

Balancing Privacy and Accountability in Blockchain Identity Management*

Ivan Damgård¹, Chaya Ganesh², Hamidreza Khoshakhlagh¹,
Claudio Orlandi¹, and Luisa Siniscalchi¹

¹Concordium Blockchain Research Center, Aarhus University, Denmark

²Indian Institute of Science, Bangalore, India

Abstract

The lack of privacy in the first generation of cryptocurrencies such as Bitcoin, Ethereum, etc. is a well known problem in cryptocurrency research. To overcome this problem, several new cryptocurrencies were designed to guarantee transaction privacy and anonymity for their users (examples include ZCash, Monero, etc.).

However, the anonymity provided by such systems appears to be fundamentally problematic in current business and legislation settings: banks and other financial institutions must follow rules such as “Know Your Customer” (KYC), “Anti Money Laundering” (AML), etc. It is also well known that the (alleged or real) anonymity guarantees provided by cryptocurrencies have attracted ill-intentioned individuals to this space, who look at cryptocurrencies as a way of facilitating illegal activities (tax-evasion, ransom-ware, trading of illegal substances, etc.).

The fact that current cryptocurrencies do not comply with such regulations can in part explain why traditional financial institutions have so far been very sceptical of the ongoing cryptocurrency and Blockchain revolution.

In this paper, we propose a novel design principle for identity management in Blockchains. The goal of our design is to maintain privacy, while still allowing compliance with current regulations and preventing exploitations of Blockchain technology for purposes which are incompatible with the social good.

*Research supported by: the Concordium Blockchain Research Center (COBRA), Aarhus University, Denmark; the Carlsberg Foundation under the Semper Ardens Research Project CF18-112 (BCM); the European Research Council (ERC) under the European Unions’s Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO); the European Research Council (ERC) under the European Unions’s Horizon 2020 research and innovation programme under grant agreement No 803096 (SPEC);

Contents

1	Introduction	3
2	Preliminaries and Building Blocks	7
2.1	Notation	7
2.2	Pseudorandom Functions	7
2.2.1	Instantiation with Dodis-Yampolskiy PRF.	7
2.3	Blind Signature Schemes	8
2.4	(Ad-hoc) Threshold Encryption Scheme	8
3	System Design	9
3.1	Entities Involved	9
3.2	Data Objects	10
3.3	Protocols	11
3.4	Informal Analysis of the Design	12
4	ID-layer Formalization	12
4.1	The UC-Model	12
4.2	The ID Layer Functionality	13
4.3	Issuing Credentials – the Functionality	15
5	Formal Protocols Specifications	16
5.1	Identity Layer Protocol	16
5.2	Proof of Security for Identity Layer	18
5.3	Credential Issue Protocol	22
5.4	Proof of Security for Issue Protocol	22
6	Putting Everything Together	23
A	Additional Definitions	30
A.1	Bilinear groups	30
A.2	Commitment Schemes	30
A.3	(Blind) Signature Scheme	31
A.3.1	Instantiation with PS Signature scheme [PS16]	31
A.4	Threshold Encryption Scheme	33
A.4.1	Construction based on Share and Encrypt paradigm.	33
A.5	Secret Sharing Scheme	34
A.6	Non-Interactive Zero-Knowledge Proofs	35
B	Implementation Details	36
B.1	Auxiliary Σ -protocols.	37
B.2	Using Fiat-Shamir for non-interactive Proof	37
B.3	Sigma protocol for proving knowledge of discrete log	37
B.4	Sigma protocol for proving equality of committed value and Elgamal encrypted value	38
B.5	Proof of Equality of Aggregated Discrete Logs & Commitments	38
B.6	Proof of Aggregated Discrete Logs	39

B.7	Proof of Equality for Commitments in Different Groups	39
B.8	Proof of multiplicative relation on committed values	40
C	Protocols for ID layer	41
C.1	Protocol for Π_{issue}	41
C.2	Protocol for $\Pi_{\text{id-layer}}$	41
D	CL Framework	43
D.1	Definition	43
D.2	Hard Subgroup Membership Problem	44
D.3	PKE scheme under HSM-CL assumption	44
E	Transaction Layer	45
E.0.1	Plaintext Transactions.	46
E.0.2	Encrypted Transactions.	46
E.0.3	Compressing an account.	46
F	Universally Composable Non-interactive Zero-Knowledge	47
F.1	Simulation Extractable NIZK	47
F.2	SNARK-Lifting transformations via $C\emptyset C\emptyset$ framework [KZM ⁺ 15]	48
G	Standard Ideal Functionalities	48
G.0.1	Registration Functionality	50
G.0.2	NIZK Functionality	50
G.0.3	Ledger Functionality	50
G.0.4	Functionality for MPC of PRF output	51

