Haze and Daze: Compliant Privacy Mixers

Stanislaw Baranski
Gdansk University of Technology
stanislaw.baranski@pg.edu.pl

Maya Dotan
The Hebrew University in Jerusalem, Israel
mayadotan@mail.huji.ac.il

Ayelet Lotem
The Hebrew University in Jerusalem, Israel
ayelet.lotem@mail.huji.ac.il

Margarita Vald
Reichman University, Israel
margarita.vald@cs.tau.ac.il

ABSTRACT

Blockchains enable mutually distrustful parties to perform financial operations in a trustless, decentralized, publicly-verifiable environment. Blockchains typically offer little privacy, and thus motivated the construction of *privacy mixers*, a solution to make funds untraceable. Privacy mixers concern regulators due to their increasing use by bad actors to illegally conceal the origin of funds. Consequently, Tornado Cash, the largest privacy mixer to date, is sanctioned by large portions of the Ethereum network.

In this work, we propose Haze and Daze, two privacy mixers that mitigate the undesired abuse of privacy mixers for illicit activities. Haze guarantees users' privacy together with *compliance*, i.e., funds can be withdrawn as long as they were deposited from a non-banned address, without revealing any information on the matching deposit. This means that once a user is flagged as non-compliant, their funds can no longer exit the mixer. However, this leads to a race condition for bad actors to withdraw funds before becoming flagged as unlawful in the *banned-addresses list*. Thus, we introduce Daze, a second mixer protocol that, in addition to compliance, enables retroactive de-anonymization of transactions of non-compliant users, making the mixer fruitless for them. To maintain privacy of compliant users, the de-anonymization procedure is performed by a committee, with privacy guaranteed as long as at least one of the committee members is honest. Moreover, Haze and Daze have an optional feature for non-compliant funds to be released from the mixer to some predetermined entity.

We empirically evaluate our solution in a proof-of-concept system, demonstrating gas consumption for each deposit and withdrawal that is comparable to Tornado Cash for compliant users, both for Haze and Daze. To the best of our knowledge, our solution is the first to guarantee compliance and privacy on the blockchain (on-chain) that is implemented via a smart contract.

1 INTRODUCTION

Blockchains and privacy. Blockchains are decentralized, publicly verifiable, and distributed append-only immutable ledgers that allow mutually distrustful parties to maintain a common state. Bitcoin [29] is the first blockchain system to go live, enabling parties to engage in money transfers using the native currency of the blockchain. Ethereum [46] is a blockchain platform that enables users, in addition to simple money transfers, to perform more complex operations in the form of a smart contract. A smart contract can be any program implemented on the blockchain. The state of the smart contract is maintained as part of the state of the blockchain. While Bitcoin and Ethereum offer users pseudonymity, in both blockchains funds are traceable. Over the years, there have been several attempts at adding various flavors of privacy to blockchains [5, 18, 34, 44]. One such flavor is untraceability of funds. A popular way to make funds untraceable in blockchains is through the use of privacy mixers [4, 17, 43]. A widely used privacy mixer in practice is Tornado Cash [31], which is decentralized and implemented via a smart contract on the blockchain. The untraceability property provided by privacy mixers aided a growing phenomenon of money being laundered via such systems. For instance, the Ethereum address $0x\ldots383E2f96$ which belongs to the hacker group Lazarus of North Korea [40] used Tornado Cash to launder millions of dollars in stolen funds. The U.S. Department of Treasury publishes the "Specially Designated Nationals And Blocked Persons List (SDN)" [41] that contains addresses of persons that the U.S. prohibits dealing with, as part as the OFAC list (Office of Foreign Assets Control). This list contains, amongst other things, blockchain addresses suspected to be involved in various types of illegal activity. In August 2022, following the Lazarus incident, the list was updated to include Tornado Cash [42]. This act changed the patterns of block-inclusion for Tornado Cash transactions by miners/validators. Today about a third of validators in the Ethereum network censor Tornado Cash transactions [24]. Currently, such a list is maintained on the Ethereum blockchain by Chainalysis [7]. The extensive usage of privacy mixers to move illicit funds and the addition of Tornado Cash to OFAC's list emphasizes the need for solutions that provide privacy only to "good" users, but do not allow access to the system to entities that do not comply with the policy. In this paper we refer to the problem of preventing addresses from OFAC's list from transferring funds through a privacy mixer as the "compliance" problem. A *compliant privacy mixer* is therefore a mixer that preserves privacy in the sense of fund untraceability for honest users, ones that are not on the banned-addresses list, and does not enable the release of funds deposited from banned addresses on the list, even if the address only becomes banned after successfully depositing funds to the mixer. To construct such mixers one needs to take into account the dynamic nature of the banned-addresses list that is constantly updated to include new addresses. For this reason, a compliant privacy mixer must verify that, at the time of withdrawal, the funds being withdrawn did not originate from a banned address. However, this requirement that is

essential for achieving compliance together with privacy induces a non-trivial combination, as at the time of withdrawal the mixer must be oblivious to the origin of the funds. Burleson et al.[6] were the first to introduce the question of compliant privacy mixers, and discuss at a high level the desired features of such a solution. However formal security definitions and implemented solutions are still missing. This raises the following question: *How can we construct a practical privacy mixer with provable compliance?*

1.1 Our contribution

In this work we construct the first compliant privacy mixer, Haze, that guarantees the following security properties: (1) correctness - compliant users can always withdraw their funds, (2) soundness - funds cannot be double spent, (3) privacy in the form of deposit-withdrawal unlinkability, and (4) compliance - users on the banned-addresses list cannot withdraw funds from the mixer. Moreover, we formalize these properties and cast them into general security definitions. We note that the mixer is considered honest, due to the fact that its code is publicly deployed and immutable on the blockchain, and can be verified for correctness prior to usage. However, illicit users may manage to exit the mixer prior to being inserted to the banned-addresses list, gaining privacy. To handle this undesirable situation we introduce Daze, a protocol that provides retroactive de-anonymization of transactions of non-compliant users, in addition to the guarantees of Haze. The de-anonymization is in the sense that the deposit and withdrawal transactions become publicly linked, which revokes the privacy of the user in the mixer. The de-anonymization in Daze is implemented as a committee based distributed procedure, such that privacy of compliant users is maintained as long as at least one committee member is honest. We further show how Haze and Daze can be extended to allow funds deposited to the mixer from banned addresses (i.e., funds that cannot be withdrawn) to be released to a predetermined trusted entity. This enables, for example, stolen funds to be returned to their rightful owner instead of being locked forever inside the mixer. We implement Haze and Daze and evaluate their performance, showing comparable costs to the most prominent privacy mixer, Tornado Cash.

Formalization of compliance for mixers. In order to formalize compliance, we consider an idealized compliant ledger. In this ledger, deposits that become non-compliant are "removed" from the ledger, alongside the funds that are associated with them. This implies that mixer protocols in the idealized compliant ledger are compliant by default, as non-compliant deposits are not inside the mixer and hence cannot be withdrawn. Informally, our compliance definition is the following: a mixer protocol is compliant if it behaves indistinguishably in the idealized compliant ledger and the standard (append-only) ledger. Concretely, any accepted withdrawal transaction by the mixer is also accepted if the ledger is replaced with the idealized compliant ledger and vice versa. This definition coincides with the intuitive notion of compliance - illicit funds can't go through the mixer.

Overview of our constructions. Both Haze and Daze are comprised of two entities, a *user* and a *mixer*. The mixer is a smart contract implemented on the blockchain, and the user is a client run locally by any user wishing to interact with the smart contract of the mixer. Users interact with the mixer by depositing and withdrawing funds, by means of transactions on the blockchain.

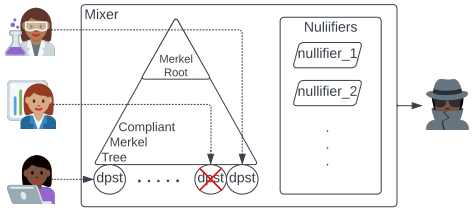

Figure 1: Haze**. Deposits from non-compliant addresses are removed from the tree (the red X). The removal is triggered by a withdrawal transaction and done by zeroing their leaf value and updating the hashes at the nodes along the path from this leaf to the root. The privacy guarantee is that a withdrawal cannot be linked to its corresponding deposit.**

Similarly to Tornado Cash, Haze and Daze utilize Merkle trees and zero-knowledge proofs. Deposits are made by submitting a leaf to the Merkle tree maintained by the mixer. Withdrawals are made by users by creating a zero-knowledge proof that asserts that they have an unspent deposit from a *compliant address* in the mixer. The proof is sent to the mixer alongside a *nullifier*, where both are based on some secret information known only to the depositor. The nullifier is then stored in the smart contract and is used to ensure funds cannot be double-spent.

In Haze, the proof is constructed and verified with respect to the *compliant* Merkle tree, a tree where leaves associated with deposits from non-compliant addresses are removed. A non-compliant address is an address on the banned-addresses list. The banned-addresses list is implemented as a smart contract on the blockchain, and maintained by a list maintainer. At withdrawal, the mixer queries this list in order to keep the compliant Merkle tree up-to-date, and if an address of a deposit had become non-compliant (banned), it removes the corresponding leaf from the Merkle tree and updates the relevant path in the tree accordingly. Funds belonging to removed leaves (equivalently, funds deposited from non-compliant addresses) are unrecoverable to the depositor, as it no longer can generate a successfully verifiable zero-knowledge proof to withdraw these funds, see fig. 1. Therefore, Haze ensures that funds can never be withdrawn once the address they deposited from becomes non-compliant. Moreover, we emphasize that our technique for achieving compliance via removing leaves in the Merkle tree that are associated with non-compliant deposits can be applied with respect to any general compliance policy that can be checked against a deposit and not only policies defined by addresses. Thus, Haze captures a richer family of compliance policies.

Haze*'s Implementation.* We implement both Haze's client (i.e., "user") in JavaScript and server (i.e., "mixer") in Solidity and evaluate the protocol's performance and gas consumption, demonstrating:

- No overhead at deposit. The gas consumption of a deposit transaction and the running time of the user are identical to Tornado Cash, i.e., $\sim$ 1M gas and $\sim$ 0.04 seconds, respectively.
- $\sim$ 1M gas consumption for the Merkle tree update per each newly non-compliant address that is associated with a leaf in the tree. The gas for the this update operation is paid by the withdrawal that triggers this tree update (i.e., the first withdrawal transaction since this address entered the banned-addresses list), and funded by non-compliant users, as explained below.
- Negligible running time overhead at withdrawal. Concretely, 0.24 seconds user running time amortized per Merkle tree node, with a $\sim$ 0.005 seconds difference from Tornado Cash. The time difference stems from fetching the banned-addresses list. The gas consumption per withdrawal is $\sim$ 0.31M for the zero-knowledge proof and nullifier validation (as in Tornado Cash), plus $\sim$ 1M gas per newly non-compliant address that requires an update of the Merkle tree, as specified above.

The gas consumption per Merkle tree update is identical to the gas consumption of a deposit transaction, as both only require updating the relevant path in the tree. Therefore, a reasonable solution to cover the additional gas consumption of withdrawal due to compliance is to charge an extra fee per deposit proportional to the gas consumption of a single deposit. The legitimacy for the fee is by having each user cover not only the cost of the deposit itself, but also the cost of preserving compliance in case its address becomes non-compliant. Concretely, the cost of withdrawal is comprised of the gas consumption of the zero-knowledge and nullifier verification and an additive factor proportional to the number of newly non-compliant addresses associated with leaves in the tree. The extra cost paid at withdrawal for the Merkle tree updates is refunded to the withdrawer by the mixer, and funded by the extra fee charged with each deposit transaction. The cost overhead in the deposit is refunded to compliant users, as described in appendix C, making this a type of limited-time escrow. This approach makes a deposit transaction in Haze cost at most twice compared to Tornado Cash for non-compliant users, while maintaining the cost of a withdrawal identical to Tornado Cash. Overall, together with this feature, our protocol does not incur, for compliant users, a cost overhead over Tornado Cash. The full implementation and experimental results are detailed in section 5.

