E_2, E_3), (X_2^*, X_3^*), (\sigma_2^*, \sigma_3^*)}$$

where $\mathbf{ct}^* = (R_2, P_2^*, D_2^*, \psi_2^*)$.

Finally, we show that **removing ψ but keeping ρ** also leads to attacks, similarly to the rogue-key attack observed in [Yu20, §2.9.3]. In this scenario, the above aggregated transaction would look as follows:

$$txo_1 = (C_1, \dots) \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_3, \pi_3, R_3, \rho_3, P_3, \chi_3}_{\mathbf{ct}, t_2+t_3, y_2+y_3, (E_2, E_3), (X_2, X_3), (\sigma_2, \sigma_3)}$$

where $\mathbf{ct} = (C_2, \pi_2, R_2, \rho_2, P_2, \chi_2, D_2)$; note that now C_2, π_2 and χ_2 need to be kept for ρ_2 to be verifiable. Consider an honest transaction tx_2

$$txo_1 = (C_1, \dots) \leftarrow \boxed{P_1, D_1, \psi_1 \mid C_2, \pi_2, R_2, \rho_2, P_2, \chi_2}_{t_2, y_2, E_2, X_2, \sigma_2}$$

A miner can steal the output as follows (again, assuming it knows the view key of its holder and thus the opening of C_2).

The miner computes a fresh output $(C_3, \pi_3, R_3, \rho_3, P_3, \chi_3)$ honestly (i.e., knowing the logarithms of R_3 and P_3), picks random R_2^* and X_3 (knowing the corresponding logarithms) and sets $D_2^* := R_2^* + R_3 - D_1 - X_2 - X_3 - y_2G$. It signs P_2 (as well as C_2 , and π_2) under R_2^* as ρ_2^* and produces σ_3 under E_3, X_3 . It then publishes the transaction

$$txo_1 = (C_1, \dots) \leftarrow \boxed{P_1, D_1 \mid C_3, \pi_3, R_3, \rho_3, P_3, \chi_3}_{\mathbf{ct}^*, t_2+t_3, y_2, (E_2, E_3), (X_2, X_3), (\sigma_2, \sigma_3)}$$

where $\mathbf{ct}^* := (C_2, \pi_2, R_2^*, \rho_2^*, P_2^*, \chi_2, D_2^*)$. Note that the transaction is valid, since we have

$$R_2^* + R_3 - D_1 - D_2^* = X_2 + X_3 + y_2G$$

and the miner knows the key to spend the output.

4.4 Replay attacks

Yu [Yu20, §2.9.2] explains a replay attack for MW that is a result of non-interactive transactions: the adversary pays Alice via some output txo , which Alice later spends. Then the adversary pays her again, creating the exact same output; if Alice accepts it, the adversary can replay Alice's previous spend, making her lose the money.

$\text{Sch.Setup}(1^\lambda)$ <hr/> $(p, \mathbb{G}, G) \leftarrow \text{GrGen}(1^\lambda)$ $\text{fix } H: \{0, 1\}^* \rightarrow \mathbb{Z}_p$ $\text{return } par := (p, \mathbb{G}, G, H)$	$\text{Sch.KeyGen}(par)$ <hr/> $(p, \mathbb{G}, G, H) := par; x \leftarrow_{\$} \mathbb{Z}_p; X := xG$ $sk := par \parallel x; pk := par \parallel X$ $\text{return } (sk, pk)$
$\text{Sch.Sign}(sk, m)$ <hr/> $(p, \mathbb{G}, G, H, x) := sk; r \leftarrow_{\$} \mathbb{Z}_p; R := rG$ $c := H(xG, R, m); s := r + cx \text{ mod } p$ $\text{return } \sigma := (R, s)$	$\text{Sch.Ver}(pk, m, \sigma)$ <hr/> $(p, \mathbb{G}, G, H, X) := pk; (R, s) := \sigma$ $c := H(X, R, m)$ $\text{return } (sG = R + cX)$

Fig. 2. Key-prefixed Schnorr signature scheme PoP[GrGen] based on a group generator GrGen

A simple defense against replay attacks is requiring users to store all outputs ever received and never accept the same output a second time. A more viable method is using time stamps (named *epochs* and denoted χ), which Yu introduces in order to prevent rogue-key attacks. “Each *Input* must attach its own proof for $[P_i]$, as a second proof for the coin ownership” [Yu20, §2.9.2].

While it is not specified which message is signed, it is crucial that the entire output specifically including the time stamp (and not just C) is signed. Otherwise, the above attack still works, as the adversary can change the time stamp in the replayed transaction (so the user accepts it) and recompute ρ , and send it to the user again. If the proof contained in the user’s spendings did not authenticate the time stamp (or ρ), then the previous spend would still be valid on the replayed transaction. This is why in MW-NIT we define ψ as a proof that involves the *entire* output.

Epochs. If a user only accepts outputs that correspond to the current epoch, she only needs to compare a new output to those received in the same epoch; she can therefore delete all outputs from previous epochs. The duration of an epoch is a global parameter of the system, where short time intervals minimize data storage, while larger intervals yield better privacy (as there are more transactions per epoch).

5 Simulation-extractability of Schnorr signatures

Before analyzing our scheme, we introduce and analyze its main building block. Key-prefixed Schnorr signatures, formally defined in Figure 2, can be reinterpreted as zero-knowledge proofs of knowledge of the secret key, with the statement also containing the message. To improve efficiency, we generalize this to a “batch” version that enables proving knowledge of the logarithms of two group elements, that is, proofs for the NP language defined w.r.t. a group description (p, \mathbb{G}, G) by the relation

$$\{((X, Y, m), (x, y)) : X = xG \wedge Y = yG \wedge m \in \{0, 1\}^*\}. \quad (16)$$

The proof system PoP is defined in Figure 3. We also use it to prove statements (X, m) with witness x by using standard Schnorr signatures, that is, PoP.P runs Sch.Sign and PoP.V runs Sch.Ver. The witness relation for PoP is thus the union of (16) and $\{((X, m), x) : X = xG\}$.

We show that PoP satisfies *strong simulation extractability* in the *algebraic group model* [FKL18] (see below) and the random oracle model (ideal models jointly used also in [FPS20] to show tight security of Schnorr signatures under the discrete-logarithm assumption).

Simulation-extractability. Strong simulation extractability for the above language means that from any adversary that returns a proof ψ^* for a statement (X^*, Y^*, m^*) , the witness $(\log X^*, \log Y^*)$

PoP.Setup(1^λ)	PoP.P($par, (X, Y, m), (x, y)$)	PoP.V($par, (X, Y, m), \psi$)
$(p, \mathbb{G}, G) \leftarrow \text{GrGen}(1^\lambda)$	$(p, \mathbb{G}, G, H) := par$	$(p, \mathbb{G}, G, H) := par$
fix $H: \{0, 1\}^* \rightarrow \mathbb{Z}_p^2$	$r \leftarrow \$ \mathbb{Z}_p; R := rG$	$(R, s) := \psi$
return (p, \mathbb{G}, G, H)	$(c, d) := H(X, Y, m, R)$	$(c, d) := H(X, Y, m, R)$
	$s := r + cx + dy \bmod p$	return $(sG = R + cX + dY)$
	return $\psi := (R, s)$	

Fig. 3. Zero-knowledge simulation-extractable Schnorr proof of two logarithms PoP[GrGen] based on a group generator GrGen

can be extracted; and this holds even if the adversary gets access to an oracle that on inputs (X_i, Y_i, m_i) returns simulated proofs ψ_i for these statements. The only restriction is that the returned pair $((X^*, Y^*, m^*), \psi^*)$ must be different from all query/response pairs $((X_i, Y_i, m_i), \psi_i)$. Thus, forging a fresh proof ψ^* on a queried statement is considered a break of *strong* simulation-extractability if the extractor fails to extract a witness from ψ^* . (Note that this notion is stronger than forms of related-key-attack security for signature schemes (like UNF-CRO as defined and used in [FOS19]), where the adversary can only query signatures under keys for which it knows the difference in secret keys w.r.t. the challenge key.)

The algebraic group model. In the algebraic group model (AGM) [FKL18], adversaries are assumed to return a *representation* of any group element that they return. This means that, after having received input group elements Z_1, \dots, Z_n , whenever the adversary returns a group element X , it must also return coefficients ζ_1, \dots, ζ_n so that $X = \sum \zeta_i Z_i$.

All our security proofs (except for the privacy notion) are reductions of solving the discrete logarithm (DL) problem to breaking the analyzed security notion of our scheme MW-NIT, assuming that PoP satisfies (strong) simulation-extractability (SE) in the AGM. The reduction thus receives a DL challenge Z and simulates the security game to an adversary \mathcal{A} , which we assume is algebraic. To leverage SE of PoP *in the AGM*, the reduction must construct an *algebraic* SE adversary, that is, one that accompanies each group-element output by their representations. However, the reduction can only return representations in basis (G, Z) , its own group-element inputs. In particular, the reduction will run the adversary on some group elements X_1, \dots, X_n , which it produces from its inputs G and Z in an “algebraic” way (i.e., knowing representations in basis (G, Z)). The adversary’s group-element outputs will thus be in basis X_1, \dots, X_n , which the reduction can then translate into the basis (G, Z) .

For our reductions make use of simulation extractability, we must therefore strengthen the notion and consider auxiliary inputs. In the SE game, the adversary receives, besides a description of the underlying group, with generator G , and possible proof system parameters, an “auxiliary” uniform group element Z . At the end of its execution, the algebraic adversary must accompany each group element queried to the simulation oracle and output to the challenger (in particular, group elements in the statements (X_i, Y_i, m) for which the extractor must extract the witness) by a representation in basis (G, Z) .

Security of PoP. Extending the techniques for showing tight security of Schnorr signatures in the AGM+ROM [FPS20], we show that in this model our proof system PoP is SE with auxiliary inputs.