1 Introduction

Early applications of blockchain to payment systems such as Bitcoin do not guarantee privacy. In the Bitcoin blockchain, blocks posted on the public ledger consist of transactions, making Bitcoin transparent – transactions are there for everybody to see. However, the identities are pseudonymous, and not tied to real world identities. Consequently, Bitcoin has the property that while the ownership of money is implicitly anonymous, the flow of money is globally visible. While this was perceived to be truly anonymous early on, there have been several works that deanonymize Bitcoin flow by analysing the payment graph [MPJ⁺13]. To overcome the problem of lack of privacy in the first generation of cryptocurrencies such as Bitcoin, Ethereum, etc., new systems were designed to guarantee transaction privacy and anonymity for their users [BCG⁺14, PBF⁺19, FMMO19]. Systems like Zerocash [BCG⁺14] fully hide both the value inside a transaction, and the sender and receiver identities. Most blockchain use-cases, however, are hindered by complete privacy as they need *accountability* and *identity management*. Privacy-preserving systems like ZCash are not designed with accountability in mind ¹. In order to conform with regulations like “Know your customer” (KYC) and “Anti-money laundering” (AML), a legal authority should be able to learn the value and identities of the parties involved in any transaction; this requirement seems to be at odds

¹Zcash considers solutions to implement AML and KYC controls [Zca], however this solution requires trust on a single party.

with privacy. The seemingly contradictory requirements of transaction privacy & user anonymity, and regulatory requirements such as KYC/AML imposed on financial and banking institutions is a major hurdle in widespread adoption of the blockchain.

Our Contribution. In this work, we address the problem of balancing accountability with privacy in blockchain-based systems. We propose a new architectural design of an “identity layer” that will provide privacy for its users – that is, no one, observing the network transactions and the status of the Blockchain should be able to learn about the identity of the owner of any account in the system. At the same time, the identity layer achieves accountability in the sense that in the presence of a reasonable suspicion, law-enforcement agencies (or other authorized parties), will be able to access the transaction history of a given user and/or block its funds, in a way similar to what is guaranteed today by traditional financial institutions. We develop cryptographic mechanisms that enhance accountability measures against misuse of the blockchain, while still providing privacy. Towards this end, we employ cryptographic techniques to design provably secure protocols, with both privacy and accountability guarantees. We prove the security of our constructions in the Universal Composability (UC) framework. We provide a high-level overview of the design of the system, and then discuss the techniques and cryptographic tools used. We believe that such an identity layer design will make Blockchain and cryptocurrencies more attractive for regulators, public institutions and traditional businesses which are interested in complying with existing legislation. In fact, the identity layer of Concordium², an upcoming major Blockchain project, is based on the design presented in this paper.

Overview of the System. In the proposed system, the identity and credentials of each participant in the network are initially verified and stored by authorized parties called *Identity Providers* (IPs). Each user can open a limited number of accounts where an account has an identifier that is derived from a PRF applied to a value that is between 1 and the maximal number of accounts, say n . The PRF key K is held by the user. When a user registers with an IP, K is encrypted through a threshold encryption scheme, and this ciphertext is stored with the IP. This is set up such that an appropriate number of *Anonymity Revokers* (ARs) would be able to decrypt. Standard anonymous credentials are used to certify additional attributes of the user.

When a user creates an account, they prepare some data to be published on the blockchain. This includes a threshold encryption of the account holder’s public key (that was also stored with the IP at registration time). It also includes zero-knowledge proofs that the attributes the user chooses to publish in the account have been signed by the IP, and that the account identifier has been correctly computed. Thus, an account may contain complete identification of the account holder, if the user chooses to include it, or it may reveal less information, for instance the citizenship and age of the account holder.

Finally, an account includes various account specific public keys. Using the corresponding secret keys, the account holder can then perform transactions anonymously in the network. Depending on the key material included in accounts, several different ways to do transactions can be realized – this is a problem orthogonal to that of implementing the identity layer, and we give some informal examples of how this could be done in Section E.

If it is suspected that an account is used for fraudulent purposes, the encrypted account information can be decrypted by a qualified set of the ARs, and an anonymous account can be linked

²concordium.com

(via the public key) to an id provided by the IP. On the other hand, if a particular user is suspected of fraud, the IP can provide its record for this user, and a qualified set of ARs can decrypt the information to learn the PRF key K . Now, they can generate the set of all values $\text{PRF}_K(x)$ for $x = 1, \dots, n$ which are all the possible values for an account identifier. One can then identify all accounts of the user by searching the blockchain for accounts with these identifiers. Privacy is therefore guaranteed for all users, except those whose anonymity is revoked by a sufficient number of ARs.