Daze: De-anonymizing non-compliant users. Due to the strong privacy guarantees of Haze, users that manage to withdraw funds prior to becoming non-compliant succeed in concealing the trace of their illicit funds. To handle this risk, we construct a solution that enables the banned-addresses list maintainer to publicly revoke the privacy of users that become non-compliant. For this we introduce Daze, a mixer protocol which supports, in addition to correctness, soundness, privacy and compliance, the retroactive de-anonymization of transactions of illicit users, even if they become banned only after withdrawing their funds from the mixer. Daze differs from Haze in the way it achieves compliance. Instead of removing non-compliant deposits from the Merkel tree, it provides compliance as follows: when depositing funds to the mixer, a user provides an encryption of its nullifier, encrypted with the public key of the banned-addresses list maintainer, together with a correctness proof. Later, if an address becomes non-compliant, the list maintainer updates the banned-addresses list with the user's address and the nullifier in plaintext. Upon a withdrawal request, the mixer compares the submitted nullifier against the nullifiers that appear on the banned-nullifier set, making funds of deposits from non-compliant addresses non-withdrawable. Publishing the nullifiers associated with non-compliant addresses publicly links withdrawals of illicit funds to their deposits, thus revoking the privacy of the withdrawal.

In Daze, preventing abuse by the banned-address list maintainer is crucial, to prevent off-chain de-anonymization of compliant users. We therefore suggest instantiating the de-anonymization process as a committee using threshold encryption. Ideally, the committee should contain a designated set of parties that are in charge of maintaining the mixer (i.e. an on-chain governance module [39]), as well as two off-chain separate entities such as the court and an executive authority (police, etc.). The governance module is incentivised to comply with de-anonymization of transactions of users that meet some threshold of bad behaviour, as not participating in de-anonymizing can render the entire mixer illegal (as in the case of Tornado Cash). Once a proof of user non-compliance is presented by the off-chain entities the on-chain entities contribute, publicly on chain, their part to the de-anonymization of this user's transactions. This guarantees the privacy of compliant users as long as at least one of the committee members is honest. We emphasize that even if all committee members are corrupted, they still cannot steal funds from the mixer.

We implemented Daze, and found that it is comparable to Tornado Cash it terms of running time and gas cost of withdrawal, and with a fixed overhead per deposit of $\sim$ 350K gas and $\sim$ 1.35 seconds. In Daze all users pay the "price of compliance", in contrast to Haze, in which the cost overhead can be relayed to non-compliant depositors. This is detailed in section 5.

Releasing non-compliant funds. Our construction ensures that funds deposited into Haze and Daze from addresses on the banned-addresses list cannot be withdrawn. However, this might mean that stolen funds deposited into them are locked forever in the smart contract and cannot be returned to their rightful owners. It is desirable to be able to release these funds to some predetermined entity, possibly a hard-coded address of a law enforcement agency that can then redistribute these funds. This entity can also implement a dispute resolution mechanism for individuals that claim to be wrongly placed on the list. In Daze, the mixer can count non-compliant funds by counting the number of "banned" nullifiers on the list that have not yet withdrawn their funds from it. This way, Daze can release these funds to this entity at any desired period during the life-time of the mixer. In Haze, due to its privacy guarantees, determining whether a deposit made by a non-compliant user has been withdrawn is infeasible. Therefore, non-compliant funds that have been block by Haze can be released only at the end of the mixer's life-cycle.

1.2 Related Work

Flavors of privacy over the blockchain. Prior to this paper, extensive work has been done towards ensuring transaction privacy on the blockchain, e.g., Hopwood *et al.*, Sasson *et al.* etc. [2, 5, 18, 28, 32, 34, 44] are blockchain solutions that utilize cryptography to anonymize transactions. Most of them utilize Merkle trees and nullifiers, as in our construction. However, these solutions tend to be slow and expensive deeming them less popular for use in practice. Other works, such as [10, 26, 33] by Malatova *et al*, Roos *et al.*, etc. provide privacy to layer 2 systems implemented on top of the blockchain. They however do not address the privacy of on-chain transactions. Another desired flavor of privacy is untraceability of funds over the blockchain, that is commonly achieved through the use of privacy mixers. These are sometimes referred to as "add-on" privacy solutions that derive privacy by mixing a user's funds with many other funds. Mixers can be either centralized, see Bonneau *et al.*, Heilamn *et al.*, etc. [4, 17, 43], and depend on a trusted central entity or decentralized, see Meiklejohn *et el.*, Pertsev *et al.*, Bunz *et al.* etc. [5, 27, 31] which means that the functionality is implemented via an on-chain smart contract. Several papers quantify the privacy achieved by existing systems. For instance, Wu *et al.* [47] and Wang *et al.* [45]. However, none of these systems provide any guarantees of compliance. Moreover, some of these systems have been prone to abuse by money launderers, as mentioned above.

Compliance with privacy over the blockchain. Several papers have studied the intersection of privacy and compliance in the blockchain setting. In particular, Goldwasser *et al.* [15] proposes a protocol that enables to prove that specific regulations are being adhered to while maintaining secrecy of recorded data. Burleson *et al.* [6] were the first to introduce the question of compliance for privacy mixer, and state it in the sense of a banned-addresses list. They proposed a solution based on exclusion proof, where the zero-knowledge proof is extended to include a proof of exclusion from the banned-addresses list. However they do not provide an implementation, or definitions of the desired properties. A concurrent work of Soleimani [35], called privacy pools, proposed a preliminary implementation of the exclusion proof approach. Their notion of compliance is however weaker as a user may prove exclusion with respect to any banned-addresses list of their choice, including the empty one, and thus funds from non-compliant addresses are not blocked from exiting the mixer. In [35] an additional Merkle tree is maintained and used for the exclusion proof, in contrast to our solution that uses a single Merkle tree and rather manipulates it directly. Our solution provides a comparable cost to Tornado Cash for compliant users, and punishes non-compliant users by charging them an extra fee. A monetary punishment of non-compliant users might be a desired feature on its own. Tomescu *et al.* [36] suggest UTT, a system for decentralized ecash with accountable privacy. They consider a notion of compliance, based on a privacy budget, which limits the volume of private transactions a user can perform. This differs from our setting, where we aim to block, indefinitely, funds originating from misbehaved users and guarantee privacy to all other users.

Paper Organization

The rest of this paper is organized as follows. Preliminary terminology and definitions appear in section 2. Our protocols for compliant privacy mixers in section 3. Details on integrating our protocols with the blockchain in section 4. The implemented system and empirical evaluation in section 5. Conclusions in section 6.

2 PRELIMINARIES

We use standard definitions for functions being *negligible* with respect to a system parameter λ called the *security parameter*, denoted $\mathrm{negl}(\lambda)$; similarly for *polynomial*, where ppt stands for *probabilistic polynomial time* in λ. See definitions in [23].

In the following we establish definitions and terminology required for the rest of the paper.

Hash functions. We call an efficiently computable family of keyed functions $\mathcal{H} = \{H_s : \{0,1\}^* \to \{0,1\}^t\}_{t=t(\lambda), s\in\{0,1\}^\lambda, \lambda\in\mathbb{N}}$ *collision resistant* hash functions, if for every ppt adversary $\mathcal{A}$, and any λ, a uniformly random function $H_s \in \mathcal{H}$ satisfies that $\mathcal{A}$ cannot find $x \neq x'$ s.t $H(x) = H(x')$, except with negligible probability.

A CPA-secure PKE scheme. for public key encryption (PKE) scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ and its properties of *correctness*, *CPA-indistinguishability experiment* against an adversary $\mathcal{A}$ denoted $\mathsf{EXP}^{cpa}_{\mathcal{A},\mathsf{pk}}(\lambda)$, and *CPA-security*. See the formal definitions in appendix A.

Commitment schemes. Commitment schemes are protocols which enable a party, known as the committer C, to commit himself to a value while keeping it secret from the (potentially cheating) receiver, R. This property is known as hiding. Additionally, upon receiving the commitment from C, R is ensured that even if C cheated, there is at most one value that C can decommit to during a later, decommitment phase (binding). We formally define a *secure commitment scheme* in appendix A. It is known how to construct a non-interactive, perfectly binding commitment scheme from any one-way permutation [3]. Pedersen [30] constructed a computationally binding and unconditionally hiding scheme based on the discrete logarithm problem.

Merkle trees. A Merkle tree $\mathcal{T}$ is a complete binary tree equipped with a collision-resistant hash function H and computed on n leaves having values $[v_1], \ldots, [v_n]$, and the value of each internal node is $H(a||b)$ where a and b are the values of its two children; (We assume n is a power of 2; if not, we fix a zero value for the missing leaves). We use a standard notion where each leaf value $[v]_i$ is a hash of some cleartext data d_i. An *authentication path* $O(\mathcal{T}, \ell)$ of a leaf with position ℓ in $\mathcal{T}$ is made up of the values of all "sibling" nodes on the path from leaf ℓ to the root denoted $R_{\mathcal{T}}$, as well as $[v]_\ell$ itself and d_ℓ. We use computationally binding and unconditionally hiding commitment scheme for the leaves value, and a collision resistant hash to compute the internal nodes of $\mathcal{T}$.

Zero-knowledge succinct non-interactive arguments of knowledge. Let R be a polynomial time decidable binary relation over pairs (ϕ, w) where ϕ is the statement and w the witness. An efficient-prover publicly verifiable non-interactive argument for R is a tuple of ppt algorithms $\Phi = (\mathsf{ZK.Setup}, \mathsf{ZK.Prove}, \mathsf{ZK.Ver}, \mathsf{ZK.Sim})$

satisfying perfect completeness, perfect zero-knowledge, computational zero-knowledge soundness, producing a proof of polynomial size in λ and having verification ZK.Ver polynomial in λ and $|\phi|$. See the formal definitions in appendix A.

3 COMPLIANT PRIVACY MIXERS HAZE & DAZE

In this section we present our two protocols for privacy mixers with compliance. Our protocols achieve correctness, soundness, privacy, and compliance.[1] First, Haze is presented in section 3.1, and fig. 4 with its security analysis. Then, in section 3.2, and fig. 6 we present Daze, the mixer protocol which supports retroactive de-anonymization of non-compliant users. Formal definitions for the properties achieved by our protocols and listed in this section are available in appendix B.

3.1 A Compliant Privacy Mixer Haze

In this section we formally describe our protocol Haze. We enhance the Tornado Cash protocol [31] to obtain compliance in the sense of preventing withdrawals of funds that belong to non-compliant deposits, without exposing information on the deposit being withdrawn, thus maintaining privacy of the user. A deposit is non-compliant if it was deposited from an address that has become non-compliant in the duration leading to the withdrawal attempt. The main difference between Haze and [31] is in the withdrawal phase, where we first manipulate the Merkle tree to remove leaves corresponding to non-compliant deposits. This treatment guarantees that even if deposits become non-compliant after entrance to the mixer, they cannot be withdrawn. More formally,

The protocol Haze = (deposit, withdraw) consists of a pair of protocols, deposit and withdrawal where any user Usr can communicate with Srv to perform the following functionality:

- deposit enables users to deposit money to Srv. Depositing is done by user Usr generating a deposit transaction of a fixed amount and communicating it to Srv.
- withdraw enables users to withdraw deposited funds from Srv. Withdrawing is done by user Usr generating a withdraw transaction (of the same fixed amount) using undisclosed data generated by Usr during the deposit phase and communicating it to Srv.

Functionality $\mathcal{F}_{bb}$ for communication. Communication between users and Srv is done via the *Bulletin board* $\mathcal{F}_{bb}$, a functionality that models the blockchain. $\mathcal{F}_{bb}$ supports the following requests from any entity in the (blockchain) system: Write a message, and Read written messages. Written messages are stored in an append-only linked list, where each list node is a tuple containing: its index in $\mathcal{F}_{bb}$, sender's and recipient's address, and the message itself. Similarly to the blockchain, entities in the system can generate and possess multiple addresses, where each address is unique (w.h.p), and there is no linkability between the addresses and the identity of the user. Read requests return the content of the linked list at the request time (i.e., up to the node with the most recent index). When a user executes either deposit or withdraw, the generated transaction is written to $\mathcal{F}_{bb}$ with the address of Srv as the recipient's address, which is hard-coded within the protocol Haze. The sender's address is recorded as well. Users can access $\mathcal{F}_{bb}$ from any address they own, but cannot use other entities' addresses (as in the blockchain, sending a message from an address requires signing the message with the secret key associated with the address being used). See fig. 2 for formal details.