Claim 1 *The proof system PoP in Figure 3 is strongly simulation-extractable with auxiliary group-element input in the algebraic group model and the random oracle model.*

Proof. The game is parametrized by a group (p, \mathbb{G}, G) and a random oracle H , all provided to the adversary, who also receives a random group element $Z \leftarrow \mathbb{S} \mathbb{G}$. When the adversary queries simulation of a proof for a statement $(X_i, Y_i, m_i) \in \mathbb{G}^2 \times \{0, 1\}^*$, the simulator chooses uniform $c_i, d_i, s_i \leftarrow \mathbb{S} \mathbb{Z}_p$, sets $R_i := s_i G - c_i X_i - d_i Y_i$ and programs the random oracle so that $H(X_i, Y_i, m_i, R_i) = (c_i, d_i)$. Since R_i is uniform and independent, the probability that the RO has already been defined for this value is negligible. If this happens then the simulator aborts and the adversary wins. As the adversary is algebraic, it needs to accompany its first query (X_1, Y_1, m_1) by coefficients $\alpha_1, \beta_1, \gamma_1, \delta_1$ with $X_1 = \alpha_1 G + \beta_1 Z$ and $Y_1 = \gamma_1 G + \delta_1 Z$. After receiving R_1 , its second query is accompanied by $\gamma_X^{(2)}, \zeta_X^{(2)}, \rho_X^{(2)}, \gamma_Y^{(2)}, \zeta_Y^{(2)}, \rho_Y^{(2)}$ with $X_2 = \gamma_X^{(2)} G + \zeta_X^{(2)} Z + \rho_X^{(2)} R_1$ and $Y_2 = \gamma_Y^{(2)} G + \zeta_Y^{(2)} Z + \rho_Y^{(2)} R_1$. Since $R_1 = s_1 G - c_1 X_1 - d_1 Y_1$, this yields $X_2 = \alpha_2 G + \beta_2 Z$ with

$$\alpha_2 := \gamma_X^{(2)} + \rho_X^{(2)}(s_1 - c_1 \alpha_1 - d_1 \gamma_1) \quad \text{and} \quad \beta_2 := \zeta_X^{(2)} + \rho_X^{(2)}(-c_1 \beta_1 - d_1 \delta_1),$$

and analogous values γ_2 and δ_2 with $Y_2 = \gamma_2 G + \delta_2 Z$. In general, for every query answered with (R_i, s_i) for $(c_i, d_i) = H(X_i, Y_i, m_i, R_i)$ we thus have

$$\begin{aligned} R_i &= s_i G - c_i X_i - d_i Y_i = (s_i - c_i \alpha_i - d_i \gamma_i) G - (c_i \beta_i + d_i \delta_i) Z \\ &= \Gamma_i G - \Delta_i Z \quad \text{with} \quad \Gamma_i := s_i - c_i \alpha_i - d_i \gamma_i \quad \text{and} \quad \Delta_i := c_i \beta_i + d_i \delta_i \end{aligned} \quad (17)$$

and from any representation

$$X_i = \gamma_X^{(i)} G + \zeta_X^{(i)} Z + \sum_{j=1}^{i-1} \rho_X^{(i,j)} R_j$$

we can recursively derive α_i, β_i so that $X_i = \alpha_i G + \beta_i Z$ and similarly γ_i, δ_i for $Y_i = \gamma_i G + \delta_i Z$.

Consider a proof (R^*, s^*) for a statement (X^*, Y^*, m^*) output by the adversary together with representations (α^*, β^*) for X^* and (γ^*, δ^*) for Y^* so that

$$(X^*, Y^*, m^*, R^*, s^*) \neq (X_i, Y_i, m_i, R_i, s_i) \quad \text{for all } i. \quad (18)$$

Validity means

$$R^* + c^* X^* + d^* Y^* = s^* G \quad \text{with} \quad (c^*, d^*) = H(X^*, Y^*, m^*, R^*). \quad (19)$$

Consider the point when $H(X^*, Y^*, m^*, R^*)$ gets defined. This must be during a random-oracle query by the adversary, since a successful adversary cannot have made a simulation query for (X^*, Y^*, m^*) answered with R^* : as there is only one valid value s^* , this would mean that the adversary returned the oracle's response, i.e., (18) does not hold.

Let q be the number of simulation queries made before the random-oracle query (X^*, Y^*, m^*, R^*) . Since the adversary is algebraic, it must accompany X^*, Y^* and R^* by representations $(\gamma_X, \zeta_X, \boldsymbol{\xi})$, $(\gamma_Y, \zeta_Y, \boldsymbol{v})$ and $(\gamma_R, \zeta_R, \boldsymbol{\rho})$, respectively with

$$\begin{aligned} X^* &= \gamma_X G + \zeta_X Z + \sum_{i=1}^q \xi_i R_i \stackrel{(17)}{=} \left(\gamma_X + \sum_{i=1}^q \xi_i \Gamma_i \right) G + \left(\zeta_X - \sum_{i=1}^q \xi_i \Delta_i \right) Z, \\ Y^* &= \left(\gamma_Y + \sum_{i=1}^q v_i \Gamma_i \right) G + \left(\zeta_Y - \sum_{i=1}^q v_i \Delta_i \right) Z \quad \text{and} \quad (20) \\ R^* &= \left(\gamma_R + \sum_{i=1}^q \rho_i \Gamma_i \right) G + \left(\zeta_R - \sum_{i=1}^q \rho_i \Delta_i \right) Z \end{aligned}$$

where the equalities follow from (17). Substituting X^* , Y^* and R^* in (19) by the above RHSs and grouping the coefficients of Z and G yields

$$\begin{aligned} & \left(\zeta_R - \sum_{i=1}^q \rho_i \Delta_i + c^* \left(\zeta_X - \sum_{i=1}^q \xi_i \Delta_i \right) + d^* \left(\zeta_Y - \sum_{i=1}^q v_i \Delta_i \right) \right) Z \\ &= \left(s^* - \left(\gamma_R + \sum_{i=1}^q \rho_i \Gamma_i \right) - c^* \left(\gamma_X + \sum_{i=1}^q \xi_i \Gamma_i \right) - d^* \left(\gamma_Y + \sum_{i=1}^q v_i \Gamma_i \right) \right) G. \end{aligned} \quad (21)$$

First consider the case where the representations of X^* and Y^* in (20) are independent of Z , that is

$$\zeta_X - \sum_{i=1}^q \xi_i \Delta_i \equiv_p 0 \equiv_p \zeta_Y - \sum_{i=1}^q v_i \Delta_i. \quad (22)$$

The extractor can thus output the witness $\log X^* = \gamma_X + \sum_{i=1}^q \xi_i \Gamma_i$ and $\log Y^* = \gamma_Y + \sum_{i=1}^q v_i \Gamma_i$.

Otherwise, at least one of the terms in (22), which are the coefficients of c^* and d^* in (21), is non-zero. The adversary chose the values $\alpha_i, \beta_i, \delta_i, \gamma_i$ (which define Γ_i and Δ_i) for all $i \in [q]$ when making simulation (or random-oracle) queries *before* making the query $H(X^*, Y^*, m^*, R^*)$. Likewise, it must have chosen the values $\zeta_X, \zeta_Y, \zeta_R$ and ξ_i, v_i, ρ_i for all $i \in [q]$ before making this query. Therefore, (c^*, d^*) is chosen uniformly at random *after* all other values in (21) are defined, and moreover at least one of c^* and d^* is not multiplied by 0. The probability that the coefficient of Z in (21) is congruent to 0 modulo p is thus $\frac{1}{p}$. From (21), the reduction can then efficiently compute $\log Z$ with overwhelming probability, and from the representations of X^* and Y^* in (20), it can compute the witness $(\log X^*, \log Y^*)$. \square

Note that for readability we assumed all queried statements and the one returned are of the form (X, Y, m) . However, the proof is easily extended to simultaneously allow for statements of the form (X, m) , noting that (a) the inputs to the random oracle of type (X, Y, m, R) and (X, m, R) are disjoint and (b) extraction from a “simple” proof is done as from a proof for two elements by setting $d^* = 0$. This immediately yields the following:

Corollary 1. *In the AGM and the ROM, Schnorr signatures (Figure 2) are proofs of secret keys that are strongly simulation-extractable with auxiliary group-element input.*

6 Security analysis of MW-NIT

6.1 Assumptions

In our security analysis of the protocol from Section 3, we assume that range proofs in RaP prove knowledge of the committed value v and the opening q . (Note that for the employed Pedersen commitment, a proof of language membership, that is not “of knowledge” is vacuous, as for any C there always exists an opening e.g. $(v = 0, q = \log C)$.) We thus assume that there exists an extractor that from (an adversary outputting) a range proof π for $C \in \mathbb{G}$ can extract the values $v \in [0, v_{\max}]$ and $q \in \mathbb{Z}_p$.

We assume the existence of strongly simulation-sound (sSS) zero-knowledge (zk) proofs of knowledge (PoK) of the discrete logarithm of group elements with tags. In Section 5, we showed that in the combination of the random-oracle model and the algebraic group model [FKL18], Schnorr signatures are *adaptive* sSS zk-PoKs of the logarithm of the public key, for which the message acts as a tag. We furthermore extended this to proofs of knowledge of two logarithms, so that the proofs are of the same size as Schnorr signatures.

Finally, we assume that the discrete logarithm (DL) problem is hard in the group underlying the system, and for transaction privacy that the decisional Diffie-Hellman (DDH) assumption holds.

6.2 Syntax

We briefly review the syntax of an *aggregate cash system* ACS [FOS19] and describe the adaptations required to capture addresses and non-interactive transactions.

The public parameters and an empty ledger are set up by $(pp, \Lambda) \leftarrow \text{Setup}(1^\lambda, v_{\max})$, which takes as input the security parameter λ in unary and a maximal coin value v_{\max} . A *ledger* Λ specifies a *supply* $\Lambda.\text{sply}$ (also denoted s) representing the value stored in Λ and a list of *transaction outputs* (TXOs) $\Lambda.\text{out}$. Users create addresses (or “wallets”) by running $(pk, vk, sk) \leftarrow \text{KeyGen}(pp)$, which returns a public key (the *address*), a *view key* and a *spending key*.

A *transaction* tx has three attributes: a *supply* tx.sply specifying the amount of money it creates, the *fee* tx.fees paid to miners and a list tx.out of *outputs* txo_i . A transaction is created by running $\text{tx} \leftarrow \text{Send}(pp, (\mathbf{txo}, \mathbf{v}, \mathbf{sk}), (\hat{\mathbf{v}}, \mathbf{pk}, \boldsymbol{\chi}), s, f)$ on vectors of transaction outputs $\mathbf{txo} = (\text{txo}_i)_i$, corresponding values $\mathbf{v} = (v_i)_i$ and spending keys $\mathbf{sk} = (sk_i)_i$; and vectors of output values $\hat{\mathbf{v}}$, destination addresses \mathbf{pk} , and epochs $\boldsymbol{\chi}$; as well as a supply s and a fee f . To aggregate transactions $\text{tx}_1, \dots, \text{tx}_n$, run $\text{tx} \leftarrow \text{Agg}(pp, (\text{tx}_1, \dots, \text{tx}_n))$ (which returns \perp if they are incompatible, e.g. having an input in common).

Using a vector of view keys \mathbf{vk} and a list of values $\hat{\mathbf{v}}$, one obtains the list of outputs of a transaction tx belonging to these keys by running $\mathbf{txo} \leftarrow \text{Rcv}(pp, \text{tx}, \hat{\mathbf{v}}, \mathbf{vk})$. If tx does not spend $\hat{v}_i \in \hat{\mathbf{v}}$ to $vk_i \in \mathbf{vk}$, then $\text{txo}_i \in \mathbf{txo}$ is set to \perp . Finally, $\Lambda' \leftarrow \text{Ldgr}(pp, \Lambda, \text{tx})$ returns an updated ledger Λ' including tx if it is valid and spends outputs present in the ledger, otherwise $\Lambda' := \perp$;

Correctness. We require the following straightforward correctness condition. Let $(pk_j, vk_j, sk_j)_j$ be key triples generated by KeyGen , and for $k \in [n]$ let tx'_k be a transaction, \mathbf{v}'_k a vector of elements in $[0, v_{\max}]$, \mathbf{vk}'_k a vector of elements from \mathbf{vk} (i.e., $\mathbf{vk}'_k = (vk_{j_i})_i$ for some j_i) and let \mathbf{sk}'_k be the corresponding spending keys (i.e., $\mathbf{sk}'_k = (sk_{j_i})_i$). For all $k \in [n]$, define $\mathbf{txo}_k \leftarrow \text{Rcv}(pp, \text{tx}'_k, \mathbf{v}'_k, \mathbf{vk}'_k)$.