Note that the above also implies that we have a mechanism for preventing a user from opening an unbounded number of accounts using a single certificate from the IP: if this were possible, it would open the door for attacks where an individual registers with an IP and then allow other individuals to open accounts in their name, perhaps after payment of a small sum of money. On the other hand, we do not want the account holder to have to interact with the IP for every new account it wants to create, as this would affect efficiency. While the concrete number of accounts allowed per user is an implementation dependent parameter, our technique allows to achieve a reasonable tradeoff. The zero-knowledge proofs force the user to compute the account id's correctly, which only allows n different id's. Thus, if one attempts to open more accounts than allowed, this must result in a pre-existing account identifier, and the Blockchain will reject it.

We prove security of the system when either any number of account holders are actively corrupt, or when the identity providers are semi-honest corrupt. Note that, similar to certification authorities in standard PKI, we need some trust in the IPs: a malicious identity provider (equivalently, a malicious account holder colluding with a semi-honest IP and therefore learning the key), could produce certificates containing false identities, therefore undermining the system. Finally, depending on which properties we want to emphasize, we could tolerate different corruption levels among the anonymity revokers. Thus our system is secure in the presence of actively corrupt users and a threshold number of passively corrupt anonymity revokers; or, in the presence of passively corrupt identity-provider and a threshold number of passively corrupt anonymity revokers. In our design it is paramount that the service provided by the anonymity revokers to be available, and we want to emphasize privacy. Thus, we opt for assuming a majority of semi-honest ARs. Using standard methods, we could instead tolerate a minority of actively corrupted ARs.

Overview of technical ideas. We use cryptographic schemes such as Pedersen commitments [Ped92], Dodis-Yampolskiy PRF [DY05], Pointcheval-Sanders (PS) signature scheme [PS16], CL encryption scheme [CL15]. We use zkSNARKs in combination with commitments and signatures in the spirit of [AGM18, CFQ19]: the PS blind signature we use is defined using groups of a certain prime order, and when a user proves knowledge of a signature, the message that is signed is committed to using a Pedersen-type commitment in such a group. Now, we can use standard sigma-protocols to provide commitments to individual attributes of the user in the same group, and finally use SNARKs on committed messages to show statements such as “the age attribute of the user is a number greater than 18”. In this way, we only need to use SNARKs on rather small circuits, and we can achieve much greater efficiency than if we had to convert large statements involving, e.g, group operations into a Boolean circuit to be evaluated inside the SNARK. In this way, creating an account requires a constant number of exponentiations (i.e., independent of the security parameter), and likewise, the number of group elements in an account is constant.

We provide a generic lifting transformation for Fiat Shamir NIZKs for DL-languages into UC NIZKs. While such a transformation by encrypting the witness under a key that is part of the CRS

(and the secret key part of the CRS trapdoor) is folklore [DP92], using the CL encryption scheme allows us to efficiently prove statements about values in the exponent, which is novel to the best of our knowledge.

Related Work. The cryptographic tools used in building our solution, like commitment schemes, blind signatures, zero-knowledge proofs, and threshold encryption are based on anonymous credentials technology. Anonymous credentials [Cha83] allow a party to prove to a verifier that one has a set of credentials without revealing anything beyond this fact. Revocable anonymity [CMS96, KTY04] allows a trusted third party to discover the identity of all otherwise anonymous participants. Conditional anonymity requires that a user’s transactions remain anonymous until certain conditions are violated [CHL05, DDP06, CHK+06]. In [DDP06], an unclonable identification scheme is introduced, that is, roughly, an identification scheme where honest users can identify themselves anonymously as members of a group, but where clones of users can be detected and have their identities revealed if they identify themselves simultaneously. This was extended from one-time authentication to n -times anonymous authentication in [CHK+06] where a certain number of unlinkable accounts are derived that can later be efficiently traced. The works of [TFS04, NS05, ASM06, TS06] addressed related problems of allowing a user to show a credential anonymously and unlinkably up to n times to a particular verifier. The potential for abuse of unconditional anonymity by misbehaving users has been articulated in the context of group signatures. In a group signature scheme, each group member can sign a message on behalf of a group such that anyone can verify that the group signature is produced by someone in the group, but not who exactly. Our idea for identifying all accounts of a user in case of revocation by using a PRF to generate account identifiers is reminiscent of the work of traceable signatures [Cho09] that enable a tracing agent to identify all signatures produced by a particular member. The idea of deriving a certain number of unlinkable accounts that can later be efficiently traced has been used in various forms in the anonymous credential literature [CHK+06, ASMC12, Cho09] for the purposes of balancing accountability and anonymity.

Unfortunately none of the previous works seem to fit our intended use case, which motivated us to design the system described in this paper. Moreover, the toolbox of efficient tools available to the protocol designer has grown in recent years (e.g., the CL encryption scheme, advances in SNARKs, etc.), which also motivates exploring new designs.