[1] We note that dishonesty of Srv is not a concern in the blockchain setting, since its code is *public and immutable* (i.e., deployed on the blockchain) and thus correctness of Srv can be verified prior to the usage of the protocol.

Functionality $\mathcal{F}_{bb}$ proceeds as follows: Set $index = 0$

Upload. Upon receiving $(\text{Write}, \text{msg}, \text{address}_B)$ from some address address_A, store the tuple $(index, \text{msg}, (\text{address}_A, \text{address}_B))$, output $(index, \text{msg}, \text{address}_A)$ to address_B, and set $index++$.

Read. On Read request from a party return all stored records in $\mathcal{F}_{bb}$.

Figure 2: The bulletin-board functionality

functionality $\mathcal{F}_{ban}^{Q}$ for non-compliant addresses. A privacy mixer is required to reject deposits and withdrawals of funds associated with banned addresses. Concretely, since the banned addresses are dynamic, in the sense that new banned addresses are added from time to time, we consider an interactive *banned-addresses functionality* $\mathcal{F}_{ban}^{Q}$, that stores the banned addresses (along with optional metadata), and is updated only by a predefined entity with a fixed address Q. We call Q the banned-addresses list maintainer. The banned-addresses list can be read by any entity in the system, and in particular, by users and Srv in Haze. The decision on which addresses are updated in $\mathcal{F}_{ban}^{Q}$ is left outside the model. See fig. 3 for formal details.

Functionality $\mathcal{F}_{ban}^{Q}$ parameterized on address Q acts as follows:

Update. Upon receiving $(\text{Ban}, \text{address}_A, \text{data})$ from address Q, record $(\text{address}_A, \text{data})$.

Read. On Read request from any party, return all stored records in $\mathcal{F}_{ban}^{Q}$.

Figure 3: The banned-addresses functionality. Records are pairs of banned address, together with an (optional) field containing data related to the address.

The protocol Haze. We present our privacy mixer protocol with compliance Haze = (deposit, withdraw) in fig. 4. Haze modifies the Tornado Cash protocol to obtain compliance in the sense that deposits from banned addresses cannot be withdrawn. That is, Haze operates in the presence of $\mathcal{F}_{ban}^{Q}$ and rejects withdrawal of funds that were deposited from addresses that are recorded in $\mathcal{F}_{ban}^{Q}$ at the moment of withdrawal. This guarantees blocking banned addresses,

that were not necessarily in $\mathcal{F}_{ban}^{\mathcal{Q}}$ when the deposit transaction communicated to Srv. The protocols deposit and withdraw are non-interactive in the sense that users communicate with Srv, but not vice versa, and Srv only performs Read requests to the $\mathcal{F}_{bb}$ and $\mathcal{F}_{ban}^{\mathcal{Q}}$ functionalities. The communication to Srv (i.e., deposit and withdrawal transactions) is done by users sending a Write request to $\mathcal{F}_{bb}$ with the transaction and $\text{address}_{\text{Srv}}$ being the recipient's address.

The difference between Haze and Tornado cash resides in the withdraw protocol. In the original Tornado Cash withdrawal protocol [31], when a user Usr wants to withdraw funds that it deposited in a deposit transaction dtxn, it proceeds as follows: (1) computes the root $R_{\mathcal{T}}$ of a Merkle tree $\mathcal{T}$, where the leaves of $\mathcal{T}$ are all deposit transactions submitted to Srv so far, where ℓ is the leaf associated with the deposit dtxn. Then, (2) Usr computes $O(\mathcal{T}, \ell)$, the authentication path of ℓ in $\mathcal{T}$ as defined in section 2. Next, (3) it computes a hash, called nullifier, over part of the randomness used to generate dtxn. Finally, (4) Usr produces a proof that it "knows" the authentication path for one of the leaves in $\mathcal{T}$ that has not been previously withdrawn. The proof in (4) is generated using a ZK-SNARK scheme Φ for a polynomial time decidable binary relation R, where the statement is ($R_{\mathcal{T}}$, nullifier) and the witness is $(\text{randomnees}_{\text{dtxn}}, \ell, O(\mathcal{T}, \ell))$. The withdrawal transaction submitted by Usr consists of (nullifier, proof).

Upon receiving the withdrawal request, Srv fetches its locally stored $\mathcal{T}$ and verifies the proof wrt its root and the received nullifier (in addition to the nullifier uniqueness assertion). In our withdrawal protocol, we modify $\mathcal{T}$ and nullify the leaves that correspond to deposit transactions associated with an address that appears in $\mathcal{F}_{ban}^{\mathcal{Q}}$. Consequently, if the ZK-SNARK proof verifies, it guarantees that the deposit transaction it withdraws is not from an address in $\mathcal{F}_{ban}^{\mathcal{Q}}$, as those do not appear in $\mathcal{T}$ anymore.

The formal description of our protocol Haze appears in fig. 4.

Our protocols deposit and withdraw in Haze provide the same time complexity as Tornado Cash except for an additive factor in withdraw for Srv, that is, for a security parameter λ:

- deposit has Usr and Srv time complexity of $\text{poly}(\lambda)$ and $\text{poly}(\lambda) \cdot O(\log(n))$, respectively.
- withdraw has a Usr time complexity of $\text{poly}(\lambda) \cdot O(n \cdot \log(n))$, and an additive factor of $\Delta \cdot \text{poly}(\lambda) \cdot O(\log(n))$ on the Srv side, compared to Tornado Cash.

where n is the maximal number of leaves in $\mathcal{T}$, and Δ is the number of added addresses to $\mathcal{F}_{ban}^{\mathcal{Q}}$, since the previous withdrawal transaction, that are associated with leaves in $\mathcal{T}$. See section 5 for performance measurements of withdraw and deposit.

Our protocol Haze provides correctness, soundness, privacy, and compliance in the following sense. Detailed formalization of these properties appears in appendix B.

Correctness is in the sense that any deposited funds can be withdrawn (once) as long as the matching deposit transaction is compliant at the time of the withdrawal, i.e., the withdrawn funds were not deposited from an address in $\mathcal{F}_{ban}^{\mathcal{Q}}$. Correctness stems from the collision resistance of $H(\cdot)$ together with the completeness property of Φ, and k being randomly sampled. Concretely, a valid deposit transaction $\text{Com}(k||r)$ is a leaf in $\mathcal{T}$ as long as it is not from an address in $\mathcal{F}_{ban}^{\mathcal{Q}}$. Therefore, on input (k, r) the withdraw protocol produces an accepting zero-knowledge proof π and a unique nullifier h.

Soundness is in the sense that a user cannot withdraw funds that it did not deposit. Soundness stems from the hiding of C, the collision resistance of $H(\cdot)$, and the computational knowledge soundness of Φ. That is, a user that produces a valid proof for the instance $(R_{\mathcal{T}}, H(k))$, without possessing (k, r) for one of the leaves, can be used to either break the hiding property of C or the collision resistance of $H(\cdot)$. The reduction uses the knowledge extractor χ_{Usr}, guaranteed by the knowledge soundness of Φ, to extract the witness $(k, r, \ell, O(\mathcal{T}, \ell))$ with non-negligible probability. Moreover, the binding property of C prevents double spending of the deposited funds.

Privacy is in the sense that a withdrawal cannot be linked to any non withdrawn deposit. Privacy stems from the hiding property of the commitment scheme C and the zero-knowledge property of Φ. Concretely, the zero-knowledge proof guarantees to hide $(k, r, \ell, O(\mathcal{T}, \ell))$ and the hiding property of C guarantees that $H(k)$ can be linked to $\text{Com}(k||r)$ w.p at most negligibly larger than a random guess.

Compliance is in the sense that funds deposited from an addresses in $\mathcal{F}_{ban}^{\mathcal{Q}}$ cannot be withdrawn. This follows immediately from the construction as leaves associated with deposit transactions from banned addresses are zeroed and do not appear in $\mathcal{T}$. Therefore, depositors from addresses in $\mathcal{F}_{ban}^{\mathcal{Q}}$ cannot produce an accepting nullifier and zero-knowledge proof to withdraw these funds.

3.2 Daze: De-anonymizing Transactions of Non-compliant Users

As shown in section 3.1, Haze maintains the privacy of users, even in the event they become non-compliant after withdrawing their funds. In such cases, a user successfully launders money and cannot be traced. This strong post-withdrawal privacy guarantee for deposits that became non-compliant might not be acceptable in practice. For instance, funds of non-compliant users are not traceable even in the event of a court order. Due to this concern, we suggest an alternative mixer protocol, called Daze, to de-anonymize withdrawals of funds that originated from non-compliant addresses, even if these funds were successfully withdrawn from the mixer.

More formally, the protocol presented in fig. 6 denoted by Daze, relies on banned-addresses list maintainer $\mathcal{Q}$ that in addition to uploading the non-compliant addresses to $\mathcal{F}_{ban}^{\mathcal{Q}}$, also provides for each such address a data field (see fig. 3) that enables to publicly disclose the trace of funds deposited from this address to the mixer. Concretely, it enables *linking the withdrawal to the address of the deposit*, and hence revoking privacy. The privacy and compliance of the protocol rely on $\mathcal{Q}$ to perform the data extraction (1) correctly and (2) only on deposits to the mixer that are from non-compliant addresses. To enforce (2), we propose a distributed realization of the de-anonymization capability of $\mathcal{Q}$ in section 4.1. The protocol

Common parameters: A security parameter λ, a function H sampled uniformly at random from a collision-resistant hash function family $\mathcal{H}$, a secure commitment scheme $\mathcal{C} = (\mathsf{Com}, \mathsf{Dcom})$, and a ZK-SNARK scheme $\Phi = (\mathsf{ZK.Setup}, \mathsf{ZK.Prove}, \mathsf{ZK.Ver}, \mathsf{ZK.Sim})$ for relation R as defined above.

Parties and addresses: A mixer Srv with public address $\mathsf{address}_{\mathsf{Srv}}$ and a user Usr with some addresses $\mathsf{address}_{\mathsf{Usr}_1}$ and $\mathsf{address}_{\mathsf{Usr}_2}$.

Storage: Srv locally stores a full binary Merkle tree $\mathcal{T}$ on n leaves, all initialized to zero, the location of the next available leaf $next = 0$, and, initially empty, nullifier set $\mathsf{S}_\emptyset$. Each leaf in $\mathcal{T}$ is associated with an address, initially set to $\perp$.

Trusted setup: Bulletin-board functionality $\mathcal{F}_{bb}$, and Banned-addresses Functionality $\mathcal{F}_{ban}^{\mathcal{Q}}$. A common reference string σ produced by running ZK.Setup.

Deposit: deposit is executed by user Usr from some address $\mathsf{address}_{\mathsf{Usr}_1}$ and mixer Srv, as follows:

(1) **User:** samples uniformly at random $k, r \leftarrow \{0,1\}^{t(\lambda)}$ for some polynomial $t(\cdot)$, computes $c = \mathsf{Com}(k||r)$, and sends $(\mathsf{Write}, c, \mathsf{address}_{\mathsf{Srv}})$ to $\mathcal{F}_{bb}$. We call c a *deposit transaction.*

(2) **Mixer:** Upon receiving $(index, c, \mathsf{address}_{\mathsf{Usr}_1})$ from $\mathcal{F}_{bb}$, perform the following steps:

(a) Check that c is not already a leaf in $\mathcal{T}$ (else output 0).

(b) Invoke the subroutine in fig. 5 on $\mathsf{Update\_Tree}(\mathcal{T}, next, c, \mathsf{address}_{\mathsf{Usr}_1})$ to update $\mathcal{T}$, set $next = next + 1$, and output 1 if successful. We call such a deposit transaction valid. [a]

Withdrawal: withdraw is executed by user Usr from some address $\mathsf{address}_{\mathsf{Usr}_2}$ and mixer Srv, as follows:

(1) **User:** On input (k, r), to withdraw a deposit transaction $c = \mathsf{Com}(k||r)$ proceed as follows:

(a) Send Read request to $\mathcal{F}_{ban}^{\mathcal{Q}}$ and denote by S_{ban} the received banned-addresses list.