Let $(\text{txo}_i)_i := \mathbf{txo}_1 \parallel \dots \parallel \mathbf{txo}_n$ and let I be an index set s.t. $\text{txo}_i \neq \perp$ for all $i \in I$. Consider $\hat{\mathbf{v}} \in [0, v_{\max}]^*$, $\hat{\mathbf{pk}}$ consisting of elements from \mathbf{pk} , some $\boldsymbol{\chi}$ and $s, f \in [0, v_{\max}]$ such that $\sum_{i \in I} v'_i = \sum \hat{\mathbf{v}} + f - s$. Then for $\text{tx} \leftarrow \text{Send}(pp, ((\text{txo}_i)_{i \in I}, (v'_i)_{i \in I}, (sk'_i)_{i \in I}), (\hat{\mathbf{v}}, \hat{\mathbf{pk}}, \boldsymbol{\chi}), s, f)$ and $\mathbf{txo} \leftarrow \text{Rcv}(pp, \text{tx}, \hat{\mathbf{v}}, \hat{\mathbf{vk}})$ where $\hat{\mathbf{vk}}$ corresponds to $\hat{\mathbf{pk}}$, we have $\perp \notin \mathbf{txo}$, that is, all outputs are accepted.

Comparison to FOS. The syntax of Send and Rcv differs from the one in [FOS19] due to the inclusion of addresses (as well as fees and epochs) and transactions being non-interactive. We moreover simplified notation by merging their algorithm Mint , used for creating money, with Send (which is now non-interactive and takes a supply as input). That is, $\text{Mint}(pp, \hat{\mathbf{v}}, \mathbf{pk}, \boldsymbol{\chi})$ is an alias for $\text{Send}(pp, (), (\hat{\mathbf{v}}, \mathbf{pk}, \boldsymbol{\chi}), \sum \hat{\mathbf{v}}, 0)$.

6.3 Inflation resistance

Definition. Inflation resistance guarantees that the only way to create money in an *aggregate cash system*, such as Mimblewimble, is explicitly via the supply contained in transactions. The notion is defined by the following game, adapted from [FOS19, Def. 10].

The adversary is given the system parameters (for MW-NIT they contain the elements G and H and potential parameters for the range proof), and its task is to produce a (valid) ledger and a transaction tx^* (accepted by the ledger) that spends an amount that exceeds the supply of the ledger (plus its own supply). In addition to the output amounts $\hat{\mathbf{v}}$ of tx^* , the adversary must also return view keys that accept the outputs of tx^* . Letting s denote the ledger supply, s^* the supply and f^* the fee of tx^* , the adversary wins if

$$s < \sum \hat{\mathbf{v}} + f^* - s^* . \quad (23)$$

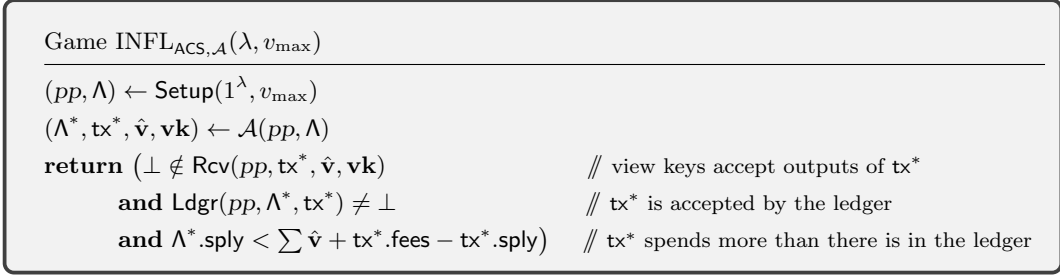


Fig. 4. Game INFL for inflation resistance

The game is formalized in [Figure 4](#).

Theorem 1. *If the range-proof system RaP and the proof-of-possession system PoP are extractable and if the discrete-logarithm assumption holds in the underlying group, then MW-NIT satisfies inflation resistance.*

Proof. Our analysis follows closely that of [[FOS19](#), Theorem 13] for MW-FOS, since a ledger (or transaction) in MW-NIT contains an MW-FOS ledger (or transaction). A small difference is that FOS did not consider fees and kernel offsets, but these are easily added to the argument.

Consider a successful adversary that returns a ledger and a transaction

$$\begin{aligned} \Lambda^* &= (\mathbf{C}, \boldsymbol{\pi}, \mathbf{R}, \boldsymbol{\rho}, \mathbf{P}, \boldsymbol{\chi}, (\mathbf{ct}, s, f, t, y, \mathbf{E}, \mathbf{X}, \boldsymbol{\sigma})) \\ \text{tx}^* &= ((P'_i, D_i, \psi_i)_{i \in [n]}, (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \chi_i)_{i \in [\hat{n}]}, (\mathbf{ct}^*, s^*, f^*, t^*, y^*, \mathbf{E}^*, \mathbf{X}^*, \boldsymbol{\sigma}^*)) \end{aligned}$$

together with output values $(\hat{v}_i)_{i \in [\hat{n}]}$ and view keys $(a_i, B_i)_{i \in [\hat{n}]}$ that accept the outputs.

The reduction extracts all values v_i and openings q_i for all the output commitments C_i in Λ^* from the respective range proofs π_i ; thus $C_i = \text{Cmt}(v_i, q_i)$. Since the ledger is valid, from [\(11\)](#) we have

$$\sum \mathbf{E} = \sum \text{Cmt}(v_i, q_i) + \text{Cmt}(f - s, 0) - \text{Cmt}(0, t) . \quad (24)$$

We first show that the ledger must be balanced w.r.t. the values v_i extracted from the output commitments C_i , that is

$$s = \sum \mathbf{v} + f . \quad (25)$$

From the kernel proofs σ_i in Λ^* , the reduction can extract $e_i := \log E_i$. Since $\sum \mathbf{E} = \text{Cmt}(0, \sum \mathbf{e})$, we have that $(0, \sum \mathbf{e})$ is an opening of the right-hand side of [\(24\)](#). On the other hand, since Cmt is homomorphic, $(\sum \mathbf{v} + f - s, \sum \mathbf{q} - t)$ is also an opening. Thus if [\(25\)](#) did not hold then we would have two openings for two different values. This represents a break of the binding property of Pedersen commitments, which holds under the DL assumption. (This is formally proven by a reduction that obtains a DL challenge H , which then plays the role of Z in the AGM proof of extractability of PoP in [Section 5](#).)

Let \hat{q}_i be openings of the commitments in the outputs of tx^* , derived from \hat{R}_i and the view key (a_i, B_i) . Since the latter accepted the i -th output of tx^* , we have $\hat{C}_i = \text{Cmt}(\hat{v}_i, \hat{q}_i)$. Since Λ^* accepted tx^* , all values P'_i are in the ledger, that is, $\mathbf{P}' = (P_i)_{i \in I}$ for some $I \subseteq \mathbb{N}$. From validity of tx^* we further have

$$\sum \mathbf{E}^* = \sum \hat{\mathbf{C}} - \sum_{i \in I} C_i + \text{Cmt}(f^* - s^*, 0) - \text{Cmt}(0, t^*) . \quad (26)$$

From σ_i^* contained in tx^* , the reduction can extract $e_i^* := \log E_i^*$ and, analogously to the above, from [\(26\)](#) we get

$$\text{Cmt}(0, \sum e^*) = \text{Cmt}(\sum \hat{\mathbf{v}} - \sum_{i \in I} v_i + f^* - s^*, \sum \hat{\mathbf{q}} - \sum_{i \in I} q_i - t^*) . \quad (27)$$

<p>Game $\text{THFT}_{\text{ACS}, \mathcal{A}}(\lambda, v_{\max})$</p> <hr/> <p>$(pp, \Lambda) \leftarrow \text{Setup}(1^\lambda, v_{\max})$ $j := 0; \text{Hon}, \text{Archv} := ()$ $\mathcal{A}^{\text{KEYGEN}, \text{SEND}, \text{LEDGER}}(pp, \Lambda)$ return $(\text{Hon} \not\subseteq \Lambda.\text{out})$</p> <hr/> <p>Oracle $\text{SEND}((\mathbf{txo}, \mathbf{v}, I), (\hat{\mathbf{v}}, \mathbf{pk}, \boldsymbol{\chi}), s, f)$</p> <hr/> <p>if $\mathbf{txo} \not\subseteq \text{Hon}$ then return \perp $\mathbf{tx} \leftarrow \text{Send}(pp, (\mathbf{txo}, \mathbf{v}, (sk_i)_{i \in I}), (\hat{\mathbf{v}}, \mathbf{pk}, \boldsymbol{\chi}), s, f)$ if $\mathbf{tx} = \perp$ then return \perp $\text{Hon} := \text{Hon} - \mathbf{txo}$ // remove spent outputs return \mathbf{tx}</p>	<p>Oracle $\text{KEYGEN}()$</p> <hr/> <p>$j := j + 1$ $(pk_j, vk_j, sk_j) \leftarrow \text{KeyGen}(pp)$ return (j, pk_j, vk_j)</p> <hr/> <p>Oracle $\text{LEDGER}(\mathbf{tx}, \hat{\mathbf{v}}, (i_j)_j)$</p> <hr/> <p>$\Lambda' := \text{Ldgr}(pp, \Lambda, \mathbf{tx})$ if $\Lambda' = \perp$ then return \perp else $\Lambda := \Lambda'$ // Check for new honest outputs in \mathbf{tx}: $\mathbf{txo} \leftarrow \text{Rcv}(pp, \mathbf{tx}, \hat{\mathbf{v}}, (vk_{i_j})_j)$ for $i \in [\mathbf{txo}]$: if $txo_i \neq \perp$ and $txo_i \notin \text{Archv}$ $\text{Hon} := \text{Hon} \parallel txo_i; \text{Archv} := \text{Archv} \parallel txo_i$ return Λ</p>
---	--

Fig. 5. Game THFT for theft resistance

If the adversary outputs \mathbf{tx}^* satisfying (23), then we have

$$\sum_{i \in I} v_i \leq \sum \mathbf{v} + f \stackrel{(25)}{=} s < \sum \hat{\mathbf{v}} + f^* - s^* ,$$

and thus $\sum \hat{\mathbf{v}} + f^* - s^* - \sum_{i \in I} v_i \neq 0$ (since all terms are bounded by v_{\max} and thus never wrap around the modular representation). This means that the two sides in (27) are two different openings of the commitment $\sum \mathbf{E}^*$, which represents a break of commitment-binding. \square

6.4 Theft resistance

We define two notions that protect users from losing money. The first one is an adaptation of the notion from [FOS19] to a scheme with non-interactive transactions. It guarantees that outputs in the ledger belonging to a user can only be spent by that user.