The zkLedger protocol [NVV18] is an asset transfer scheme that hides transaction amounts and sender-receiver relationship, and supports auditing. The protocol is for a setting where the transacting parties are banks, and requires the participation of the banks for an audit to take place. The work of [ACC+20] presents a privacy-preserving token management system that supports auditing in permissioned blockchains. The system of [ACC+20] is in the UTXO framework, where users own tokens that are certified, and prove ownership of tokens in a privacy-preserving manner. In contrast, we work in the account-based model; and our design is modular – the identity layer is separate from the transaction layer. One main difference of our work from the works of [NVV18] and [ACC+20] is that while both these works assume that the entire system is permissioned, again our design is more modular: our ID layer obviously assumes that IPs and ARs are known (and trusted to some extent) and is therefore in some sense permissioned. However, the ID layer can work on top of the consensus mechanism of a permissionless blockchain, i.e., any blockchain that can be abstracted using the ledger functionality $\mathcal{F}_{\text{ledger}}$.

Finally, Solidus [CZJ+17] is a privacy-preserving system that allows customers of financial institutions (e.g., banks) to transfer assets and ensures that only the banks of the sender and

receiver can learn the transaction details. While there is no explicit audit functionality in Solidus, banks can reveal the content of a suspicious transaction to the authorized auditors. However this approach requires to trust a single party (i.e., the bank).

2 Preliminaries and Building Blocks

This section defines our notation and introduces the cryptographic schemes we use in our construction.

2.1 Notation

For any positive integer n , $[n]$ denotes the set $\{1, \dots, n\}$. We write $f(\lambda) \approx_\lambda g(\lambda)$ if the difference between f and g is negligible in λ . We use DPT (resp. PPT) to mean a deterministic (resp. probabilistic) polynomial time algorithm. We denote by $Y \leftarrow_{\$} F(X)$ a probabilistic algorithm F that on input X outputs Y . Similarly, notation $Y \leftarrow F(X)$ is used for a deterministic algorithm with input X and output Y . All adversaries will be stateful. We use the identifier **AH** for account holder, **IP** for identity provider and **AR** for anonymity revoker. By an identifier, we mean an arbitrary string that uniquely identifies a party. Throughout the paper, \mathbb{F}_q will denote the field with q elements.

2.2 Pseudorandom Functions

We recall the standard notion of pseudorandom functions.

Definition 2.1 (PRF). *Let $PRF: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}$ be a family of functions and let Γ be the set of all functions $\mathcal{D} \rightarrow \mathcal{R}$. We say that PRF is a pseudorandom function (PRF) (family) if it is efficiently computable and for all PPT distinguishers \mathcal{D}*

$$\left| \Pr \left[\mathcal{K} \leftarrow_{\$} \mathcal{K}, \mathcal{D}^{PRF_{\mathcal{K}(\cdot)}}(1^\lambda) \right] - \Pr \left[g \leftarrow_{\$} \Gamma, \mathcal{D}^{g(\cdot)}(1^\lambda) \right] \right| \approx_\lambda 0.$$

We define a weak notion of PRF robustness, meaning that it should be hard to find a key that produce collisions with the PRF evaluation of an honest user. Our definition is similar to the one in [FOR17], but here one of the two keys is chosen honestly.

Definition 2.2 (Weakly Robust PRF). *A PRF is weakly robust if:*

$$\Pr[\mathcal{K} \leftarrow_{\$} \text{Gen}(1^\lambda), (x^*, \mathcal{K}^*) \leftarrow_{\$} \mathcal{A}^{PRF_{\mathcal{K}(\cdot)}}(1^\lambda) : \exists (x, y) \in \mathcal{Q}, PRF_{\mathcal{K}^*}(x^*) = y] \approx_\lambda 0$$

where \mathcal{Q} is the set of inputs/outputs of the oracle available to the adversary.

2.2.1 Instantiation with Dodis-Yampolskiy PRF.

We use the PRF of Dodis and Yampolskiy [DY05] that operates in a group \mathbb{G} of order q with generator g . On input x and the PRF key $\mathcal{K} \leftarrow_{\$} \mathbb{F}_q$, $PRF_{\mathcal{K}}(x) = g^{1/\mathcal{K}+x}$. This is shown to be pseudorandom under the Decisional Diffie-Hellman Inversion assumption in group \mathbb{G} . Note that the security holds only for small domains, namely inputs that are slightly superlogarithmic in the security parameter, but this is sufficient for our work, as the maximum number of accounts a user can open is less than a constant Max_{ACC} . It can also be easily shown that the Dodis-Yampolskiy

PRF is weakly robust: using the PRF assumption, we can replace the output of the PRF oracle with random group elements. If the adversary outputs an input x^* and key K^* that are compatible with one of the output of the oracles, we can compute the discrete logarithm of that element as $1/(K^* + x^*)$.

2.3 Blind Signature Schemes

We adapt the notation of [SU12] to two-round blind signature schemes.