(b) Send Read request to $\mathcal{F}_{bb}$ and denote by $\mathsf{S}_{leaves} = \{(index_i, c_i, \mathsf{address}_i)\}_{i\in[k]}$ the subset of the returned records from $\mathcal{F}_{bb}$ where c_i is a valid deposit transaction, and $\mathsf{address}_i$ is the address associated with it. For each $i \in [k]$ such that $\mathsf{address}_i \in \mathsf{S}_{ban}$ set $c_i = 0$ and $\mathsf{address}_i = \perp$.

(c) Construct a Merkle tree $\mathcal{T}$ with $(c_1, \ldots, c_k)$ being the leaves and let $R_{\mathcal{T}}$ be the root of $\mathcal{T}$.

(d) Compute the authentication path $O(\mathcal{T}, \ell)$, where ℓ is the leaf index of c in the computed $\mathcal{T}$ (if no leaf c, abort).

(e) Compute $h = H(k)$ and $\pi \leftarrow \mathsf{ZK.Prove}(\sigma, (R_{\mathcal{T}}, h), (k, r, \ell, O(\mathcal{T}, \ell)))$.

(f) Send $(\mathsf{Write}, (h, \pi), \mathsf{address}_{\mathsf{Srv}})$ to $\mathcal{F}_{bb}$. We call (h, π) a *withdrawal transaction* and h its *nullifier.*

(2) **Mixer:** Upon receiving $(index, (h, \pi), \mathsf{address}_{\mathsf{Usr}_2})$ from $\mathcal{F}_{bb}$, perform the following steps:

(a) Send Read request to $\mathcal{F}_{ban}^{\mathcal{Q}}$ and denote by S_{ban} the received banned-addresses list.

(b) Zero leaves associated with deposits from banned addresses: For each leaf ℓ in $\mathcal{T}$ and an address $\mathsf{address}_\ell$ associated with it, check if $\mathsf{address}_\ell \in \mathsf{S}_{ban}$, and if so invoke the subroutine in fig. 5 on $\mathsf{Update\_Tree}(\mathcal{T}, \ell, 0, \perp)$ to update $\mathcal{T}$.

(c) Verify that h did not appear in any previous withdrawal transaction (output 0 otherwise).

(d) Output $b \leftarrow \mathsf{ZK.Ver}(\sigma, (R_{\mathcal{T}}, h), \pi)$ and if $b = 1$ set $\mathsf{S}_\emptyset = \mathsf{S}_\emptyset \cup h$.

[a] in the blockchain implementation the mixer also verifies that the deposit transaction is funded.

Figure 4: Compliant privacy mixer protocol Haze

Daze instructs the user to encrypt its nullifier as part of the deposit, and in case its address becomes non-compliant $\mathcal{Q}$ decrypts and publishes the nullifier. The idea in Daze is to guarantee compliance by blocking withdrawals that present a nullifier that is associated with a non-compliant address. If a user manages to withdraw their funds prior to becoming non-compliant, their nullifier appears on the nullifier set in the smart contract. Once becoming non-compliant, their decrypted nullifier is published alongside with their address, and can be linked to the withdrawal transaction, thus revoking the transaction privacy provided by the mixer to non-compliant users. Concretely:

- In deposit, Usr samples k, r and computes $\mathsf{Com}(k||r)$, and an encryption of the nullifier under the public-key $pk_{\mathcal{Q}}$ of the banned-addresses list maintainer $\mathcal{Q}$, together with a "consistency" proof. The proof is generated using a ZK-SNARK scheme Φ for a polynomial time decidable binary relation $R_{consist}$, where the statement is $\left(\mathsf{Com}(k||r), pk_{\mathcal{Q}}, \text{ciphertext}\right)$ and the witness is $(k, r, \text{ciphertext randomness})$. The deposit transaction submitted by Usr consists of: $\left((\mathsf{Com}(k||r), \text{ciphertext}), \text{proof}\right)$.
- The withdraw protocol of Daze differs from Tornado Cash only in its nullifier treatment. That is, the Srv compares the nullifier in the withdraw transaction not only to nullifiers of prior withdrawals but also to the nullifiers in the data field in $\mathcal{F}_{ban}^{\mathcal{Q}}$, and rejects withdrawal if appears in either.

Subroutine Update_Tree executed by Srv on $(\mathcal{T}, \ell, c, \text{address})$, and shared parameters as in fig. 4, where $\mathcal{T}$ is a Merkle tree on n leaves (and height $\log(n)$).

For a node $v \in \mathcal{T}$ we denote by $[v]$ its value, and similarly by $[v]_{\text{sb}}$ and $[v]_{\text{pr}}$ its sibling and parent values, respectively.

The subroutine proceeds as follows:

(1) Set the value of the ℓ'th leaf to c and denote this leaf by v.
(2) while $v \neq R_{\mathcal{T}}$ compute:
   (a) $[v]_{\text{pr}} = H([v]||[v]_{\text{sb}})$ for left child v and
   $[v]_{\text{pr}} = H([v]_{\text{sb}}||[v])$ for right child v.
   (b) $v = \text{parent}(v)$

Figure 5: The subroutine Update_Tree **updates the hashes along the path from the ℓ'th leaf to the root in $\mathcal{T}$.**

We emphasize that in Daze the mixer is not required to maintain a compliant Merkle tree, and in particular does not perform any tree updates for non-compliant leaves. This reduces the gas consumption of Daze compared to Haze. See fig. 6 for formal description of Daze.

The protocol Daze provides correctness, privacy, soundness and compliance in the same sense as in section 3.1, where privacy is guaranteed only for compliant users. Correctness and soundness stems from the same reasons as in Haze. *Privacy* of compliant users stems from the same reasons as in Haze, combined with the CPA-security of $\mathcal{E}$. Moreover, due to the knowledge soundness of Φ, the decryption functionality provided by the banned-address list maintainer Q cannot be leveraged to violate privacy of compliant user by mounting a malleability attack. *Compliance* follows from the construction of Daze. In particular, due to the binding of C, the collision resistance of $H(\cdot)$, and Daze's updates to include nullifiers that appear in the banned-addresses list it guaranteed that depositors from banned-addresses cannot produce a nullifier and zero knowledge proof that is accepted by Daze.

4 INTEGRATING WITH THE BLOCKCHAIN

In this section we address the necessary adjustments needed to make our protocols in section 3 securely deployed on a blockchain. In particular, they need to remain correct, compliant, private, and sound on the blockchain. Both protocol Haze and Daze use primitives that are available on many contemporary platforms, and in particular all EVM blockchains, and therefore are broadly applicable. In section 4.3 we introduce additions to Haze to enable balanced distribution of the cost overhead induced by the compliance maintenance mechanism. Additionally, in appendix C, we propose a solution that enables releasing funds deposited from non-compliant addresses back to some predetermined entity. We implement both Haze and Daze and empirically evaluate their performance, comparing them to state of the art in section 5.

4.1 Deployment on the Blockchain

In the blockchain deployment of Haze and Daze, the *mixer* Srv is an on-chain smart contract that implements the logic of fig. 4 and fig. 6, respectively, and can be publicly audited. The Merkle tree $\mathcal{T}$ is stored in the smart contract's storage. Users interact with Srv by sending transactions to the blockchain. The address of a transaction sender is publicly visible and therefore a deposit to Srv can be linked to the address from which it originated. However, a user may send many transactions to Srv from multiple addresses.

In Haze, a *deposit transaction* is an on-chain transaction transferring M coins to the smart contract Srv (we assume for ease of notations that $M = 1$, but any amount will work as long as it is the same amount across all users and all transactions). The transaction will also have as auxiliary data $c = \text{Com}(k||r)$ as described in deposit of both protocols. In Daze, the encrypted nullifier and the zero knowledge proof are also included, as described in fig. 6, deposit Step 1. The user must also include in the transaction additional funds to pay the gas fees.

In both Haze and Daze, a *withdrawal transaction* is an on-chain transaction from the user to Srv. In order to maintain privacy, the transaction should originate from a previously unused address. The transaction does not transfer any funds into Srv, and includes in the auxiliary data the user inputs as described in withdraw of both protocols. When the withdrawal is processed by Srv, M coins will be released to the address designated in the transaction. The transaction should again include the gas fee needed to execute the withdrawal. Paying for gas of a withdrawal is preferably done via a relay in order to preserve privacy, see section 4.2.

The banned-addresses list and its maintainer. The banned-addresses list functionality $\mathcal{F}_{ban}^{Q}$ is an on-chain smart contract which receives queries from users and asserts whether an address has been included in the banned address list (sanctions list). Today, the company Chainalysis maintains such a contract [7] on the Ethereum blockchain and reflects the sanctions designations listed on economic/trade embargo lists from governments and organizations including the US, EU, and the UN. In off-chain settings, there are defense mechanisms in place to prevent the corruption of a single entity from compromising the system, i.e. search warrants, that require authorities to reach some threshold of permissions. For this reason, the list maintainer should ideally consist of at least two off-chain authorized and separate entities (i.e. law enforcement agency and court) which participate together in signing the transaction which adds a user to the banned-addresses list.

Instantiating the de-anonymization process. In Daze, it is vital that users are not "quietly" de-anonymized, without a public record of the decrypted nullifiers being posted on-chain. The de-anonymization process begins once an address being published on the banned-addresses list. We instantiate the de-anonymization process as a committee, utilizing a CPA-secure threshold encryption scheme. We aim for a committee that contains a designated set of parties that are in charge of maintaining the mixer (i.e. a governance module [39]). Each member of the committee holds a key share and together the members decrypt nullifiers associated with deposits from non-compliant addresses by running threshold

Common parameters: A security parameter λ, a function H sampled uniformly at random from a collision-resistant hash function family $\mathcal{H}$, and a ZK-SNARK scheme $\Phi = (\mathsf{ZK.Setup}, \mathsf{ZK.Prove}, \mathsf{ZK.Ver}, \mathsf{ZK.Sim})$ for relation R as defined in section 3.1 and $R_{consist}$ as defined above. A CPA-secure PKE scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, and a a secure commitment scheme $C = (\mathsf{Com}, \mathsf{Dcom})$.
Parties and addresses: A mixer Srv with public address $\mathsf{address}_{\mathsf{Srv}}$ and a user Usr with some addresses $\mathsf{address}_{\mathsf{Usr}_1}$ and $\mathsf{address}_{\mathsf{Usr}_2}$.
Storage: Srv locally stores a full binary Merkle tree $\mathcal{T}$ on n leaves, all initialized to zero, the location of the next available leaf $next = 0$, and, initially empty, nullifier set $\mathsf{S}_\emptyset$. Each leaf in $\mathcal{T}$ is associated with an address, initially set to $\perp$.
Trusted setup: Bulletin-board functionality $\mathcal{F}_{bb}$ and a Banned-addresses Functionality $\mathcal{F}_{ban}^{Q}$, containing records $(\mathsf{address}_\ell, \mathsf{data}_\ell)$, where $\mathsf{data}_\ell = \mathsf{Dec}_{sk_Q}(e)$ for e appearing in a deposit from $\mathsf{address}_\ell$. A common reference string σ produced by running ZK.Setup. A key pair (sk_Q, pk_Q) produced by running Gen, where $\mathcal{F}_{ban}^{Q}$ is parameterized on pk_Q, and sk_Q is securely stored by the predefined entity with address Q.
Deposit: deposit is executed by user Usr from some address $\mathsf{address}_{\mathsf{Usr}_1}$ and mixer Srv, as follows:

(1) **User:** samples uniformly at random $k, r, r_e \leftarrow \{0,1\}^{t(\lambda)}$ for some polynomial $t(\cdot)$ and performs the following:
   (a) computes $c = \mathsf{Com}(k||r)$ and $e = \mathsf{Enc}_{pk_Q}(H(k), r_e)$.
   (b) $\pi_{\text{in}} \leftarrow \mathsf{ZK.Prove}(\sigma, (c, pk_Q, e), (k, r, r_e))$.
   (c) Sends $(\mathsf{Write}, ((c, e, \pi_{\text{in}}), \mathsf{address}_{\mathsf{Srv}})$ to $\mathcal{F}_{bb}$.
(2) **Mixer:** Upon receiving $(index, (c, e, \pi_{\text{in}}), \mathsf{address}_{\mathsf{Usr}_1})$ from $\mathcal{F}_{bb}$, perform the following steps:
   (a) Check that c is not already a leaf in $\mathcal{T}$ (else output 0).
   (b) Compute $b \leftarrow \mathsf{ZK.Ver}(\sigma, (c, pk_Q, e), \pi_{\text{in}})$ and if $b = 0$ output 0.
   (c) Invoke the subroutine in fig. 5 on $\mathsf{Update\_Tree}(\mathcal{T}, next, c, \mathsf{address}_{\mathsf{Usr}_1})$ to update $\mathcal{T}$, set $next = next + 1$, and output 1 if successful. We call such a deposit transaction valid.