The main difference between MW-FOS and MW-NIT is that the former relies on the coin keys (the opening of the commitments) being kept secret, while in MW-NIT, the spender knows (and defines) the openings of the receivers' commitments.¹² In MW-NIT, the security relies on the secrecy of the “*spend key*” for the user's stealth address. We assume that the *view key* is known to the adversary (as delegating scanning for transactions should not endanger the security of these transactions).

Definition. Resistance to theft means that for any output belonging to a user in the ledger (that is, it was accepted by the user's view key), no matter how it was received (e.g., sent by the adversary), as long as the output has never been spent before and the user keeps her spend key safe, no one except her can spend it (even if her view key is publicly known). This is formalized via the following game, which is an adaptation of [FOS19, Fig. 8].

Honest users are simulated by the experiment and the adversary can create new users by calling KEYGEN ; the adversary can instruct users spend outputs they own (stored in the list Hon) by calling SEND ; moreover it can submit any transactions to the ledger using its oracle LEDGER . If the transaction contains outputs that are accepted by honest users, these are added to Hon . As

¹² Note that this is unavoidable for non-interactive transactions: knowing the (sum of) the receivers' keys is necessary to compute the excess proof σ .

users are not supposed to accept outputs they have already owned (cf. Section 4.4), the experiment stores all outputs ever received in a list *Archv*; if an output is already in *Archv*, then LEDGER does not add it to *Hon*. The adversary wins the game if it spends any of the honest user’s outputs, that is, if some output in *Hon* is not in the ledger. The game is formalized in Figure 5.

Remark. Note that the game in Figure 5 for a scheme with non-interactive transactions is a lot simpler than [FOS19, Fig. 8], which had to take care of the interactive spending protocol. In particular, this involves an instruction for the honest users to *receive* coins, and the definition of the coins an honest user owns in [FOS19] is cumbersome, whereas for non-interactive transactions, anything accepted by the honest user’s view key (and not previously owned) is considered belonging to her. Note also that we do not require an oracle MINT, since the adversary can run SEND with an empty input list.

Theorem 2. *If PoP is a simulation-sound proof of knowledge of discrete logarithms (cf. Section 5) and the DL assumption holds in the underlying group, then MW-NIT satisfies theft-resistance.*

Proof. Consider a user owning an output $txo = (C, \pi, R, \rho, P, \chi)$ with amount v in the ledger, that is, txo is accepted (see Section 3.4) by her view key (a, B) , meaning $P = B + kG$ where $(k, q) = H(aR)$. Spending this coin requires proving possession of P with tag txo , but an honest user, unless she spends that output, never proves possession of P with tag txo .

We formally use a theft to break the DL assumption assuming simulation-extractability of PoP. The reduction receives a DL challenge $B^* \in \mathbb{G}$, chooses $a^* \leftarrow_{\$} \mathbb{Z}_p$, sets a random honest user’s address to $(A^* := a^*G, B^*)$ and gives the adversary the view key (a^*, B^*) .

Whenever the user is asked to spend an output $txo' = (C', \pi', R', \rho', P', \chi')$ belonging to the user, the reduction computes the transaction as specified, choosing a doubling key $D' := d'G$ for a $d' \leftarrow_{\$} \mathbb{Z}_p$. As it does not know the logarithm of P' (since this requires knowledge of $\log B^*$), it runs the zero-knowledge simulator for a proof of possession ψ for (P', D') with tag txo' .

Assume an output $txo = (C, \pi, R, \rho, P, \chi)$ belonging to the user is spent by the adversary. (If the adversary attacked a different user, the reduction aborts.) Then the corresponding transaction input must contain a proof ψ^* of possession of (P, D^*) with tag txo for some D^* . By the definition of the security game, the honest user has never spent txo before. This means that the reduction has never simulated a proof for (P, D^*, txo) . Therefore, by simulation-extractability (SE) of PoP, the reduction can extract $p = \log P$, and since $P = B^* + kG$, for a value k known to the reduction, it can compute the solution $p - k$ to its DL challenge B^* .

Note that if extraction fails, we can construct a reduction that breaks SE of PoP, which is algebraic if the adversary is: from the adversary’s representation of (P, D^*) , the reduction against PoP can derive a representation in basis (G, B^*) , where B^* is the auxiliary group element input; likewise, it can give representations $P' = B^* + k'G$ and $D' = d'G$ of the elements for which it queries simulated proofs. \square

6.5 Transaction-binding

While the previous notion states that once a user owns an output in the ledger it cannot be stolen, we also need to guarantee that nothing can be stolen from a transaction *before* it is even added to the ledger (this protects against malicious miners, for example). In particular, if a user produces a transaction tx then no one should be able to create a transaction that contains *one* of the inputs of tx while not containing *all* its outputs (except for the owners of the outputs). Since transactions can be aggregated, further inputs and outputs can also be added to the original transaction. Thus, while *theft-resistance* protects the outputs of a transaction, *transaction-binding* protects the inputs.

Definition. We define transaction-binding via the following game. The experiment simulates all honest users, which the adversary can create by calling KEYGEN, which creates a new address,

Game $\text{TXBND}_{\text{ACS}, \mathcal{A}}(\lambda, v_{\max})$	
<hr/> $(pp, \Lambda) \leftarrow \text{Setup}(1^\lambda, v_{\max})$ $j := 0; HT_{\text{XS}} := ()$ // honest transactions $ST_{\text{XO}} := ()$ // spent outputs $(\Lambda, \text{tx}^*, \widehat{txo}, \hat{v}, i) \leftarrow \mathcal{A}^{\text{KEYGEN}, \text{SEND}, \text{LEDGER}}(pp, \Lambda)$ // Λ should contain the TXOs spent by tx^* $\Lambda^* := \text{Ldgr}(pp, \Lambda, \text{tx}^*)$ return $(\exists \text{tx} \in HT_{\text{XS}} : \text{Ldgr}(pp, \Lambda, \text{tx}) \neq \perp$ // tx accepted by Λ , ... and $\Lambda^* \neq \perp$ and $\text{Ldgr}(pp, \Lambda^*, \text{tx}) = \perp$ // ... but not after tx^* was (accepted and) added and $\widehat{txo} \in \text{Rcv}(pp, \text{tx}, \hat{v}, vk_i)$ // \widehat{txo} belongs to an honest user, ... and $\widehat{txo} \notin ST_{\text{XO}}$ and $\widehat{txo} \notin \text{tx}^*.\text{out}$) // ... and was not spent, and it is not in tx^*	
<hr/> Oracle $\text{KEYGEN}()$ $j := j + 1$ $(pk_j, vk_j, sk_j) \leftarrow \text{KeyGen}(pp)$ return (j, pk_j, vk_j)	<hr/> Oracle $\text{SEND}((\text{txo}, \mathbf{v}, I), (\hat{\mathbf{v}}, \hat{\mathbf{pk}}, \hat{\chi}), s, f)$ // Honest users never spend the same txo twice: if $\text{txo} \cap ST_{\text{XO}} \neq \emptyset$ then return \perp $\text{tx} \leftarrow \text{Send}(pp, (\text{txo}, \mathbf{v}, (sk_i)_{i \in I}), (\hat{\mathbf{v}}, \hat{\mathbf{pk}}, \hat{\chi}), s, f)$ $HT_{\text{XS}} := HT_{\text{XS}} \parallel \text{tx}; ST_{\text{XO}} := ST_{\text{XO}} \parallel \text{txo}$ return tx

Fig. 6. Game TXBND for transaction-binding

for which the adversary receives the view key. The adversary can instruct honest users to spend outputs of the adversary’s choice to addresses of the adversary’s choice; the experiment computes the corresponding transaction using the users’ spending keys and gives it to the adversary. The adversary’s goal is to create a transaction tx^* which *spends the same input as an honest transaction* tx , but *one of the outputs in tx for an honest user is missing in tx^** (and that user was not instructed to spend it).

We formalize this by requiring the adversary to return a ledger Λ that must accept the attacked honest transaction tx , but after tx^* is included in Λ , the latter does not accept tx anymore (thus tx and tx^* have an input in common); the adversary also returns values and indices of honest users, so that their view keys accept tx but do not accept tx^* (thus one of the outputs of tx is missing in tx^*). The game is formalized in [Figure 6](#).

Remark. While it might seem restrictive that only stealing outputs of *honest* users is considered a break of the notion, it is not. In aggregate cash systems like Mimblewimble an adversary can always “steal” outputs it owns from any transaction tx : simply create a transaction that spends these outputs and then merge it with tx ; the result is a transaction in which some of the outputs of tx have been replaced by new ones. Thus, we do not consider this an attack and transaction-binding gives no guarantees against it.

Theorem 3. *If PoP is a strongly simulation-sound proof of knowledge of discrete logarithms (cf. [Section 5](#)) and the DL assumption holds in the underlying group, then MW-NIT satisfies transaction-binding.*

Proof intuition. As this is the most complex notion to prove, we start with a simplified scenario. Consider an adversary \mathcal{A} that creates two users via KEYGEN , computes an output $txo = (C, \pi, R, \rho, P, \chi)$ for User 1 and calls SEND so txo is spent to User 2, which creates a transaction

$$\text{tx} = ((P, D, \psi), \widehat{txo}, (s, f, t, y, \sigma)) \quad \text{with} \quad \widehat{txo} = (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi}) .$$

\mathcal{A} makes no further oracle calls and returns a transaction $\text{tx}^* = ((P, D^*, \psi^*), \text{txo}^*, (s^*, f^*, t^*, y^*, \sigma^*))$ that spends txo and creates an output $\text{txo}^* = (C^*, \pi^*, R^*, \rho^*, P^*, \chi^*) \neq \widehat{\text{txo}}$ (so tx^* is a *simple*, i.e., non-aggregate transaction).

We first show that conditions (I)–(III) below must hold with overwhelming probability, and in (IV) we show that when (I)–(III) hold, \mathcal{A} can only win with negligible probability. We show all four claims by reducing the DL problem to them, assuming strong simulation-extractability (SE) of the proof-of-possession system PoP.

(I) $D = D^*$. Assume $D \neq D^*$. By SE of PoP, this can be used to break DL by simulating the game as in the proof of theft-resistance (Theorem 2):

Given a DL challenge B , the reduction embeds it into User 1’s address (aG, B) for $a \leftarrow_{\$} \mathbb{Z}_p$. It computes tx as prescribed, except that it simulates the proof ψ (as its witness depends on $\log B$). When \mathcal{A} outputs a transaction with input (P, D^*, ψ^*) , the witness (p, d^*) can be extracted from ψ^* , as the only simulated proof was for a different statement (P, D, txo) . Since $P = B + kG$, with k derived from User 1’s view key per Equation (3), the reduction can return $\log B = p - k$.