Definition 2.3 (Blind Signature Schemes). *An interactive signature scheme between a signer S and user U consists of a tuple of efficient algorithms $BS = (\text{Setup}, \text{KeyGen}, \text{Sign}_1, \text{Sign}_2, \text{Unblind}, \text{VerifySig})$ where*

- $\text{Setup}(1^\lambda)$, on input the security parameter 1^λ outputs pp , which is given implicitly as input to all other algorithms, even when omitted.
- $\text{KeyGen}(\text{p})$, on input the public parameter p generates a key pair (sk, pk) for security parameter λ .
- $\text{Sign}_1(\text{pk}, m)$, which is run by U , takes as input pk and a message $m \in \{0, 1\}^*$ and outputs sign_1 and ω (w.l.o.g. ω can be thought of as the randomness used to run Sign_1).
- $\text{Sign}_2(\text{sk}, \text{sign}_1)$, which is run by S , takes as input sk and sign_1 and outputs sign_2 .
- $\text{Unblind}(\text{sign}_2, \omega)$, which is run by U , takes as input sign_2, ω and outputs σ .
- $\text{VerifySig}(\text{pk}, m, \sigma)$ outputs a bit.

Remark 2.4. Note that a blind signature scheme implicitly defines a normal signature scheme as well, where the signing algorithm $\text{Sign}(\text{sk}, m)$ simply emulates a blind signature protocol and outputs the resulting signature σ .

The correctness property of the scheme requires that the following holds: for any $(\text{pp}) \leftarrow \$ \text{Setup}(1^\lambda)$, $(\text{sk}, \text{pk}) \leftarrow \$ \text{KeyGen}(\text{pp})$, any message $m \in \{0, 1\}^*$, if $(\text{sign}_1, \omega) \leftarrow \$ \text{Sign}_1(\text{pk}, m)$, $\text{sign}_2 \leftarrow \$ \text{Sign}_2(\text{sk}, \text{sign}_1)$, $\sigma = \text{Unblind}(\text{sign}_2, \omega)$ then $\text{VerifySig}(\text{pk}, m, \sigma) = 1$ with overwhelming probability over $\lambda \in \mathbb{N}$.

We require the standard notion of *existential unforgeability under chosen message attacks* (EUF-CMA) [GMR88]. The blind signature scheme we use should additionally satisfy two properties, namely *Blindness* and *simulatability*, where the second is an ad-hoc definition required for our UC proof of security that ensures the existence of an additional simulation algorithm Sim that can simulate sign_2 . The formal definition of these properties can be found in Appendix A.3.

2.4 (Ad-hoc) Threshold Encryption Scheme

We recall the definition of an ad-hoc threshold encryption scheme here.

Definition 2.5. *A (n, d) -threshold encryption scheme $TE = (\text{TKeyGen}, \text{TEnc}, \text{ShareDec}, \text{TCombine})$ over message space M consists of the following algorithms:*

- $\text{TKeyGen}(1^\lambda)$ is a randomized key generation algorithm that takes the security parameter λ as input and returns a private-public key pair (sk, pk) .

- $\text{TEnc}_{\vec{\text{pk}}_R}^{n,d}(m)$, a probabilistic encryption algorithm that encrypts a message $m \in M$ to a set of public keys $\vec{\text{pk}}_R = \{\text{pk}_i\}_{i \in R}$ in such a way that any size $d+1$ subset of the recipient set should jointly be able to decrypt. We sometimes write $\text{TEnc}_{\vec{\text{pk}}_R}^{n,d}(m; r)$ when we want to be able to fix the value of the randomness r to a specific value.
- $\text{ShareDec}_{\vec{\text{pk}}_R, \text{sk}_i}^{n,d}(ct)$, on input a ciphertext ct and a secret key sk_i , outputs a decryption share μ_i .
- $\text{TCombine}_{\vec{\text{pk}}_R}^{n,d}(ct, \{\mu_i\}_{i \in I})$, a deterministic algorithm that takes a subset $I \subset [n]$ with size $d+1$ of decryption shares $\{\mu_i\}_{i \in I}$ and outputs either a message $m \in M$ or \perp .

We use the static security definition of Reyzin et al. [RSY18] for threshold encryption schemes which requires two properties, namely *static semantic security* and *partial decryption simulatability* as defined in Appendix A.4.

Definition 2.6. A threshold encryption scheme $\text{TE} = (\text{KeyGen}, \text{TEnc}, \text{ShareDec}, \text{TCombine})$ is (n, d) -statically secure if it is both (n, d) -statically semantically secure (Definition A.5) and (n, d) -partial decryption simulatable (Definition A.6)

The definitions of Commitments scheme, Secret Sharing and Zero-Knowledge can be found in Appendix A.

3 System Design

We give a high-level overview of the design of the identity layer in terms of the entities involved, data objects and protocols between the entities.