Withdrawal: withdraw is executed by user Usr from some address $\mathsf{address}_{\mathsf{Usr}_2}$ and mixer Srv, as follows:

(1) **User:** On input (k, r), to withdraw a deposit transaction $c = \mathsf{Com}(k||r)$ proceed as follows:
   (a) Send Read request to $\mathcal{F}_{bb}$ and denote by $\mathsf{S}_{leaves} = \{(index_i, c_i, \mathsf{address}_i)\}_{i\in[k]}$ the subset of the returned records from $\mathcal{F}_{bb}$ where c_i is a valid deposit transaction, and $\mathsf{address}_i$ is the address associated with it.
   (b) Construct a Merkle tree $\mathcal{T}$ with $(c_1, \ldots, c_k)$ being the leaves and let $R_{\mathcal{T}}$ be the root of $\mathcal{T}$.
   (c) Compute the authentication path $O(\mathcal{T}, \ell)$, where ℓ is the leaf index of c in the computed $\mathcal{T}$ (if no leaf c, abort).
   (d) Compute $h = H(k)$ and $\pi_{\text{out}} \leftarrow \mathsf{ZK.Prove}(\sigma, (R_{\mathcal{T}}, h), (k, r, \ell, O(\mathcal{T}, \ell)))$.
   (e) Send $(\mathsf{Write}, (h, \pi_{\text{out}}), \mathsf{address}_{\mathsf{Srv}})$ to $\mathcal{F}_{bb}$. We call (h, π) a *withdrawal transaction* and h its *nullifier*.
(2) **Mixer:** Upon receiving $(index, (h, \pi_{\text{out}}), \mathsf{address}_{\mathsf{Usr}_2})$ from $\mathcal{F}_{bb}$, perform the following steps:
   (a) Send Read request to $\mathcal{F}_{ban}^{Q}$ and denote by $\mathsf{S}_{ban} = \{(\mathsf{address}_\ell, \mathsf{data}_\ell)\}_{\ell\in[m]}$ the received banned-addresses list.
   (b) Add nullifiers associated with deposits from banned addresses: For each leaf ℓ in $\mathcal{T}$ and an address $\mathsf{address}_\ell$ associated with it, check if $\mathsf{address}_\ell \in \mathsf{S}_{ban}$, and if so $\mathsf{S}_\emptyset = \mathsf{S}_\emptyset \cup \mathsf{data}_\ell$.
   (c) Verify that $h \notin \mathsf{S}_\emptyset$ (output 0 otherwise).
   (d) Output $b \leftarrow \mathsf{ZK.Ver}(\sigma, (R_{\mathcal{T}}, h), \pi_{\text{out}})$ and if $b = 1$ set $\mathsf{S}_\emptyset = \mathsf{S}_\emptyset \cup h$.

Figure 6: Compliant privacy mixer protocol Daze**, with de-anonymization of non-compliant users.**

decryption. This guarantees the privacy of compliant users as long as at least one of the committee members is honest, i.e., it refuses to decrypt a nullifier of a non-banned address. The committee is incentivised to perform de-anonymization, as if they do not, the entire mixer risks being added to the banned-addresses list and censored by the network. We emphasize that even if all committee members are corrupt, they still cannot break the soundness of the mixer, i.e. cannot steal funds. We note that due to the two-stage de-anonymization process, there is a wait time between the publication of an address on the banned-addresses list and the time the cleartext nullifier is published. Therefore, if blockage of these funds during the wait time is desired, one can combine Daze with Haze and block the funds as soon as the address is included in the banned-addresses list.

4.2 Security Concerns over the Blockchain

Care needs to be taken in order to deploy our protocols on top of the blockchain, due to the asynchronous nature of the blockchain and the fact that messages are not written to the blockchain directly by users. The blockchain is indeed an append-only linked list, as per our theoretical model. However, messages may arrive at the blockchain simultaneously and several messages may be included in a single block (a single state update). Messages are sent to the blockchain either via broadcast, through a trusted relay or directly to validators through private channels. Block builders\validators choose how to order transactions inside a block according to their own best interest (typically according to a tip-maximizing order) [9]. This means that in addition to the security properties mentioned above, over blockchains, several other security concerns arise. For example, the deployed protocol needs to guarantee resilience to

hijacking and front-running.[2] We therefore have to take extra precautions when implementing Haze and Daze as a smart contract to mitigate these concerns. We emphasize that the honesty of the mixer is a assumed due to the fact that the code of the mixer is publicly auditable on the blockchain.

Hijacking resilience. In the idealized bulletin-board model, messages are written to $\mathcal{F}_{bb}$ directly without the possibility of interception. On the blockchain, messages are sent either via broadcast to the entire network, or through private channels to builders\validators. This introduces a previously undiscussed risk of messages being hijacked and modified in order to steal funds headed out of Srv. To prevent an attacker from replacing the recipient with its address after seeing withdraw transaction, a "non-malleability" property is required. This is achieved by using the SnarkJS and Circom implementation of Groth16 [16, 20, 22] which encodes non-malleability into the implemented circuit. This way, Haze and Daze prevent hijacking by including the recipient address in the zero-knowledge proof in the withdrawal transaction.

Front-running resilience. The front-running problem arises when Alice issues a withdrawing transaction w.r.t some state, and before Alice's withdrawal is included in a block, a new deposit or withdrawal by Bob is made to the mixer and is included in a block, thus changing the state of the mixer and potentially deeming Alice's withdrawal invalid [25] (meaning Bob's transaction front-ran Alice's). Due to this concern, we extend Haze and Daze to be front-running resilient. Both protocols enable a withdrawal to reference any previous root, as long as there were no updates to the banned-addresses list since that root was valid. In Haze, to accommodate this the previous root is simulated ("zeroing" the relevant deposits) which can be done in $O(\log(n))$ time, where n is the number of leaves in $\mathcal{T}$. This means the withdrawal transaction will be processed correctly.

Withdrawal gas fees. In order to pay the gas fee of a withdrawal transaction, the initiator of the transaction needs to have sufficient funds. The recipient address of a withdrawal needs to be a fresh address to maintain the privacy of our protocols. If the withdrawal were initiated by a fresh address, that address would need to somehow have these funds. However, the withdrawal address needs to be unlinkable to the address of the depositor. So the depositor can't simply fund this fresh address to pay the gas fees. This raises the question of how the gas fees can still be paid. For this reason, users should utilize relays to minimize the privacy loss. Relays exist in the original Tornado Cash implementation [31]. A relay receives a withdrawal transaction from the depositor via a private secure channel. In the recipient field of the withdrawal, the depositor will list a fresh address. The relay funds the gas costs for the withdrawal and forwards the withdrawal transaction to Srv. The relay cannot alter the recipient field since Haze and Daze are hijacking resilient. Srv processes the withdrawal and releases the funds to the fresh recipient address, minus a fee paid to the relay and the gas cost for the withdrawal transaction which are sent to the relay. This way, the withdrawal request cannot be linked to the depositor's address on-chain. However, this requires trust between the depositor and the relay, as the relay can link the depositor to the withdrawal request. See fig. 7.

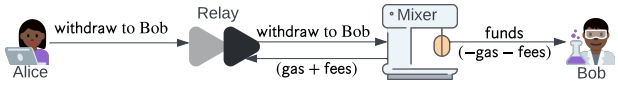

Figure 7: In order to pay the gas costs of the withdrawal transaction withdraw **addressed to Bob, the depositor Alice should utilize a relay. Alice and the relay communicate through a private channel. The relay forwards the withdrawal transaction to Srv. Srv processes the withdrawal and sends the released funds to Bob, minus the gas cost and fee that is sent back to the relay. This way Alice remain unlinkable to Bob on the blockchain.**

4.3 Haze: Economic Concerns over the Blockchain

In addition to the concerns addressed above, when implemented over the blockchain, Haze needs to not fail due to gas limitations. Concretely, the number of tree updates required per withdrawal depend on the newly banned addresses, thus making the cost of withdrawal non uniform per withdrawal and a priori unpredictable. In this section we provide a twofold treatment: once at the feasibility level, enforcing that such updates do not fail due to technical limitation of the host blockchain, and at the user level: a single user should not be made to cover the cost of a large update due to "bad timing". We implement and evaluate the cost of deposit and withdraw in section 5.

Balancing the gas costs associated with withdraw**.** As mentioned in section 3, enforcing compliance introduces a cost proportional to the number of newly non-compliant deposits times the cost of a single path update in $\mathcal{T}$. This implies that the cost of withdrawal depends on the number of newly non-compliant deposits and hence is a priori unpredictable and also non uniform on all withdrawals. In order to resolve this issue and create uniformity in the cost of withdrawals, we suggest creating a fund in the smart contract of Srv which will be funded by a depositor fee for every deposit to refund these users that happen to bear the cost of updates induced by changes in the banned-addresses list. As an update costs the same amount of gas as a deposit, the maximal amount of fee Haze needs to charge in order to cover these costs (per deposit) is bounded by the gas fee per a deposit transaction. This makes the price of a deposit increase by at most a factor of 2. In appendix C we show how this overhead can be refunded to compliant users, making the overall cost of using Haze comparable to Tornado Cash. This treatment also mitigates the risk of spanning the banned-addresses list by bad actors, as they would cover the price of the spamming upon the deposit to the mixer.

Overcoming block gas limit for large updates to $\mathcal{T}$. Recall that if a withdrawal transaction invokes a number of path updates in $\mathcal{T}$ that require gas that exceeds the gas limit for a single block, it

[2]We note that the protocol is replay resilience due to its soundness guarantee.

could potentially fail. To mitigate this issue, in the blockchain implementation of Haze, the withdraw function is split into two functions - withdraw and update. withdraw handles the withdrawal logic, while update handles the logic for updating $\mathcal{T}$ according to the banned-addresses list. This way, if the backlog of updates is too big for a single withdrawal transaction to be processed, any user in the system can call the update function to relieve the backlog. The gas for this altruistic transaction can be funded by the fees collected upon deposit, and will be refunded to the caller of the update function. This solution is compatible with the incentives of users of Haze who want to be able to withdraw their funds.

5 EMPIRICAL EVALUATION

In this section we implement and empirically evaluate the performance of our protocols Haze and Daze and compare them to Tornado Cash [31]. In our experiments, we used an Apple M1 Pro chip machine with 8-core CPU and 16GB RAM to run client Usr. We deploy and manage the mixer Srv on an Ethereum local blockchain using Ganache [8].

5.1 Implementation

Our implementation comprises of two components:

- Server Srv: the mixer, implemented as a smart contract in Solidity [11].
- Client Usr: the user, implemented in JavaScript.

Both are implemented using a fork of the Tornado Cash mixer [37] and client [38] extended to our protocols. As in Tornado Cash, for practical reasons we use Pedersen hash function [19] for the leaves instead of a commitment scheme and the MiMC hash function [1], which are implemented in the circomlib library [21]. The SNARK keypair and the Solidity verifier code are generated using SnarkJS [22], It uses a non-malleable implementation of the Groth16 [16] Protocol, PLONK [13] and FFLONK [12].

Haze*'s implementation.* We implemented Haze by forking the Tornado Cash code and introducing the following changes:

- The smart contract of the mixer now calls an external smart contract that manages the banned-addresses list.
- We store the Merkle tree $\mathcal{T}$ in the smart contract as a map with the node indices as keys. The map grows gradually as deposits enter our system. In addition we maintain a map of depositors' addresses to the indices of the tree leaves representing their deposits. This map is used to efficiently locate leaves in the tree associated with non-compliant deposits.
- We maintain a queue of the indices of the Merkle tree leaves associated with all non-compliant addresses that are currently known and not yet zeroed in the Merkle tree, based on banned-addresses list obtained from the external contract.
- Path updates of the Merkle tree are realized as follows: We implement an update function, that on input n zeroes the leaves associated with the first n indices in the queue and updates the values along their path to the root accordingly. We allow $n \leq 35$ as this is the maximum possible number of leaves that can be handled in a single transaction within the block gas limit of 30M. We change the implementation of the withdraw function to call the update function as part of its internal logic.