(II) $R^* \neq \hat{R}$. Assume $R^* = \hat{R}$. Since $\text{txo}^* \neq \widehat{\text{txo}}$, either the tag $C^* \parallel \pi^* \parallel P^* \parallel \chi^*$ for ρ^* , or ρ^* itself is different. As the statement/proof pair is different from the one created by User 1, $\log R^*$ can be extracted by strong SE (“strong”, since possibly only the proof differs):

Given a DL challenge \hat{R} , the reduction embeds it in the output of tx . Not knowing $\log \hat{R}$, the reduction can still compute $(\hat{k}, \hat{q}) = H(a'\hat{R})$ from User 2’s view key (a', B') (who must be honest). From this, the reduction computes $\hat{C}, \hat{\pi}$ and \hat{P} . The proofs $\hat{\rho}$ and σ , whose witnesses depend on $\log \hat{R}$, are simulated. If the adversary returns a transaction tx^* with $R^* = \hat{R}$, then from ρ^* the reduction can extract $\log \hat{R}$ (since the tag or ρ^* must be different), solving the DL challenge.

(III) $X^* \neq X$, that is, the stealth excess of tx^* is different from that of tx . Assume $X^* = X$. Then from the definition of stealth excesses in Equation (2) we get $R^* - D^* - y^*G = \hat{R} - D - yG$. By (I) we have $D^* = D$ and thus $\hat{R} = R^* + (y - y^*)G$. Since ρ^* proves knowledge of $\log R^*$, this means the adversary must also know $\log \hat{R}$.

Formally, as in case (II), the reduction embeds a DL challenge as \hat{R} and simulates the proofs $\hat{\rho}$ and σ to compute tx . Since, $R^* \neq \hat{R}$ by (II), the statement of ρ^* is different from that of the simulated proof $\hat{\rho}$ and thus the reduction can extract $r^* = \log R^*$ from ρ^* . If $X^* = X$ then it can compute $\log \hat{R} = r^* + y - y^*$.

(IV) Finally, we show that \mathcal{A} cannot win the game when (I)–(III) hold. From (I) and (2), we have $X^* := R^* - D - y^*G$. Since σ^* contained in tx^* proves knowledge of $x^* := \log X^*$ (since $X^* \neq X$ by (III)) and ρ^* proves knowledge of $r^* := \log R^*$, this means that the adversary must know $\log D = r^* - y^* - x^*$.

The reduction embeds its DL challenge as D , the value in the input of tx , and simulates proofs ψ and σ (whose DL depend on $\log D$). All other components of tx are computed as prescribed. If the adversary is successful, then from ρ^* the reduction extracts r^* (no ρ proofs were simulated), and from σ^* the reduction extracts x^* (the simulated proof σ was under $X \neq X^*$) and returns $r^* - y^* - x^* = \log D$.

Proof (of Theorem 3). We generalize the above arguments, now considering the actual security game. In particular, there can now be many honest transactions created by the game and the adversary’s output tx^* can be an aggregated transaction, whose stealth excesses might contain X values from different honest transactions. Moreover, tx^* can contain cut-through terms (in particular, the “stolen” output $\widehat{\text{txo}}$ could be in the cut-through of tx^*).

Let tx^* and $\widehat{\text{txo}}$ be the adversary’s outputs and tx be the attacked transaction, that is, tx and tx^* have an input in common, and txo is an output of tx (accepted by an honest user) but not of tx^* . We proceed by a sequence of hybrid games, each one introducing new abort conditions which make the game return 0. The conditions in H_1 correspond to (I) above and $\widehat{\text{txo}}$ being in the cut-through of tx^* . The conditions in H_2 correspond to (II) and (III), which we combine in the same hybrid as they can be proved in one reduction.

H_0 This is the original transaction-binding game.

H_1 The first hybrid game aborts if one of the inputs in tx^* spends the same txo as an honest transaction, i.e., $\text{txo} \in \text{HTxs}$, but the associated D value is different (abort condition (A1a)); or if $\widehat{\text{txo}} \in \text{ct}^*$, the cut-through of tx^* (A1b).

H_2 In addition, the second hybrid aborts if either the value \hat{R} in $\widehat{\text{txo}}$ appears in tx^* (that is, either in an output or in its cut-through) (A2a); or \hat{R} does not appear in tx^* and among the stealth excesses of tx^* is the stealth excess X of tx (A2b).

Note to win H_2 , tx^* must share an input with tx (the attacked transaction), and moreover the corresponding D values in tx^* and tx must be the same (otherwise the game would abort because of (A1a)). Also note that if tx^* contains the stealth excess X of tx then H_2 returns 0: if it does not abort in (A2b) then \hat{R} must appear in tx^* , and thus it aborts in (A2a). Therefore, if the adversary wins H_2 then

(W1) one of the D values in tx^* is from an input of tx

(W2) the stealth excess list \mathbf{X} of tx^* does not contain the stealth excess X of tx

Hybrid H_0 to H_1 . We start by showing that the probability that H_0 returns 1 but H_1 does not (because it aborts) is negligible by reducing the DL problem to the event (A1a) or (A1b) occurring, assuming simulation-extractability (SE) of PoP.

Given a DL challenge $Z \in \mathbb{G}$, the reduction embeds it into the addresses of *all* honest users: it picks $a_j, b_j \leftarrow \mathbb{Z}_p$ and defines the j -th user's address as $(a_j G, B_j := Z + b_j G)$. Note that this is correctly distributed and the corresponding view key (a_j, B_j) can be given to the adversary. When the adversary calls SEND, the reduction proceeds as prescribed, except that runs the zero-knowledge simulator for the proof ψ (the only value that depends on $\log B_j$). Let $(\text{txo}_i)_i = \text{STxo}$ be the outputs queried to SEND and let (P_i, D_i, ψ_i) be the corresponding inputs of the transactions created by SEND. Thus, ψ_i are the simulated proofs for (P_i, D_i, txo_i) . Since SEND never accepts the same txo twice, all txo_i are distinct.

(A1a). Let tx^* be the transaction returned by the adversary, and assume (A1a) occurs, i.e., tx^* spends txo_{i^*} for some i^* , and the corresponding input is of the form (P_{i^*}, D^*, ψ^*) with $D^* \neq D_{i^*}$. Since $(P_{i^*}, D^*, \text{txo}_{i^*}) \neq (P_{i^*}, D_{i^*}, \text{txo}_{i^*})$ and all txo_i are distinct, we have $(P_{i^*}, D^*, \text{txo}_{i^*}) \neq (P_i, D_i, \text{txo}_i)$ for all i . This means that no proof for $(P_{i^*}, D^*, \text{txo}_{i^*})$ has been simulated and by SE of PoP, the reduction can extract the witness (p_{i^*}, d^*) from ψ^* . Let j be the index of the user that spent txo_{i^*} (i.e., the value in I associated to Tx_{i^*} in the call to SEND); then $P_{i^*} = B_j + kG$ for k derived the view key (a_j, B_j) as per (3). Since $B_j = Z + b_j G$, we get $\log Z = p_{i^*} - b_j - k$, the solution of the DL challenge.

(A1b). Now assume that (A1b) occurs, that is, $\widehat{\text{txo}}$, the value returned by the adversary, is an output of an honest transaction tx , and $\widehat{\text{txo}} \notin \text{STxo}$ (otherwise the adversary did not win), and the cut-through of tx^* contains $\widehat{\text{txo}} \parallel (\hat{P}, D^*, \psi^*)$ for \hat{P} in $\widehat{\text{txo}}$ and some D^*, ψ^* . Since P values are derived from R values, which are uniform, the probability that \hat{P} is contained in some $\text{txo}_i \in \text{STxo}$ is negligible. As the entries in STxo are precisely those for which ψ proofs were simulated and $\widehat{\text{txo}} \notin \text{STxo}$, the reduction can extract from ψ^* the witness (\hat{p}, d^*) . Since the i -th user (with i from the adversary's output) accepted $\widehat{\text{txo}}$, we have $\hat{P} = B_i + \hat{k}G$ for \hat{k} derived from (a_i, B_i) as per (3). Since $B_i = Z + b_i G$, the reduction can thus again compute $\log Z = \hat{p} - b_i - \hat{k}$.

Hybrid H_1 to H_2 . We next show that in a winning run of H_2 the abort conditions (A2a) and (A2b) can only occur with negligible probability. If any of them occurs, then there exist i and \hat{i} so that tx will be the i -th transaction and its \hat{i} -th output will be $\widehat{\text{txo}}$, and one of the following hold:

(A2a) \hat{R} from $\widehat{\text{txo}}$ is contained in some txo^* in tx^* (either in an output or in its cut-through list ct^*). (Note that we have $\widehat{\text{txo}} \neq \text{txo}^*$: otherwise, if $\widehat{\text{txo}}$ is an output then the adversary did not win, and if $\widehat{\text{txo}}$ is in ct^* then the game would have aborted because of (A1b)); or

(A2b) \hat{R} from \widehat{txo} does not appear in tx^* and the stealth excess X of tx is among the stealth excess list of tx^* .

We define a reduction that makes a *guess* (i, \hat{i}) , so that if the adversary wins H_1 but not H_2 (i.e., either (A2a) or (A2b) occur) and the guess is correct, then it breaks DL.

The reduction receives a DL instance $Z \in \mathbb{G}$ and guesses i and \hat{i} . The reduction simulates H_1 as prescribed, except when creating the i -th transaction tx_i , in its \hat{i} -th output it sets $R_{\hat{i}} := Z$ (and thus does not know its discrete logarithm). It completes the transaction as follows: if the receiver of $tx_{o_{\hat{i}}}$ is not honest, it aborts (and the guess (i, \hat{i}) was wrong). Otherwise, let (a_j, B_j) be receiver's view key. The reduction computes

$$(k_{\hat{i}}, q_{\hat{i}}) := H(a_j R_{\hat{i}})$$

and, from this, $C_{\hat{i}}$, $P_{\hat{i}}$ and $\pi_{\hat{i}}$ as prescribed. It simulates a proof $\rho_{\hat{i}}$ for $(R_{\hat{i}}, C_{\hat{i}} \parallel \pi_{\hat{i}} \parallel P_{\hat{i}} \parallel \chi_{\hat{i}})$, which completes $tx_{o_{\hat{i}}}$. To complete tx_i , it also simulates σ (whose witness depends on $\log R_{\hat{i}}$).