3.1 Entities Involved

The following entities are involved in our design:

- **Account Holders (AH):** those are individuals who hold accounts on the block-chain. We assume AHs possess some mean for performing legal identification (e.g., a passport), in the country where they live. They are interested in opening accounts and performing transactions on the blockchain but, before doing so, they have to register with an Identity Provider (IP).
- **Identity Provider (IP):** an identity provider is an entity that, as the name suggests, can provide a digital identity to an AH. The identity provider “authorizes” a user to open accounts on the blockchain, and therefore to perform transactions. Jumping ahead, when observing transactions on the blockchain, it should not be possible to find out the identity of an AH (not even for the IP itself), while everyone should be able to see which IP has authorized a given account, thus creating trust in the account.
- **Anonymity Revoker (AR):** anonymity revokers are parties which are involved in case where law-enforcement or other authorized entities need to be able to extract the identity of the owner of some account on the blockchain. We can make threshold assumptions on the AR and e.g., require that at least $d+1$ ARs must give an approval before the anonymity of a user is revoked.

3.2 Data Objects

We now describe the data objects that are held by the entities.

Account Holder Certificate (AHC). After an account holder registers with an identity provider, the AH obtains a certificate containing:

- A public identity credential $IDcred_{PUB}$ and a secret identity credential $IDcred_{SEC}$.
- A key K for a pseudorandom function PRF.
- One or more attribute lists AL such as some identifier, age, citizenship, expiration date, etc.
- A signature on $(IDcred_{SEC}, K, AL)$ that can be checked using pk_{IP} . A valid signature proves that an AH with attributes as in AL has registered with IP and has proved knowledge of $IDcred_{SEC}$ corresponding to $IDcred_{PUB}$.

Account Creation Information (ACI). Given an AHC, an account holder can create new accounts and post the corresponding ACI on the ledger, containing:

- $RegID_{ACC}$, an account registration ID. This is defined to be $RegID_{ACC} = PRF_K(x)$ where K is a key held by AH and signed by the IP, and where the account in question is the x 'th account opened by the AH based on a given AHC. If AH behaves honestly, then $RegID_{ACC}$ is unique for the account, and $x \leq Max_{ACC}$. The latter condition is enforced by the proof below, the former can be checked publicly.
- Anonymity revocation data: this is a threshold encryption $E_{ID} = TEnc_{PK_{AR}}^{n,d}(IDcred_{PUB})$, where any subset of size $d + 1$ of anonymity revokers are able to decrypt E_{ID} and obtain $IDcred_{PUB}$.
- The identity IP of the identity provider who did the signature in the AHC used for this account.
- An account specific public key pk_{ACC} . It will be used, for instance, to verify transactions related to the account.
- A policy P , which asserts some information about the attribute list AL .
- A proof π that can be checked using pk_{IP} and verifies that ACI can only be created by an AH that has obtained an AHC from IP, such that $P(AL) = \top$, where AH knows the secret keys corresponding to pk_{ACC} , as well as $IDcred_{SEC}$ corresponding to the $IDcred_{PUB}$ that was presented to the IP, and where $RegID_{ACC}$, E_{ID} and $E_{RegID} = TEnc_{PK_{AR}}^{n,d}(K)$ are correctly generated.

Identity Provider's information on Account Holder (IPIAH). This is the data record that the IP stores after an AH has registered. It contains:

- The name AH of the account holder and its public identity credential $IDcred_{PUB}$.
- A set of anonymity revokers AR_1, \dots, AR_n with public keys PK_{AR} and an encryption $E_{RegID} = TEnc_{PK_{AR}}^{n,d}(K)$. Here, K is the PRF key chosen by the AH at registration time.

3.3 Protocols

The following are the main protocols in our design.

Account Holder Registration. The protocol takes place between an IP and an AH who owns a key pair $(\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}})$ and an attribute list AL . At the end of the protocol, the AH receives an AHC and the IP obtains a IPIAH as described above. The AH sends their attribute list AL to IP and proves (via non-cryptographic means) their identity to IP. More concretely, this means that the IP must verify that the entity it is talking to indeed has the name AH and hence it received AL from the correct entity. It should also verify that the attributes in AL are correct w.r.t. the AH. The AH also sends to IP their public key $\text{IDcred}_{\text{PUB}}$ and an encryption $E_{\text{RegID}} = \text{TEnc}_{\text{PK}_{\text{AR}}}^{n,d}(\text{K})$ where K is a PRF key. Next, AH and IP engage in a blind signature scheme, which allows AH to receive a signature on $(\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL})$ that is generated under the secret key sk_{IP} of the IP. In addition, AH proves (cryptographically, in ZK) that they know $\text{IDcred}_{\text{SEC}}$ corresponding to $\text{IDcred}_{\text{PUB}}$, that the same $\text{IDcred}_{\text{SEC}}$ was input to the blind signature, and that the encryption contains the same K that was input to the blind signature scheme. IP stores $\text{IPIAH} = (\text{ID}_{\text{AH}}, \text{IDcred}_{\text{PUB}}, \text{AL}, E_{\text{RegID}}, \text{AR}_1, \dots, \text{AR}_n)$.