Daze's implementation. We implemented Daze by augmenting the Tornado Cash code with the de-anonymization functionality using ElGamal encryption [14] via the Circom library. Again, the smart contract of Haze calls an external contract that manages the banned-addresses list.

5.2 Experiments and Results

We evaluate the performance of Haze and Daze in terms of gas consumption and running time of the client, for deposits and withdrawals in different scenarios. We compare our measurements to Tornado Cash.

In all our experiments, we measure the actual gas consumption using the transaction receipt "gasUsed" field of corresponding requests. We set the block gas limit to $30 \cdot 10^6$ which is the gas limit used by the Ethereum mainnet today. We repeat each experiment 20 times and present the average result of all repetitions.

The purpose of our first experiment is to show that when there are no non-compliant addresses the gas consumption of the smart contract and running time of the client of Haze and Daze is comparable to Tornado Cash, which is widely used.

Experiment A: comparison of our protocols to Tornado Cash. We measure the cost of deployment of Haze and Daze to the blockchain, as well as the cost of a single deposit and withdrawal when there are no non-compliant addresses (we denote it the baseline setting). We run the deposit and withdrawal measurements on the client side as well. We compare the baseline gas costs and running time of our protocol with the deposit and withdrawal of Tornado Cash.

The results appear in Table 1. We see that for Haze, on average the gas cost has increased by 1.1% for our mixer compared to Tornado Cash for a withdrawal transaction, by 4.1% for a deposit transaction, and by 7.2% for the deployment. These increases are explained by the additional storage in Haze compared to Tornado Cash. The running time of the deposit on the client side of our protocol is identical to Tornado Cash and the withdrawal running time increases by 4.6% on average. This difference stems from the fact that in our implementation, in each withdrawal, the client checks for updates of the banned-addresses list. In Daze, we find that the overhead in deposit imposed by the new functionality is $410K$ gas, which implies approximately 40% increase in gas cost compared to Tornado Cash. The increase in gas cost stems from the zero-knowledge proof verification, that has a constant cost of approximately $300K$. We measured the running time of the deposit in Daze, and found it to be 1.4 seconds. The increase in the deposit running time stems from encryption and preparation of the zero-knowledge proof. We also ran experiments to measure the cost of withdrawal, both in terms of gas and running time on the client side.

	Server (gas)			Client (sec)		
Action	Tornado Cash	Haze	Daze	Tornado Cash	Haze	Daze
Deploy	1960209	2099721	2332409	—	—	—
Deposit	957037	996139	1367056	0.043	0.045	1.415
Withdraw	312254	315782	314529	3.040	3.035	3.123

Table 1: Comparison Haze **in the baseline setting vs. the Tornado Cash vs. our de-anonymization protocol** Daze**, for smart contract deployment (gas units), deposit (seconds), and withdrawal (gas units). Measurements taken when the mixer is populated with 1K deposits.**

Experiment B: Gas cost vs. number of non-compliant addresses. This experiment is relevant only to Haze, as the cost of a withdrawal depends on the number of newly non-compliant addresses. After populating the mixer with 1K deposits we measure the gas consumption of a deposit transaction and a withdrawal transaction as the number of deposits associated with newly non-compliant addresses increases.

We find that the gas consumption of withdrawal increases linearly with the number of deposits associated with non-compliant addresses. For completeness, we also present the gas consumption of a deposit operation, which does not change with the number of non-compliant addresses. Moreover, we find that the maximal number of tree updates due to non-compliant address that can be supported in one transaction is at most 35, since after that the gas required surpasses the block gas limit of 30M gas. The results are summarized in fig. 8. We note that the results of this experiment are independent of the number of deposits that were made to the mixer prior to the measurement. We verified this independence by repeating the experiment when populating the mixer with different numbers of deposits in $\{1, \ldots, 1000\}$, checking at increments of 100, and obtained the same results. The gas consumption of a withdrawal transaction consists of the zero-knowledge proof verification and tree updates, which result from each newly non-compliant address. The number of deposits in the mixer does not influence either of these components, as the updates depend only on tree height which is fixed throughout the lifetime of the mixer. Similarly, a deposit transaction depends on the tree height and is oblivious of the number of deposits inside the mixer, as well the number of non-compliant addresses. See fig. 8.

Experiment C: Running time vs. number of deposits in the mixer in Haze *and* Daze. We measure the client's withdrawal running time as the number of deposits in the mixer increases in the baseline setting (i.e., there are no banned addresses). Results for Haze are summarized in fig. 9. We see that the running time of the client for the withdrawal increases linearly with the number of deposits in the mixer. The increase in running time stems from increasing number of nodes in the constructed tree by the client.

Next, we assert that the number of newly non-compliant addresses does not influence the running of the client both for deposit and withdrawal transactions. More formally,

Experiment D: Running time vs. number of non-compliant addresses in Haze *and* Daze. After populating the mixer with 1K deposits we measure the running time of the client for a deposit

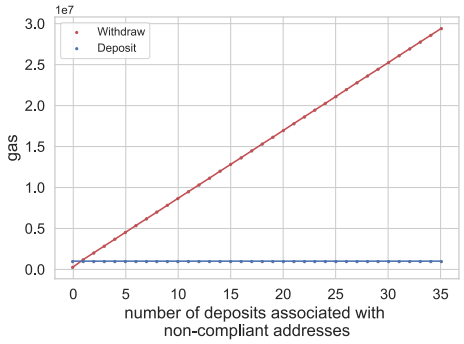

Figure 8: The deposit cost (blue) and withdrawal cost (red) of Haze **in gas units vs. the number of newly banned addresses as part of the upcoming withdrawal transaction. Measurements taken when the mixer is populated with 1K deposits.**

transaction and a withdrawal transaction as the number of deposits associated with newly non-compliant addresses increases.

We find that the running time of withdrawals in Haze and Daze is unaffected by the number of non-compliant addresses, per fixed number of deposit populating the mixer. We verified this by repeating the measurements for different deposit populations in the mixer, for the same values presented in the previous experiments. This is due to the fact that the client, similar to Tornado Cash, rebuilds the entire tree in each withdrawal request. Moreover, the deposit running time does not change with the amount of non-compliant addresses and the number of deposits in the mixer, similarly as for gas cost of deposits.

6 CONCLUSIONS

In this work we presented two compliant privacy mixers, Haze and Daze that attain correctness, soundness, privacy, and compliance. Daze additionally supports de-anonymization of non-compliant users. Our protocols can be deployed and used over the blockchain guaranteeing resiliency against: transaction hijacking, front-running, and banned-addresses list spamming. In addition, we propose a solutions for responsible release of banned funds due to non-compliance in both protocols. We implemented both Haze and Daze using Solidity for the mixer Srv and JavaScript for the client Usr. We ran extensive experiments demonstrating efficient user running time

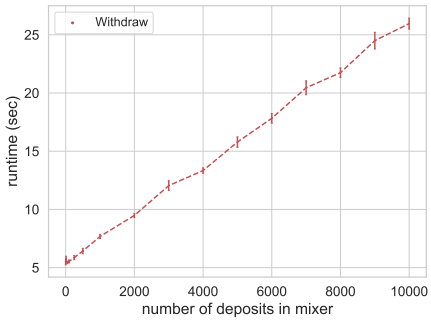

Figure 9: The time in seconds it takes Haze's client to prepare a withdraw transaction. This measurement is taken in the baseline setting when there are no banned addresses.

and realistic gas requirements comparable to the popular Tornado Cash mixer.

REFERENCES

[1] Martin Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 191–219. Springer, 2016.

[2] Kurt M. Alonso and Jordi Herrera-Joancomartí. Monero - privacy in the blockchain. *IACR Cryptol. ePrint Arch.*, 2018:535, 2017.

[3] Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *ACM SIGACT News*, 15(1):23–27, 1983.

[4] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A Kroll, and Edward W Felten. Mixcoin: Anonymity for bitcoin with accountable mixes. In *Financial Cryptography and Data Security: 18th International Conference, FC 2014, Christ Church, Barbados, March 3-7, 2014, Revised Selected Papers 18*, pages 486–504. Springer, 2014.

[5] Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In *Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers*, pages 423–443. Springer, 2020.

[6] Joseph Burleson, Michele Korver, and Dan Boneh. Privacy-protecting regulatory solutions using zero-knowledge proofs. https://api.a16zcrypto.com/wp-content/uploads/2022/11/ZKPs-and-Regulatory-Compliant-Privacy.pdf, 2022.

[7] Chainalysis. Chainalysis oracle for sanctions screening. https://go.chainalysis.com/chainalysis-oracle-docs.html, 2023.

[8] ConsenSys Software Inc. Ganache. https://www.trufflesuite.com/docs/ganache/overview, 2021.

[9] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (S&P)*, pages 910–927. IEEE, 2020.

[10] Maya Dotan, Saar Tochner, Aviv Zohar, and Yossi Gilad. Twilight: A differentially private payment channel network. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 555–570, 2022.

[11] Ethereum Foundation. Solidity programming language. https://docs.soliditylang.org/en/latest/, 2019.

[12] Ariel Gabizon and Zachary J Williamson. fflonk: a fast-fourier inspired verifier efficient version of plonk. *Cryptology ePrint Archive*, 2021.

[13] Ariel Gabizon, Zachary J Williamson, and Oana Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *Cryptology ePrint Archive*, 2019.

[14] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Inf. Theory*, 31(4):469–472, 1985.

[15] Shafi Goldwasser and Sunoo Park. Public accountability vs. secret laws: Can they coexist? *Cryptology ePrint Archive*, 2018.

[16] Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 305–326. Springer, 2016.

[17] Ethan Heilman, Leen Alshenibr, Foteini Baldimtsi, Alessandra Scafuro, and Sharon Goldberg. Tumblebit: An untrusted bitcoin-compatible anonymous payment hub. In *Network and distributed system security symposium*, 2017.

[18] Daira Hopwood, Sean Bowe, Taylor Hornby, Nathan Wilcox, et al. Zcash protocol specification. *GitHub: San Francisco, CA, USA*, 4(220):32, 2016.

[19] Iden3. edersen hash. https://iden3-docs.readthedocs.io/en/latest/iden3_repos/research/publications/zkproof-standards-workshop-2/pedersen-hash/pedersen.html, 2019.

[20] iden3. Circom. https://github.com/iden3/circom, 2022.

[21] iden3. Circomlib/circuits. https://github.com/iden3/circomlib/tree/master/circuits, 2022.

[22] iden3. Javascript and pure web assembly implementation of zksnark and plonk schemes. https://github.com/iden3/snarkjs, 2022.

[23] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2020.

[24] Labrys. Mev watch. https://web.archive.org/web/20230428094150/https://www.mevwatch.info/, 2023.

[25] Duc V Le and Arthur Gervais. Amr: Autonomous coin mixer with privacy preserving reward distribution. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 142–155, 2021.

[26] Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. Silentwhispers: Enforcing security and privacy in decentralized credit networks. In *NDSS*, 2017.

[27] Sarah Meiklejohn and Rebekah Mercer. Möbius: Trustless tumbling for transaction privacy. *Proceedings on Privacy Enhancing Technologies*, 2018(2):105–121, 2018.

[28] Ian Miers, Christina Garman, Matthew Green, and Aviel D Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *2013 IEEE symposium on security and privacy*, pages 397–411. IEEE, 2013.

[29] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, page 21260, 2008.

[30] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Annual international cryptology conference*, pages 129–140. Springer, 1991.

[31] Alexey Pertsev, Roman Semenov, and Roman Storm. Tornado cash privacy solution version 1.4. https://berkeley-defi.github.io/assets/material/Tornado%20Cash%20Whitepaper.pdf, 2019.

[32] Antoine Rondelet and Michal Zajac. Zeth: On integrating zerocash on ethereum. *arXiv preprint arXiv:1904.00905*, 2019.

[33] Stefanie Roos, Pedro Moreno-Sanchez, Aniket Kate, and Ian Goldberg. Settling payments fast and private: Efficient decentralized routing for path-based transactions. *arXiv preprint arXiv:1709.05748*, 2017.