(A2a). Assume that the reduction guessed correctly and that (A2a) occurs, that is, tx^* contains $(C^*, \pi^*, R_{\hat{i}}, \rho^*, P^*, \chi^*) \neq \widehat{txo} = tx_{o_{\hat{i}}}$. Thus, ρ^* is valid for $(R_{\hat{i}}, C^* \parallel \pi^* \parallel P^* \parallel \chi^*)$ and $(C^* \parallel \pi^* \parallel P^* \parallel \chi^*, \rho^*) \neq (C_{\hat{i}} \parallel \pi_{\hat{i}} \parallel P_{\hat{i}} \parallel \chi_{\hat{i}}, \rho_{\hat{i}})$; otherwise we would have $txo^* = \widehat{txo}$. The statement for ρ^* (or ρ^* itself) is thus different from the statement for the simulated $\rho_{\hat{i}}$ (or $\rho_{\hat{i}}$ itself). By strong SE of PoP, the reduction can extract from ρ^* the witness $\log R_{\hat{i}} = \log Z$, the DL solution. (Recall that σ is also simulated, but that different types of proofs are assumed to be domain-separated, cf. Section 2)

(A2b). Now assume that the reduction guessed correctly and that (A2b) occurs. This means that $tx^* = ((\mathbf{P}^*, \mathbf{D}^*, \psi^*), (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi}), ((\mathbf{C}', \pi', \mathbf{R}', \rho', \mathbf{P}', \chi', \mathbf{P}', \mathbf{D}', \psi'), s^*, f^*, t^*, y^*, \mathbf{E}, \mathbf{X}, \sigma))$,

returned by the adversary, is such that \mathbf{X} contains the stealth excess X of tx and neither \hat{R} nor \mathbf{R}' contain $R_{\hat{i}}$. By validity of tx , it satisfies the balance equation (12) and since $R_{\hat{i}} = Z$, we have

$$X = \sum_{j \neq \hat{i}} R_j + Z - \sum \mathbf{D} - yG . \quad (28)$$

All values $r_j = \log R_j$, for $j \neq \hat{i}$, and $d_j = \log D_j$ have been chosen by (and are thus known to) the reduction. Moreover, since (A2b) occurred, for some J^* with $|J^*| =: m \geq 1$ we have $X_j = X$ for all $j \in J^*$ (the adversary could have included X multiple times in \mathbf{X}). By validity of tx^* (recall that for aggregated transactions, Equation (12) also ranges over the cut-through terms) this yields

$$\sum_{j \notin J^*} X_j + mX = \sum \hat{R} + \sum \mathbf{R}' - \sum \mathbf{D}^* - \sum \mathbf{D}' - y^*G . \quad (29)$$

From the proofs σ_j in tx^* , for $j \notin J^*$, the reduction can extract the corresponding values $x_j = \log X_j$ (since only σ for X was simulated and $X \notin (X_j)_{j \notin J^*}$). From the proofs ψ_j^* and ψ_j' the reduction can extract the values $d_j^* = \log D_j^*$ and $d_j' = \log D_j'$, since no ψ proofs are simulated (if the adversary copied D and ψ from an honest transaction, the reduction knows $\log D$). From the proofs $\hat{\rho}_j$ and ρ_j' , the reduction can extract the values $\hat{r}_j = \log \hat{R}_j$ and $r_j' = \log R_j'$ (and it knows the logarithms of the ones copied from honest transactions). Recall that $R_{\hat{i}}$ (for which $\rho_{\hat{i}}$ was simulated) is not among \hat{R} nor \mathbf{R}' .

Plugging (28) into (29), taking the logarithm to base G and substituting all logarithms of group elements known to the reduction thus yields

$$\log Z \equiv_p \frac{1}{m} \left(- \sum_{j \notin J^*} x_j + \sum \hat{r} + \sum r' - \sum d^* - \sum d' - y^* \right) - \sum_{j \neq \hat{i}} r_j + \sum d + y ,$$

meaning the reduction can solve the DL challenge.

Hybrid H_2 . It remains to show that H_2 can only be won with negligible probability, where we leverage that winning the game implies (W1) and (W2).

Again, we reduce solving DLs to winning H_2 , assuming SE of PoP. The reduction now embeds a DL challenge $Z \in \mathbb{G}$ into *all* the D values chosen by the oracle SEND. Whenever the adversary

calls SEND, the reduction randomly samples $d \leftarrow \mathbb{Z}_p$, sets $D := dZ$ and simulates proofs ψ for (P, D, txo) , as well as the excess proof σ (whose witnesses depend on $\log D$). Let

$$\mathbf{tx}^* = ((P, D, \psi), (\hat{C}, \hat{\pi}, \hat{R}, \hat{\rho}, \hat{P}, \hat{\chi}), ((C', \pi', R', \rho', P', \chi', P', D', \psi'), s, f, t, y, E, X, \sigma)) ,$$

be the transaction returned by the adversary. By (W1) and (W2), D contains an input D from an honest transaction \mathbf{tx} , whose stealth excess X is not contained in X . In \mathbf{tx}^* we distinguish two types of values in $D \parallel D'$ and two types of stealth-excess elements:

- let \bar{D}_i be the values in D and D' that appear in honest transactions and D_i^* be those that don't
- let \bar{X}_i be the stealth excesses that appear in honest transactions and X_i^* be those that don't

For all \bar{X}_i appearing in an honest transaction \mathbf{tx}_i , let $D_{i,j}$ and $R_{i,j}$ denote the values contained in the inputs and outputs of \mathbf{tx}_i and let y_i be its stealth offset. By validity of \mathbf{tx}_i we have

$$\bar{X}_i = \sum_j R_{i,j} - \sum_j D_{i,j} - y_i G .$$

As these values $D_{i,j}$ and the values \bar{D}_i contained in \mathbf{tx}^* were created by the reduction, we have $D_{i,j} = d_{i,j} \cdot Z$ and $\bar{D}_i = \bar{d}_i \cdot Z$ for values $d_{i,j}$ and \bar{d}_i known to the reduction. Moreover, the reduction knows the values $r_{i,j} := \log R_{i,j}$, as it chose them. From the above, we thus have

$$\log \bar{X}_i = \sum_j r_{i,j} - \sum_j (d_{i,j} \cdot \log Z) - y_i \qquad \log \bar{D}_i = \bar{d}_i \cdot \log Z \qquad (30)$$

From the proofs ψ_i and ψ'_i in \mathbf{tx}^* corresponding to the values D_i^* and from σ_i corresponding to X_i^* , the reduction can extract $\log D_i^*$ and $\log X_i^*$ (since, by definition, D_i^* and X_i^* are values that have not been created by the reduction and thus no proofs have been simulated for them). Moreover, if any values \hat{R}_i or \hat{R}'_i in \mathbf{tx}^* were copied from an honest transaction, the reduction knows their logarithms. For all other \hat{R}_i and \hat{R}'_i , from the proofs $\hat{\rho}_i$ and $\hat{\rho}'_i$ contained in \mathbf{tx}^* , the reduction can extract $\log \hat{R}_i$ and $\log \hat{R}'_i$ (as no ρ proofs are simulated). The reduction thus knows the values

$$\log D_i^* =: d_i^* \qquad \log X_i^* =: x_i^* \qquad \log \hat{R}_i =: \hat{r}_i \qquad \log \hat{R}'_i =: \hat{r}'_i \qquad (31)$$

Since \mathbf{tx}^* is valid, it satisfies the balance equation (12) for aggregated transactions, that is $X = \sum \hat{R} + \sum R' - \sum D - \sum D' - yG$. Since the lists X and $X^* \parallel \bar{X}$ contain the same elements, as do $D \parallel D'$ and $D^* \parallel \bar{D}$, this yields

$$\sum X^* + \sum \bar{X} = \sum \hat{R} + \sum R' - \sum D^* - \sum \bar{D} - yG .$$

Taking the logarithm to base G , and using (30) and (31) this yields:

$$\sum x_i^* + \sum \sum r_{i,j} - \sum \sum (d_{i,j} \cdot \log Z) - \sum y_i \equiv_p \sum \hat{r}_i + \sum \hat{r}'_i - \sum d_i^* - \sum (\bar{d}_i \cdot \log Z) - y . \qquad (32)$$

We argue that with overwhelming probability we have

$$\sum \bar{d}_i - \sum \sum d_{i,j} \not\equiv_p 0 . \qquad (33)$$

Indeed, by (W1), at least one value \bar{d}_i^* comes from \mathbf{tx} . Moreover, by (W2) the stealth excess of \mathbf{tx} does not appear in \mathbf{tx}^* , and thus \mathbf{tx} is not among the transactions (\mathbf{tx}_i) whose stealth excesses were copied, and from which the values $(d_{i,j})_{i,j}$ come. The value \bar{d}_i^* is thus not among $(d_{i,j})_{i,j}$ (except with negligible probability if the reduction chose the same value twice). Since all values are chosen at random by the reduction, the probability that (33) is not satisfied is thus $1/p$, with p the order of the group \mathbb{G} . Whenever (33) is satisfied, the reduction can compute $\log Z$ from (32). \square

6.6 Transaction privacy

We now consider an attacker that passively observes blocks of the underlying blockchain and attempts to deduce information about the transaction graph or the transacted amounts. In Jedusor’s original proposal [Jed16], no guarantee of privacy was given besides hiding transaction amounts, and this was reflected in prior definitions [FOS19, Def. 12]. In the initial version of Mumblewimble one can “disaggregate” transactions [Dev20a], that is, determine which inputs and outputs come from the same original transaction. As a consequence, one could infer how money was being transferred across the network.

GRIN introduced *transaction offsets*, which enable stronger anonymity guarantees by preventing disaggregation. To reflect this, we present a stronger privacy notion than the one in [FOS19]. Our notion is tailored to non-interactive transactions. We stress that our analysis is limited to the cryptographic properties of the scheme and it does not provide network-level privacy guarantees.

Definition. Our scheme provides three basic anonymity guarantees:

- a transaction hides the amounts contained in its inputs and outputs, as well as
- the destination addresses of the inputs and outputs, *except to the receivers of the transaction*;
- in an aggregated transaction, it is not possible to tell which inputs and outputs belonged to the same component transaction *except for what can be deduced via the receiver’s keys and the epochs of the outputs*.

The above are implied by the following simulation-based definition. The adversary has access to an oracle that creates honest users and a challenge oracle that produces transactions, taking as input:

- lists $\mathbf{ref}_i = (ref_{i,j})_j$, for $i \in [\ell]$, where $ref_{i,j}$ is either a tuple $(txo_{i,j}, v_{i,j}, sk_{i,j})$ or an identifier $id_{i,j}$ of a previously generated output, within the same or a previous oracle call (\mathbf{ref}_i thus specifies the list of inputs of the i -th transaction)
- lists $(\mathbf{id}_i, (\hat{v}_i, \hat{\mathbf{pk}}_i, \hat{\chi}_i))$, where $id_{i,j}$ serves as the unique identifier of the corresponding triple value/address/epoch (this specifies the outputs of the i -th transaction)
- supplies s_i and fees f_i (for the i -th transaction)

In the “real” world, the oracle does the following: for all $i \in [\ell]$, it creates \mathbf{tx}_i by running **Send** on the inputs referenced by $(ref_{i,j})_j$: these are either explicit output/value/spending-key triples; or $(ref_{i,j})_j = id_{i,j}$, where $id_{i,j}$ was associated to a previous output argument $(\hat{v}_{i,j}, \hat{\mathbf{pk}}_{i,j}, \hat{\chi}_{i,j})$ for which $\hat{\mathbf{pk}}_{i,j}$ was generated by the “honest-user” oracle; in this case the output $txo_{i,j}$ of the transaction previously generated for $id_{i,j}$ is used together with $\hat{v}_{i,j}$ and $\hat{sk}_{i,j}$, the spending key associated to $\hat{\mathbf{pk}}_{i,j}$. (If the adversary wants to spend outputs for an address it controls, it can provide an explicit output/value/spending-key triple.) The remaining arguments for **Send** are $(\hat{v}_i, \hat{\mathbf{pk}}_i, \hat{\chi}_i)$, s_i and f_i . Using **Agg**, the oracle then aggregates $\mathbf{tx}_1, \dots, \mathbf{tx}_\ell$ to \mathbf{tx} and returns \mathbf{tx} .