Create New Account. An account holder AH wants to create an account that satisfies some policy P (e.g., above 18, resident in country X, etc.). They take as input an AHC, a policy P and the public key pk_{AR} of one (or more) anonymity revoker(s) with name AR. At the end, AH produces some ACI that can be posted to the blockchain. They also need to store secret key sk_{ACC} that is specific to the account. The protocol works as follows: AH generates an account key pair $(\text{pk}_{\text{ACC}}, \text{sk}_{\text{ACC}})$ and an encryption of their public identity credential $\text{IDcred}_{\text{PUB}}$ under the public key of the anonymity revokers' PK_{AR} , i.e., $E_{\text{ID}} = \text{TEnc}_{\text{PK}_{\text{AR}}}^{n,d}(\text{IDcred}_{\text{PUB}})$. Next, AH calculates $\text{RegID}_{\text{ACC}} = \text{PRF}_{\text{K}}(x)$, where we assume this is the x 'th account that is opened using the AHC that is input. At last, AH produces a non-interactive zero-knowledge (NIZK) proof of knowledge π for statement

$$st = (\text{P}, E_{\text{ID}}, \text{RegID}_{\text{ACC}}, \text{IP}, \text{pk}_{\text{ACC}})$$

using secret witness

$$w = (\sigma, \text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}, \text{sk}_{\text{ACC}})$$

for the relation $R(st, w)$ that outputs \top if:

1. σ is a valid signature under pk_{IP} for a message of form $(\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL})$.
2. AL satisfies the policy i.e., $\text{P}(\text{AL}) = \top$.
3. $\text{RegID}_{\text{ACC}} = \text{PRF}_{\text{K}}(x)$ for some $x \leq \text{Max}_{\text{ACC}}$.
4. $E_{\text{ID}} = \text{TEnc}_{\text{PK}_{\text{AR}}}^{n,d}(\text{IDcred}_{\text{PUB}})$.
5. $(\text{pk}_{\text{ACC}}, \text{sk}_{\text{ACC}})$ is a valid key pair.

Let $\text{ACI} = (\text{RegID}_{\text{ACC}}, E_{\text{ID}}, \text{AR}_1, \dots, \text{AR}_n, \text{IP}, \text{pk}_{\text{ACC}}, \text{P}, \pi)$.

Revoke Anonymity of Account. Revocation of the anonymity of an account can be done by at least $d+1$ of the n ARs involved in the set-up of the account, working together with the IP with whom the AH registered. The input is an account identifier $\text{RegID}_{\text{ACC}}$ and the output is the name AH of the account holder. The protocol proceeds as follows: Given an account $\text{RegID}_{\text{ACC}}$ whose anonymity needs to be revoked, the ARs find the ACI containing $\text{RegID}_{\text{ACC}}$ on the blockchain, collaborate

to decrypt E_{ID} and learn $IDcred_{PUB}$. The registration information also contains the public name IP of the identity provider who registered $IDcred_{PUB}$. The ARs contact this IP who then locates the $IPIAH = (AH, IDcred_{PUB}, AL, E_{RegID}, AR_1, \dots, AR_n)$ record that contains the $IDcred_{PUB}$ that was decrypted. This record also includes AH, thus IP and the set of ARs have now identified the AH.

Trace accounts of User. If a user with a given name AH is suspected of engaging in illegal activities, the IP and a set of at least $d + 1$ ARs can identify all accounts of that user. The IP searches its database to locate the $IPIAH = (AH, IDcred_{PUB}, AL, E_{RegID}, AR_1, \dots, AR_n)$ containing the relevant AH. This record also contains the names of the relevant ARs. A qualified set of these could decrypt the E_{RegID} to learn the PRF key K and generate all values $PRF_K(x)$ for $x = 1, \dots, Max_{ACC}$ in public. However, due to technicalities in the security reduction, this would require the PRF to satisfy some form of “selective opening attack” security. Instead, we let the ARs decrypt the ciphertext and evaluate the PRF on $x = 1, \dots, Max_{ACC}$ inside an MPC protocol, so that K is never revealed to anyone. Either way, the produced values are all the possible values for $RegID_{ACC}$ that the AH could have used to form valid accounts, so one can now search the blockchain for accounts with these registration IDs.