[34] Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*, pages 459–474. IEEE, 2014.

[35] Ameen Soleimani. Privacy pools with opt-in or opt-out anonymity sets. https://github.com/ameensol/privacy-pools, 2023.

[36] Alin Tomescu, Adithya Bhat, Benny Applebaum, Ittai Abraham, Guy Gueta, Benny Pinkas, and Avishay Yanai. Utt: Decentralized ecash with accountable privacy. Cryptology ePrint Archive, Paper 2022/452, 2022. https://eprint.iacr.org/2022/452.

[37] Tornado Cash. Tornado cash privacy solution. https://github.com/tornadocash/tornado-core, 2019.

[38] Tornado Cash. Tornado-cli. https://github.com/tornadocash/tornado-cli, 2019.

[39] Tornado Cash. Solidity programming language. https://github.com/tornadocash/tornado-governance, 2021.

[40] TRM Labs. North korea's lazarus group moves funds through tornado cash. https://www.trmlabs.com/post/north-koreas-lazarus-group-moves-funds-through-tornado-cash, 2022.

[41] U.S. Department Of Treasury. Specially designated nationals and blocked persons list (sdn) human readable lists. https://home.treasury.gov/policy-issues/financial-sanctions/specially-designated-nationals-and-blocked-persons-list-sdn-human-readable-lists, 2022.

[42] U.S. Department Of Treasury. U.s. treasury sanctions notorious virtual currency mixer tornado cash. https://home.treasury.gov/news/press-releases/jy0916, 2022.

[43] Luke Valenta and Brendan Rowan. Blindcoin: Blinded, accountable mixes for bitcoin. In *Financial Cryptography and Data Security: FC 2015 International Workshops, BITCOIN, WAHC, and Wearable, San Juan, Puerto Rico, January 30, 2015, Revised Selected Papers*, pages 112–126. Springer, 2015.

[44] Nicolas Van Saberhagen. Cryptonote v 2.0. 2013.

[45] Zhipeng Wang, Stefanos Chaliasos, Kaihua Qin, Liyi Zhou, Lifeng Gao, Pascal Berrang, Benjamin Livshits, and Arthur Gervais. On how zero-knowledge proof blockchain mixers improve, and worsen user privacy. Cryptology ePrint Archive, Paper 2023/341, 2023. https://eprint.iacr.org/2023/341.

[46] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

[47] Mike Wu, Will McTighe, Kaili Wang, Istvan A Seres, Nick Bax, Manuel Puebla, Mariano Mendez, Federico Carrone, Tomás De Mattey, Herman O Demaestri, et al. Tutela: An open-source tool for assessing user-privacy on ethereum and tornado cash. *arXiv preprint arXiv:2201.06811*, 2022.

A PRELIMINARIES (FORMAL DEFINITIONS FROM SECTION 2)

A.1 CPA-Secure Public Key Encryption

A public key encryption scheme has the following syntax and correctness requirement.

Definition A.1 (Public-Key Encryption (PKE)). A *public-key encryption (PKE) scheme* with message space M is a triple $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ of ppt algorithms satisfying the following conditions:

- Gen (key generation) takes as input the security parameter 1^λ, and outputs a pair (pk, sk) consisting of a public key pk and a secret key sk; denoted: $(pk, sk) \leftarrow \mathsf{Gen}(1^\lambda)$.
- Enc (encryption) takes as input a public key pk and a message $m \in \mathsf{M}$, and outputs a ciphertext e; denoted: $e \leftarrow \mathsf{Enc}_{pk}(m)$.
- Dec (decryption) takes as input a secret key sk and a ciphertext e, and outputs a decrypted message m'; denoted: $m' \leftarrow \mathsf{Dec}_{sk}(e)$.

Correctness. The scheme is *correct* if for every $(pk, sk) \leftarrow \mathsf{Gen}(1^\lambda)$ and every message $m \in \mathsf{M}$,

$$\Pr[\mathsf{Dec}_{sk}(\mathsf{Enc}_{pk}(m)) = m] \geq 1 - \mathrm{negl}(\lambda)$$

where the probability is taken over the random coins of the encryption algorithm.

Security against chosen plaintext attack. A PKE $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is CPA-secure if no ppt adversary $\mathcal{A}$ can distinguish between the encryption of two equal length messages x_0, x_1 of his choice. This is formally stated using the following experiment between a challenger Chal and the adversary $\mathcal{A}$.

The CPA indistinguishability experiment $\mathsf{EXP}^{cpa}_{\mathcal{A},\mathcal{E}}(\lambda)$**:**

(1) $\mathsf{Gen}(1^\lambda)$ is run by Chal to obtain keys (pk, sk).
(2) Chal provides the adversary $\mathcal{A}$ with pk $\mathcal{A}$ sends to Chal two messages $x_0, x_1 \in \mathsf{M}$ s.t. $|x_0| = |x_1|$.
(3) Chal chooses a random bit $b \in \{0, 1\}$, computes a ciphertext $e \leftarrow \mathsf{Enc}_{pk}(x_b)$ and sends e to $\mathcal{A}$. We call e the challenge ciphertext.
(4) $\mathcal{A}$ outputs a bit b'.
(5) The output of the experiment is defined to be 1 if $b' = b$ (0 otherwise).

*Definition A.2 (*CPA*-security).* A public key encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has indistinguishable encryptions under chosen-plaintext attacks (or is CPA-*secure*) if for all ppt adversaries $\mathcal{A}$ there exists a negligible function negl such that:

$$\Pr[\mathsf{EXP}^{cpa}_{\mathcal{A},\mathcal{E}}(\lambda) = 1] \leq \frac{1}{2} + \mathrm{negl}(\lambda)$$

where the probability is taken over the random coins of $\mathcal{A}$ and Chal.

A.2 Commitment Schemes

Definition A.3 (Commitment Scheme). Let $C = (\mathrm{Com}, \mathrm{Decom})$ be an non-interactive protocol between a committer C and a receiver R. We say that C is a *secure commitment scheme* if the following properties hold:

Correctness: If C and R do not deviate from the protocol, then R should accept during the decommit phase with probability 1.

Binding: For every ppt C^*, there exists a negligible function $\mathrm{negl}(\cdot)$ such that C^* succeeds in the following game with probability at most $\mathrm{negl}(\lambda)$: On security parameter 1^λ: C^* first produces a commitment c. Then C^* outputs two decommitments (c, m_0, d_0) and (c, m_1, d_1), and succeeds if $m_0 \neq m_1$ and R accepts both decommitments.

Hiding: For every ppt receiver R^* and every two messages m_0, m_1, the view of R^* after receiving a commitment to m_0 is indistinguishable from its view after receiving a commitment to m_1.

It is known how to construct a non-interactive, perfectly binding commitment scheme from any one-way permutation [3]. Pedersen [30] constructed a computationally binding and unconditionally hiding scheme based on the discrete logarithm problem.

A.3 Zero Knowledge

Zero-knowledge succinct non-interactive Argument of Knowledge [Groth [16]]. Let R be a polynomial time decidable binary relation over pairs (ϕ, w) where ϕ is the statement and w the witness.

An efficient-prover publicly verifiable non-interactive argument $\Phi = (\mathsf{ZK.Setup}, \mathsf{ZK.Prove}, \mathsf{ZK.Ver}, \mathsf{ZK.Sim})$ for R is a quadruple of ppt algorithms as follows:

- $(\sigma, \tau) \leftarrow \mathsf{ZK.Setup}(R)$: The setup produces a common reference string σ and a simulation trapdoor τ for R.
- $\pi \leftarrow \mathsf{ZK.Prove}(R, \sigma, \phi, w)$: The prover algorithm takes as input a common reference string σ and $(\phi, w) \in R$ and returns an argument π.
- $0/1 \leftarrow \mathsf{ZK.Ver}(R, \sigma, \phi, \pi)$: The verification algorithm takes as input a common reference string σ, a statement ϕ and an argument π and returns 0 (reject) or 1 (accept).
- $\pi \leftarrow \mathsf{ZK.Sim}(R, \tau, \phi)$: The simulator takes as input a simulation trapdoor and statement ϕ and returns a simulated argument π.

Definition A.4 (Succinct non-interactive zero-knowledge argument of knowledge). We say $\Phi = (\mathsf{ZK.Setup}, \mathsf{ZK.Prove}, \mathsf{ZK.Ver}, \mathsf{ZK.Sim})$ is a *perfect succinct non-interactive zero-knowledge argument of knowledge (ZK-SNARK)* for R if it has:

- perfect completeness: Given any true statement, an honest prover should be able to convince an honest verifier to accept it. Formally, for all $(\phi, w) \in R$

$$\Pr\left[\mathsf{ZK.Ver}(R, \sigma, \phi, \pi) = 1 \middle|^{(\sigma,\tau) \leftarrow \mathsf{ZK.Setup}(R);}_{\pi \leftarrow \mathsf{ZK.Prove}(R,\sigma,\phi,w)}\right] = 1$$

- perfect zero-knowledge: An argument that does not leak any information besides the truth of the statement. Formally, for all $(\phi, w) \in R$ and all adversaries $\mathcal{A}$

$$\Pr[\mathcal{A}(R,\sigma,\tau,\pi)=1|^{(\sigma,\tau)\leftarrow \mathsf{ZK.Setup}(R);}_{\pi\leftarrow \mathsf{ZK.Prove}(R,\sigma,\phi,w)}]$$

$$=\Pr[\mathcal{A}(R,\sigma,\tau,\pi)=1|^{(\sigma,\tau)\leftarrow \mathsf{ZK.Setup}(R);}_{\pi\leftarrow \mathsf{ZK.Sim}(R,\tau,\phi)}]$$

- computational knowledge soundness: There exists an extractor that extracts a witness whenever the adversary produces a valid argument (given access to its internal state). Formally, for all non-uniform polynomial time adversaries $\mathcal{A}$ there exists a non-uniform polynomial time extractor $\chi_{\mathcal{A}}$, and a negligible function $\mathrm{negl}(\cdot)$ such that,

$$\Pr[{}^{(\phi,w)\notin R \text{ and}}_{\mathsf{ZK.Ver}(R,\sigma,\phi,\pi)=1}|^{(\sigma,\tau)\leftarrow \mathsf{ZK.Setup}(R);}_{((\phi,\pi);w)\leftarrow(\mathcal{A}||\chi_{\mathcal{A}})(R,\sigma)}] < \mathrm{negl}(\lambda)$$

- The proof π is of polynomial size in λ and ZK.Ver is polynomial in $\lambda + |\phi|$.

B FORMAL DEFINITIONS

In this section we formalize the properties *correctness, privacy, soundness* and *compliance*.

First, we formally define the following terms: *Correctness* is in the sense that any deposited funds can be withdrawn (once) as long as the matching deposit transaction is compliant at the time of the withdrawal, i.e., the withdrawn funds were not deposited from an address that is banned in $\mathcal{F}^{\mathcal{Q}}_{ban}$. *Soundness* is in the sense that no user can withdraw more than it deposited. *Privacy* is in the sense that a withdrawal cannot be linked to any non withdrawn deposit. *Compliance* is in the sense that funds belonging to deposit transactions associated with an address in $\mathcal{F}^{\mathcal{Q}}_{ban}$ cannot be withdrawn.

To formally state these properties we first set some notations. We denote the view of user Usr in an execution of the deposit and withdraw protocols in Figure 4 and Figure 6, by

$$(r, \mathsf{dtxn}, \mathsf{address}_A) \leftarrow \mathsf{view}^{\mathsf{deposit}}_{\mathsf{Usr}}(\lambda) \text{ and}$$
$$(\mathsf{wtxn}, \mathsf{address}_B) \leftarrow \mathsf{view}^{\mathsf{withdraw}}_{\mathsf{Usr}}(r, \lambda)$$

respectively, where the view consists of the party's randomness, the generated (deposit/withdrawal) transaction, and the address associated with the transaction. We note that $\mathsf{view}^{\mathsf{withdraw}}_{\mathsf{Usr}}(r,\lambda)$ is defined w.r.t the first execution of withdraw on input r. We denote the output of Srv by

$$\mathsf{out}^{\mathsf{deposit}}_{\mathsf{Srv}}(\mathsf{dtxn}, \lambda) = b \text{ and}$$
$$\mathsf{out}^{\mathsf{withdraw}}_{\mathsf{Srv}}(\mathsf{wtxn}, \lambda) = b$$

where the bit b is the output of Srv. We call a transaction valid if $b = 1$. We remark that $\mathcal{F}^{\mathcal{Q}}_{ban}$ and $\mathcal{F}_{bb}$ are accessible to any party in the system, and the transactions as well as the output of Srv may depend on their content.