Transaction privacy requires that the “real” oracle is indistinguishable from an oracle that creates an aggregate transaction \mathbf{tx} using a *simulator* **Sim**, which is only given:

- a list \mathbf{in} of length the total number of inputs, whose elements are either an output txo_i or an identifier id_i where $id_i \in \hat{\mathbf{id}}$ from $\hat{\mathbf{id}}$ below (that is, it points to an output of a component transaction); moreover, \mathbf{in} is sorted lexicographically; it contains neither values nor keys
- a (sorted) list $\mathbf{out} = (\hat{\mathbf{id}}, (\hat{v}, \hat{\mathbf{pk}}, \hat{\chi}))$ of output identifiers and triples which are either for adversarial keys or which are of the form $(\perp, \perp, \hat{\chi}_i)$ if the output is for an honest user
- the total supply and fee $s, f \in [0, v_{\max}]$ and the number of transactions ℓ .

In addition to the simulated \mathbf{tx} , the simulator also provides a list of the transaction outputs in \mathbf{tx} for the honest users together with their respective identifiers (to be used by the oracle when compiling inputs in a subsequent call of the simulator). The game is formalized in [Figure 7](#). The

<p>Game $\text{TXPRV}_{\text{ACS}, \mathcal{A}}(\lambda, v_{\max})$</p> <hr/> <p>$b \leftarrow_{\\$} \{0, 1\}$ $Pks, Txos := ()$ // hash tables $(pp, \Lambda) \leftarrow \text{Setup}(1^\lambda, v_{\max})$ $b' \leftarrow \mathcal{A}^{\text{Tx}_b}(pp, \Lambda)$ return $(b = b')$</p> <p>$\text{Tx}_0((\text{ref}_i, (\hat{\text{id}}_i, (\hat{v}_i, \hat{pk}_i, \hat{\chi}_i)), s_i, f_i))_{i \in [\ell]}$</p> <hr/> <p>// “real” game // check for double and already used id’s: if $\exists (i, j) \neq (i', j') : \hat{\text{id}}_{i,j} = \hat{\text{id}}_{i',j'}$ or $Txos(\hat{\text{id}}_{i,j}) \neq \varepsilon$ then abort for $i \in [\ell]$ // get ref’s to previous txo’s $\text{in}_i := ()$ for $j \in [\text{ref}_i]$ if $\text{ref}_{i,j} =: id$ // $\text{ref}_{i,j}$ is an id if $Txos(id) =: (txo, v, sk) \neq \varepsilon$ $\text{in}_i := \text{in}_i \parallel (txo, v, sk)$ if $\text{ref}_{i,j} =: (txo, v, sk)$ $\text{in}_i := \text{in}_i \parallel (txo, v, sk)$ else abort $\text{tx}_i \leftarrow \text{Send}(pp, \text{in}_i, (\hat{v}_i, \hat{pk}_i, \hat{\chi}_i), s_i, f_i)$ // store outputs for honest users in Txos: for $j \in [\text{id}_i]$ if $Pks(\hat{pk}_{i,j}) =: (\hat{vk}_{i,j}, \hat{sk}_{i,j}) \neq \varepsilon$ // j-th recipient is honest $\text{txo}_{i,j} \leftarrow \text{Rcv}(pp, \text{tx}_i, \hat{v}_{i,j}, \hat{vk}_{i,j})$ // if multiple transactions are ... // ... found, store the subsequent ... // ... one at each iteration $Txos(\hat{\text{id}}_{i,j}) := (txo_{i,j}, \hat{v}_{i,j}, \hat{sk}_{i,j})$ $\text{tx} \leftarrow \text{Agg}(pp, \text{tx}_1, \dots, \text{tx}_\ell)$ return tx</p>	<p>Oracle $\text{KEYGEN}()$</p> <hr/> <p>$(pk, vk, sk) \leftarrow \text{KeyGen}(pp)$ // add values to table Pks: $Pks(pk) := (vk, sk)$ return pk</p> <p>Oracle $\text{Tx}_1((\text{ref}_i, (\hat{\text{id}}_i, (\hat{v}_i, \hat{pk}_i, \hat{\chi}_i)), s_i, f_i))_{i \in [\ell]}$</p> <hr/> <p>// “simulated” game if $\exists (i, j) \neq (i', j') : \hat{\text{id}}_{i,j} = \hat{\text{id}}_{i',j'}$ or $Txos(\hat{\text{id}}_{i,j}) \neq \varepsilon$ then abort $\text{in} := (); \text{out} := ()$ // simulator gets “aggregated” lists for $i \in [\ell]$ $\text{val}_i := 0; \hat{\text{val}}_i := \sum_j \hat{v}_{i,j}$ // sums of in and out values for $j \in [\text{ref}_i]$ if $\text{ref}_{i,j} =: id$ if $Txos(id) =: (txo, v) \neq \varepsilon$ $\text{in} := \text{in} \parallel txo; \text{val}_i := \text{val}_i + v$ elseif $\exists i' \in [i-1], j : id = \hat{\text{id}}_{i',j'}$ $\text{in} := \text{in} \parallel id; \text{val}_i := \text{val}_i + \hat{v}_{i',j'}$ if $\text{ref}_{i,j} =: (txo, v, sk)$ $\text{in}_i := \text{in}_i \parallel (txo, v, sk); \text{val}_i := \text{val}_i + v$ else abort if $\hat{\text{val}}_i + f_i \neq \text{val}_i + s_i$: return $\text{tx} := \perp$ // tx unbalanced for $j \in [\text{id}_i]$ // remove honest keys and values: if $Pks(\hat{pk}_{i,j}) =: \varepsilon$: out := $\text{out} \parallel (\hat{\text{id}}_{i,j}, (\hat{v}_{i,j}, \hat{pk}_{i,j}, \hat{\chi}_{i,j}))$ else out := $\text{out} \parallel (\hat{\text{id}}_{i,j}, (\perp, \perp, \hat{\chi}_{i,j}))$ $\text{in}' := \text{Sort}(\text{in}); \text{out}' := \text{Sort}(\text{out})$ $s := \sum_i s_i; f := \sum_i f_i$ $(\text{tx}, (\text{id}', \text{txo}')) \leftarrow \text{Sim}(\text{in}', \text{out}', s, f, \ell)$ // store outputs for honest users in Txos: let i', j' such that $\text{id}'_i = \hat{\text{id}}_{i',j'}$ $Txos(\text{id}'_i) := (txo'_{i'}, v_{i',j'})$ return tx</p>
--	---

Fig. 7. Game TXPRV for transaction privacy

adversary gets access to an oracle that either returns real or simulated (aggregated) transactions, and the adversary has to decide which is the case. If in a “real” query, the input and output values are unbalanced, `Send` would return \perp and after aggregation, the oracle response would be \perp . Since the simulator does not get any values, the oracle in the “simulated” world checks balancedness explicitly. The complexity of the formal game mainly stems from bookkeeping.

Remark. Since owners of addresses must be able to identify payments to them, privacy can only hold for addresses for which the adversary does not know the view key. Hence, for outputs that do not belong to honest users, the simulator receives the output addresses (and the corresponding values) created by the adversary.

The simulator does not receive any amounts related to inputs or outputs for honest users; therefore the transaction (which is indistinguishable from a simulated one) does not reveal any such amounts (to a computationally bounded adversary). The simulator does not receive any destination address apart from the ones owned by the adversary; the transaction can therefore not reveal any honest recipients. Finally, the simulator does not get the input/output matching of the constituent transactions, thus the aggregated transaction cannot be “disaggregated”.

Theorem 4. *If the proof systems RaP and PoP are zero-knowledge and if DDH is hard in \mathbb{G} , then MW-NIT satisfies transaction privacy in the random oracle model.*

Proof. We first define the **Sim**. It takes as input: a list \mathbf{in} of length n , whose elements are of the form txo_i or id_i ; an \hat{n} -element list for the outputs with elements of the form $(\hat{id}_i, (\hat{v}_i, \hat{p}k_i, \hat{\chi}_i))$ where possibly $(\hat{v}_i, \hat{p}k_i) = (\perp, \perp)$; total supply and fees s, f and the number of transactions ℓ .

Let $I_{\text{id}} \subseteq [n]$ denote the set of indices for which $in_i = id_i$, that is, the references to outputs of the transaction to be cut through; for $i \in [n] \setminus I_{\text{id}}$, the entries in_i are explicit TXO’s. Further, let \hat{I}_{hon} be the indices $i \in [\hat{n}]$ for which $(\hat{v}_i, \hat{p}k_i) = (\perp, \perp)$, that is, the indices of outputs for honest users.

- (a) for each **output** index $i \in [\hat{n}] \setminus \hat{I}_{\text{hon}}$, sample $\hat{r}_i \leftarrow \mathbb{Z}_p$ and use it to compute $\hat{R}_i = \hat{r}_i G$ and \hat{P}_i (using $\hat{p}k_i = (\hat{A}_i, \hat{B}_i)$) and \hat{C}_i (using \hat{v}_i) as prescribed in **Send**
- (b) for each $i \in \hat{I}_{\text{hon}}$, pick random values $\hat{C}_i, \hat{R}_i, \hat{P}_i \leftarrow \mathbb{G}$
- (c) for each $i \in [\hat{n}]$, simulate the range proofs $\hat{\pi}_i$ for statements \hat{C}_i , and the proofs ρ_i for \hat{R}_i with tag $\hat{C}_i \parallel \hat{\pi}_i \parallel \hat{P}_i \parallel \chi_i$. Let $\widehat{txo}_i := (\hat{C}_i, \hat{\pi}_i, \hat{R}_i, \hat{\rho}_i, \hat{P}_i, \hat{\chi}_i)$
- (d) define the honest outputs via $\mathbf{id}' := (\hat{id})_{i \in \hat{I}_{\text{hon}}}$ and $\mathbf{txo}' := (\widehat{txo}_i)_{i \in \hat{I}_{\text{hon}}}$
- (e) for each **input** index $i \in [n]$, pick a random value $D_i \leftarrow \mathbb{G}$
- (f) for each $i \in I_{\text{id}}$, let k_i be such that $in_i = \hat{id}_{k_i}$ (if there is none, abort); set $P_i := \hat{P}_{k_i}$ with \hat{P}_{k_i} from (a) or (b); simulate ψ_i for the statement $(P_i, D_i, \widehat{txo}_{k_i})$ with \widehat{txo}_{k_i} from (c)
- (g) for each $i \in [n] \setminus I_{\text{id}}$, with in_i containing an explicit output $txo_i := (C_i, \pi_i, R_i, \rho_i, P_i, \chi_i)$, simulate ψ_i for the statement (P_i, D_i, in_i)
- (h) define the cut-through \mathbf{ct} as the (lexicographical) ordering of $(\widehat{txo}_{k_i}, P_i, D_i, \psi_i)_{i \in I_{\text{id}}}$
- (i) define \mathbf{txi}' as the ordering of $(P_i, D_i, \psi_i)_{i \in [n] \setminus I_{\text{id}}}$ (i.e., cut-through terms removed)
- (j) define \mathbf{txo}' as the ordering of $(\widehat{txo}_i)_{i \in [\hat{n}] \setminus \{k_i\}_{i \in I_{\text{id}}}}$ (i.e., cut-through terms removed)
- (k) compute the **kernel**: pick random values $E_2, \dots, E_\ell, X_2, \dots, X_\ell \leftarrow \mathbb{G}$, as well as $t, y \leftarrow \mathbb{Z}_p$
- (l) set $E_1 := \sum \hat{C} - \sum C + (f - s)H - tG - \sum_{i=2}^{\ell} E_i$
and $X_1 := \sum \hat{R} - \sum D - yG - \sum_{i=2}^{\ell} X_i$ (note that we could have $\ell = 1$)
- (m) for each $i \in [\ell]$, simulate σ_i for (E_i, X_i) with tag ε