Definition B.1 (Correctness). A mixer protocol $\Pi = (\mathsf{deposit}, \mathsf{withdraw})$ is *correct* if for every $\lambda \in \mathbb{N}$ and

$$(r, \mathsf{dtxn}, \mathsf{address}_A) \leftarrow \mathsf{view}^{\mathsf{deposit}}_{\mathsf{Usr}}(\lambda) \text{ and}$$
$$(\mathsf{wtxn}, \mathsf{address}_B) \leftarrow \mathsf{view}^{\mathsf{withdraw}}_{\mathsf{Usr}}(r, \lambda)$$

the following holds with probability $\geq 1 - \mathrm{negl}(\lambda)$:

- $\mathsf{out}^{\mathsf{deposit}}_{\mathsf{Srv}}(\mathsf{dtxn}, \lambda) = 1$
- if $\mathsf{address}_A$ is not recorded in $\mathcal{F}^{\mathcal{Q}}_{ban}$ prior to wtxn generation then $\mathsf{out}^{\mathsf{withdraw}}_{\mathsf{Srv}}(\mathsf{wtxn}, \lambda) = 1$

where the probability is over the randomness of Usr and Srv.

Intuitively, the soundness property needs to capture that any user cannot withdraw more funds than it deposited. We formalize this by requiring that for any user in any point of time, represented by *index* in $\mathcal{F}_{bb}$, the number of valid withdrawals made from addresses belonging to a user must not exceed the number of successful deposits made by the user.

Definition B.2 (Soundness). A mixer protocol $\Pi = (\mathsf{deposit}, \mathsf{withdraw})$ is *sound* if for every $index \in \mathbb{N}$, and any ppt user Usr, associated with address set $\mathsf{S}_{\mathsf{Usr}}$, the following holds with probability $\geq 1 - \mathrm{negl}(\lambda)$ over the randomness of Usr and Srv:

$$\left|\{(i, \mathsf{dtxn}_i, (\mathsf{address}_i, \mathsf{address}_{\mathsf{Srv}}))\}_{i<index \text{ and } \mathsf{address}_i \in \mathsf{S}_{\mathsf{Usr}}}\right|$$
$$\geq \left|\{(i, \mathsf{wtxn}_i, (\mathsf{address}_i, \mathsf{address}_{\mathsf{Srv}}))\}_{i\leq index \text{ and } \mathsf{address}_i \in \mathsf{S}_{\mathsf{Usr}}}\right|$$

where the tuples are recorded in $\mathcal{F}_{bb}$, and for every i it holds that: dtxn_i and wtxn_i are valid deposit and withdrawal transactions, respectively.

Intuitively, the definition of privacy captures the idea that an adversary should not be able to, given two deposit transactions and a withdrawal transaction belonging to one of the deposits, distinguish which of the deposits the withdrawal belongs to. This should hold true even if the adversary gets to freely interact with system.

Privacy. We define privacy for a mixer protocol $\Pi = (\mathsf{deposit}, \mathsf{withdraw})$ using the following experiment between a challenger Chal and an adversary $\mathcal{A}$ with access to $\mathcal{F}^{\mathcal{Q}}_{ban}$:

The privacy experiment $\mathsf{EXP}_{\mathcal{A},\Pi,\mathcal{F}^{\mathcal{Q}}_{ban}}(\lambda)$**:**

(1) The adversary $\mathcal{A}$ can send to Chal a deposit and withdrawal requests that are processed by Chal as follows:
   - *Honest deposit generation:* Upon receiving a (deposit) request from $\mathcal{A}$ it executes deposit and simulates $\mathcal{F}_{bb}$ using the stored values. Then Chal stores $(r, \mathsf{dtxn}, \mathsf{address}_A) \leftarrow \mathsf{view}^{\mathsf{deposit}}_{\mathsf{Usr}}(\lambda)$ and sends $(\mathsf{dtxn}, \mathsf{address}_A)$ to $\mathcal{A}$.
   - *Honest withdrawal generation:* Upon receiving a (withdrawal, dtxn) request from $\mathcal{A}$ it fetches r that is associated with dtxn (if no such exists return $\perp$ to $\mathcal{A}$), it executes withdraw on input r and simulates $\mathcal{F}_{bb}$ using the stored values. Then Chal stores $(\mathsf{wtxn}, \mathsf{address}_B) \leftarrow \mathsf{view}^{\mathsf{withdraw}}_{\mathsf{Usr}}(r, \lambda)$ and sends wtxn to $\mathcal{A}$.

- *Adversarial deposit/withdrawal submission:* In addition, $\mathcal{A}$ can submit a deposit or withdrawal transaction of its choice to Chal that records it in the simulated $\mathcal{F}_{bb}$ with appropriate *index*.

(2) $\mathcal{A}$ outputs a pair of deposit transactions $\mathsf{dtxn}_0, \mathsf{dtxn}_1$ that correspond to two deposit transactions previously generated by Chal and were not requested to be withdrawn in item 1.
(3) Chal chooses a random bit $b \in \{0,1\}$ and fetches r_b that is associated with dtxn_b. Then it executes withdraw on input r_b and send to $\mathcal{A}$ the generated withdrawal transaction i.e., $(\mathsf{wtxn}, \mathsf{address}_B) \leftarrow \mathsf{view}_{\mathrm{Chal}}^{\mathsf{withdraw}}(\mathsf{dtxn}_b)$. We call wtxn the challenge transaction. $\mathcal{A}$ continues to have access to Chal, interacting as in item 1.
(4) The adversary $\mathcal{A}$ outputs a bit b'. The experiment's output is defined to be 1 if $b' = b$, and 0 otherwise.

Definition B.3 (Privacy). A protocol $\Pi = (\mathsf{deposit}, \mathsf{withdraw})$ is *private* if for all ppt adversaries $\mathcal{A}$, there exists a negligible function $\mathrm{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\Pr[\mathsf{EXP}_{\mathcal{A},\Pi,\mathcal{F}_{ban}^{Q}}(\lambda) = 1] \leq \frac{1}{2} + \mathrm{negl}(\lambda)$$

where the probability is taken over the random coins used by $\mathcal{A}$ and Chal.

Intuitively, a compliant protocol should not allow the flow of illicit funds through the mixer. Additionally, compliance is somewhat meaningless for non sound protocols, i.e., ones that release funds without an appropriate assurance of their deposit by the withdrawing entity. Therefore, we focus our attention on compliance for sound protocols, according to definition B.2, and define compliance as follows: Our definition considers an idealized world where compliance is enforced by an ideal compliant ledger that "magically" deletes deposits from non compliant addresses, as if they never happened. Our compliance definition requires the protocol to behave indistinguishably when executed in our standard (append-only) ledger and the idealized world. In particular, we require that any valid withdrawal transaction is also valid in the idealized world, where no deposits from non-compliant addresses reside in the mixer. Combined with soundness this guarantees that the protocol enables withdraw funds only for compliant deposits. Formally,

Definition B.4 (Compliance). Let $\mathcal{F}_{bb}$ and $\mathcal{F}_{ban}^{Q}$ be the functionalities from fig. 2 and fig. 3, respectively, and let $\Pi = (\mathsf{deposit}, \mathsf{withdraw})$ be a *sound* mixer protocol as defined in definition B.2, with all entities having access to $\mathcal{F}_{bb}$ and $\mathcal{F}_{ban}^{Q}$. We say that Π is *compliant* if the following holds:

- The mixer Srv is stateless (i.e., it does not maintain state between executions of deposit or withdraw and it only performs Read requests to $\mathcal{F}_{bb}$).
- The user Usr only performs Upload requests to $\mathcal{F}_{bb}$ in deposit.
- for every tuple $(\mathit{index}, \mathsf{wtxn}, (\mathsf{address}_B, \mathsf{address}_{\mathsf{Srv}}))$ in $\mathcal{F}_{bb}$, it holds that:

  wtxn is a valid withdrawal transaction if and only if

  $$\Pr[\mathsf{out}_{\mathsf{Srv}}^{\mathsf{withdraw}(\mathcal{F}_{bb}^{*})}(\mathsf{wtxn}, \lambda) = 1] = 1$$

  where $\mathsf{withdraw}(\mathcal{F}_{bb}^{*})$ is the withdraw protocol of Π, where calls to $\mathcal{F}_{ban}^{Q}$ are ignored and $\mathcal{F}_{bb}$ is replaced with the following functionality $\mathcal{F}_{bb}^{*}$:
  - Write and Read requests are treated as in $\mathcal{F}_{bb}$.
  - every request $(\mathsf{Ban}, \mathsf{address}_A, \mathsf{data})$ from address Q to $\mathcal{F}_{ban}^{Q}$ is forwarded to $\mathcal{F}_{bb}$ and treated by overwriting every tuple

    $$(\mathit{index}, \mathrm{msg}, (\mathsf{address}_A, \mathsf{address}_{\mathsf{Srv}}))$$

    in $\mathcal{F}_{bb}$ to $(\mathit{index}, 0, (\perp, \mathsf{address}_{\mathsf{Srv}}))$.[3]

We note that though Srv in our protocol in fig. 4 is not stateless, it can be equivalently defined as stateless that does not store the entire tree $\mathcal{T}$, but rather recreates it with each withdrawal call by reading the bulletin-board $\mathcal{F}_{bb}$. We avoid this in order to increase efficiency. Therefore compliance of our non-stateless protocol follows.

C RELEASING NON-COMPLIANT FUNDS

While both Haze and Daze are compliant in the sense that funds deposited from banned addresses cannot be withdrawn, there exists the problem of these funds being permanently locked in the mixer. For this reason, we propose a mechanism that enables releasing these funds to a predetermined entity. We deal with this problem separately for Haze and for Daze.

releasing non-compliant funds in Haze. Recall that due to the privacy property of Haze, it is indistinguishable whether a non-compliant user withdrew its funds or not, and thus counting the amount of non withdrawn funds from banned addresses inside Haze is difficult. Our solution to release non-compliant funds works as follows: the mixer will have a limited life-cycle of some predetermined amount of time[4]. At the end of its life-cycle, the mixer will no longer take new deposits and there will be a period in which all users are allowed to withdraw their remaining funds. At the end of this period, all funds that are not withdrawn are transferred to the predetermined entity. This entity can implement a dispute process in which users with compliant funds that for some reason did not withdraw their funds in time can request their funds by exposing the (k, r) that are associated with the disputed funds (this causes a privacy loss for the user). At the end of the mixer's life-cycle, Haze also enables refunding compliant users the fee overhead they paid at the time of deposit. We remind that the purpose of this overhead is to cover the cost of the tree update in the event a deposit would ever become non compliant. Since at the end of the mixer's life-cycle the deposit is still compliant, this cost will never be realized. Therefore, Haze can safely refund the remaining fees to each compliant address that deposited funds to the mixer. Note that non-compliant users are punished by not being refunded their fee overhead, as their funds were used to fund the tree updates. This holds even in the case that the user managed to withdraw their deposit before becoming non-compliant.

[3] A similar treatment could suggest removing the records from $\mathcal{F}_{bb}$ instead of overwriting, however for the ease of presentation we define it as above.
[4] Time is measured in blocks.

Recall that on the blockchain, a smart contract can only be triggered by a transaction signed by a user. We therefore design Haze in a way that enables anyone, including the trusted entity, to trigger the end-of-life of the mixer. There will be a hard-coded condition in the smart contract to verify that enough blocks have been added since the creation of the contract, and only upon meeting that requirement, can the end-of life be triggered. The refund to compliant users guarantees the incentive to trigger the end-of-life of the mixer.

This construction maintains the properties of Haze: compliance, correctness, privacy, and soundness.

Releasing non-compliant funds in Daze. In Daze we can distinguish if a non-compliant deposit has already been withdrawn or not, since non-compliant users become de-anonymized. This makes it so that the mixer can count non-compliant funds inside the mixer at any desired period during the life-time of the mixer. This amount can be released to some predetermined trusted entity.