Set

$$\mathbf{tx} = (\mathbf{txi}', \mathbf{txo}', (\mathbf{ct}, s, f, t, y, \mathbf{E}, \mathbf{X}, \boldsymbol{\sigma})) ,$$

and remove \mathbf{ct} if it is empty, and remove \mathbf{E} and \mathbf{X} if $\ell = 1$. Return $(\mathbf{tx}, (\mathbf{id}', \mathbf{txo}'))$.

We now show that in the random oracle model, assuming indistinguishability of simulated RaP and PoP proofs and the DDH assumption, a real transaction is indistinguishable from a simulated transaction. We start with a real transaction (as returned by Tx_0) and consider a sequence of hybrid games that turns it into a simulated transaction (as returned by Tx_1).

H_0 This is the “real” game, in which the adversary interacts with the oracle Tx_0 . The component transactions are created as described (in [Section 3.2](#)) using the spending keys and they are then aggregated ([Section 3.3](#)).

H_1 In the first hybrid, all RaP proofs π and PoP proofs (cf. [Section 5](#)) ψ, ρ and σ_i are simulated. By the zero-knowledge property of both primitives, this change is indistinguishable from the honest computation of the range proofs and the proofs of possession.

H_2 In this hybrid, when determining the outputs of the honest users to be added to $Txos$, instead of running Rcv , the oracle simply sets $txo_{i,j}$ as the j -th output of the i -th transaction. By the definition of the scheme, this does not change anything.

Note that in this hybrid, the part $a_{i,j}$ of the view key $vk_{i,j}$ of an honest user is only used to compute $(k_{i,j}, q_{i,j}) := H(a_{i,j}R_{i,j})$ when running $Send$ on an honest input; moreover the value $\hat{r}_{i,j}$ for honest outputs is only used to compute $(\hat{k}_{i,j}, \hat{q}_{i,j}) := H(\hat{r}_{i,j}A_{i,j})$ (since H_1). Also, the spending key sk_ℓ is never used (since H_1).

H_3 For every output for an honest user in the generation of the key shares in Equation (3), the argument $\hat{r}_{i,j}A_{i,j}$ of the hash function is replaced by a random value $Z_{i,j} \leftarrow \mathbb{G}$; when the output is later spent then analogously $a_{i,j}\hat{R}_{i,j}$ is replaced by the same $Z_{i,j}$. This is indistinguishable by the DDH assumption, noting that all $\hat{R}_{i,j}$ and $A_{i,j}$ for such outputs are created by the reduction and their logarithms are not used anywhere else in the game (the former because $\rho_{i,j}$ and σ_i are simulated; the latter because of the modification in H_2 and since the address oracle does not reveal any secret keys).

H_4 The game aborts if the adversary queries $Z_{i,j}$ for any i, j to the random oracle. Since the adversary has no information on $Z_{i,j}$, the probability of aborting is negligible. Note that in H_4 , the values $\hat{k}_{i,j}$ and $\hat{q}_{i,j}$, derived from $H(Z_{i,j})$ are uniformly random and independent.

Now that we have showed that the adversary's return value can only change negligibly between H_0 and H_4 , it remains to argue that a transaction generated in H_4 is distributed equivalently to a transaction computed by the simulator, when interacting with the oracle TX_1 . The transactions produced in H_4 and in TX_1 (by the simulator) are both of the form

$$tx = (\mathbf{txi}, \widehat{\mathbf{txo}}, (\mathbf{ct}, s, f, t, y, \mathbf{E}, \mathbf{X}, \boldsymbol{\sigma})) .$$

The supply and fee are defined as $s := \sum s_i$ and fee $f := \sum f_i$ both in H_4 (when running Agg) and in TX_1 . The values t, y are distributed uniformly at random both in H_4 (where $t := \sum_i t_i$ with each $t_i \leftarrow \mathbb{Z}_p$) and by the simulator in (k) (where $t \leftarrow \mathbb{Z}_p$). All range proofs RaP and the proof of possession PoP are simulated both in H_4 and by the simulator. We are left proving that all the group elements in tx follow the same distribution.

Within a query to the transaction oracle, denote by I_{id} the set of indices (i, j) of the inputs given as id values (and not an explicit triple) and let I_{hon} be the set of indices (i, j) for which $\hat{p}k_{i,j}$ is honest, that is, it was generated by the oracle $KEYGEN$.

In a transaction produced in H_4 , for honest outputs, the coin openings $\hat{q}_{i,j}$, for $(i, j) \in I_{hon}$, are uniformly random and independent, and thus so are the elements $\hat{C}_{i,j}$ for $(i, j) \in I_{hon}$, which is how the simulator creates them in (b). Similarly, the values $\hat{k}_{i,j}$, for $(i, j) \in I_{hon}$ are uniformly random and independent and thus so are the corresponding values $\hat{P}_{i,j}$, which is how the the simulator produces them in (b). For adversarial outputs, the values $\hat{C}_{i,j}$ and $\hat{P}_{i,j}$ for $(i, j) \notin I_{hon}$ are computed by the simulator as prescribed by $Send$, which is what H_4 does as well. Finally, all values $D_{i,j}$ and $\hat{R}_{i,j}$ are sampled uniformly random and independently both in H_4 and by the simulator.

By the definition of Agg , we have that in H_4 the excesses E_1, \dots, E_ℓ are uniformly random elements (since all t_i 's are uniformly random values) conditioned on

$$\sum E_i = \sum \sum \hat{C}_{i,j} - \sum \sum C_{i,j} + (f - s)H - (\sum t_i)G .$$

This is exactly how the simulator produces these values in (k) and (l). Likewise, the values X_1, \dots, X_ℓ are uniformly random conditioned on

$$\sum X_i = \sum \sum \hat{R}_{i,j} - \sum \sum D_{i,j} - (\sum y_i)G ,$$

which is exactly how the simulator picks them in (k) and (l).

Finally, the output/input pairs in the cut-through list of the simulator are those that Agg would put there, and the ordering (lexicographic) is the same used in input, outputs, and cut-through lists as per Agg . \square

Acknowledgements. The first author is supported by the Vienna Science and Technology Fund (WWTF) through project VRG18-002. We thank MWC and David Burkett for the fruitful collaboration; we are also grateful to the anonymous reviewers of ASIACRYPT’22 for their helpful comments.

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 *IEEE S&P 2018*, pages 315–334. IEEE, 2018.
- Bur20. David Burkett. Offline transactions in Mumblewimble, 2020. <https://gist.github.com/DavidBurkett/32e33835b03f9101666690b7d6185203>.
- Bur21. David Burkett. One-sided transactions in Mumblewimble (consensus layer), 2021. <https://github.com/DavidBurkett/lips/blob/master/lip-0004.mediawiki>.
- Dev20a. Grin Developers. Grin documentation: Intro, 2020. Available at: <https://github.com/mumblewimble/grin/blob/master/doc/intro.md>.
- Dev20b. Grin Developers. Grin documentation: Mumblewimble, 2020. Available at: <https://docs.grin.mw/wiki/introduction/mumblewimble/mumblewimble/#kernel-offsets>.
- FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. *CRYPTO 2018, Part II, LNCS 10992*, pages 33–62. Springer, 2018.
- FOS19. Georg Fuchsbauer, Michele Orrù, and Yannick Seurin. Aggregate cash systems: A cryptographic investigation of Mumblewimble. *EUROCRYPT 2019, Part I, LNCS 11476*, pages 657–689. Springer, 2019.
- FPS20. Georg Fuchsbauer, Antoine Plouviez, and Yannick Seurin. Blind Schnorr signatures and signed ElGamal encryption in the algebraic group model. *EUROCRYPT 2020, Part II, LNCS 12106*, pages 63–95. Springer, 2020.
- Jed16. Tom Elvis Jedusor. Mumblewimble, 2016. Available at <https://download.wpsoftware.net/bitcoin/wizardry/mumblewimble.txt>.
- Max13a. Gregory Maxwell. CoinJoin: Bitcoin privacy for the real world, August 2013. BitcoinTalk post, <https://bitcointalk.org/index.php?topic=279249.0>.
- Max13b. Gregory Maxwell. Transaction cut-through, August 2013. BitcoinTalk post, <https://bitcointalk.org/index.php?topic=281848.0>.
- Max15. Gregory Maxwell. Confidential Transactions, 2015. Available at https://people.xiph.org/~greg/confidential_values.txt.
- Nak08. Satoshi Nakamoto. Bitcoin: a peer-to-peer electronic cash system, 2008. Available at <http://bitcoin.org/bitcoin.pdf>.
- Poe16. Andrew Poelstra. Mumblewimble, 2016. Available at <https://download.wpsoftware.net/bitcoin/wizardry/mumblewimble.pdf>.
- PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- Seu12. Yannick Seurin. On the exact security of Schnorr-type signatures in the random oracle model. *EUROCRYPT 2012, LNCS 7237*, pages 554–571. Springer, 2012.
- Tod14. Peter Todd. Stealth addresses, 2014. <http://www.mail-archive.com/bitcoin-development@lists.sourceforge.net/msg03613.html>.
- VFV17. Shaileshh Bojja Venkatakrisnan, Giulia C. Fanti, and Pramod Viswanath. Dandelion: Redesigning the bitcoin network for anonymity. *CoRR*, abs/1701.04439, 2017.
- vS13. Nicolas van Saberhagen. CryptoNote v 2.0, 2013. <https://cryptonote.org/whitepaper.pdf>.
- Wag02. David Wagner. A generalized birthday problem. *CRYPTO 2002, LNCS 2442*, pages 288–303. Springer, 2002.
- Yu20. Gary Yu. Mumblewimble non-interactive transaction scheme. Cryptology ePrint Archive, Report 2020/1064, 2020. <https://ia.cr/2020/1064>.