3.4 Informal Analysis of the Design

If an AH misbehaves and opens more accounts than they are allowed to, this must result in two or more accounts with the same $RegID_{ACC}$. This can be publicly detected by the blockchain, and the second account will be discarded. We note that for this to work, we assume that incentives have been created so that some parties will indeed observe the duplicates and alert the relevant entities. Moreover, the construction satisfies *revocability* and *traceability*, meaning a malicious AH cannot create a valid account such that the anonymity revokers together with the identity provider are unable to revoke its anonymity or trace it. This follows from the soundness of the underlying zero-knowledge proofs which imply $E_{ID} = TEnc_{PK_{AR}}^{n,d}(IDcred_{PUB})$ and $E_{RegID} = TEnc_{PK_{AR}}^{n,d}(K)$. Thus, any subset of size $d + 1$ of anonymity revokers can decrypt E_{ID} (resp. E_{RegID}) and revoke the AH’s anonymity (resp. trace all the AH’s accounts). Lastly, due to the security of the underlying PRF and the threshold encryption scheme and also the ZK property of the proof $\pi \in ACI$, our design supports *anonymity* of the account holders, in the sense that a malicious identity provider even by cooperating with d anonymity revokers and other dishonest account holders cannot link a valid account to an account holder. Since we are using a Blockchain e.g., an imperfect bulletin board, we also need to worry that a malicious AH can’t “rush” and steal an honest user account number by maliciously choosing a PRF key K which “hits” some of the account numbers of the honest users which have not yet been finalized by the Blockchain. In order to do this we define and use a *weakly robust* PRF.

4 ID-layer Formalization

4.1 The UC-Model

We use the UC-security [Can01] framework with static corruption. In the following, the reader is assumed to be familiar with the basic concepts of UC security and is referred to [Can01] for a more detailed description.

Let \mathcal{Z} denote the environment. For a protocol Π and an adversary \mathcal{A} , we write $REAL_{\Pi, \mathcal{A}, \mathcal{Z}}$ to denote the ensemble corresponding to the protocol execution. For an ideal functionality \mathcal{F}

and a simulator \mathcal{S} , we write $\text{IDEAL}_{\mathcal{F},\mathcal{S},\mathcal{Z}}$ to denote the distribution ensemble of the ideal world execution. We say that a protocol Π UC-realizes a functionality \mathcal{F} if for all PPT adversaries \mathcal{A} corrupting a subset of parties, there exists a simulator \mathcal{S} such that for all environments \mathcal{Z} , the ensembles $\text{REAL}_{\Pi,\mathcal{A},\mathcal{Z}}$ and $\text{IDEAL}_{\mathcal{F},\mathcal{S},\mathcal{Z}}$ are computationally indistinguishable. In the rest of the section we describe the main ideal functionalities in our construction. We defer other (standard) ideal functionalities used by our protocol in Appendix G.

4.2 The ID Layer Functionality

The functionality $\mathcal{F}_{\text{id-layer}}$ captures the security properties offered by the design of our identity layer while hiding the implementation details. After the functionality is initialized, it allows identity providers IP to issue credentials to account holders AH based on their attribute lists AL. At the level of the ideal functionality, a credential is just a pointer to a record storing the tuple (AH, IP, AL). Armed with a credential, an AH can create up to Max_{ACC} accounts. When creating an account, the AH can choose a predicate P of their attributes to be made public (e.g., “I am over 18, I am resident of country X” etc.) which, together with the IP who authorized this account, are the only information which are made public. We capture this by having the functionality leak only the fact that an account was created and not the identity of the AH.³ Moreover, when creating an account, the AH also registers a key-pair associated to this account. The functionality is parametrized by any key-pair relation, which allows our ideal functionality to be used as a building block in more complex protocols, where the AH then can use those keys for authentication, encryption, etc. Our functionality also exposes some of the details about the underlying ledger on top of which it is implemented, thus new accounts are added to a buffer which can be permuted by the adversary before becoming finalized. This is inevitable as we run this on top of a ledger which has the same properties. The final two commands of the ideal functionality, revoke and trace, allow a qualified set of anonymity revokers AR and an IP to respectively disclose the AH behind a given account, or to find all accounts belonging to a certain AH.

Functionality Identity layer $\mathcal{F}_{\text{id-layer}}$

We assume that $\{\text{IP}_1, \dots, \text{IP}_m, \text{AR}_1, \dots, \text{AR}_n\}$ is the set of identifiers for identity providers and anonymity revokers. The functionality is parameterized by values m, n and threshold d , together with an NP (key-pair) relation \mathcal{R}_{ACC} such that when parties input `CreateACC`, they also specify a key-pair and the functionality verifies if the key-pair satisfies \mathcal{R}_{ACC} . Moreover, the functionality maintains the following initially empty records: `Count`, where `Count[cid]` counts the number of accounts created by certificate `cid`, and two records `Cert` and `ACC`, respectively for keeping track of certificates and accounts and a list L of public account information.

³Note that the environment provides all inputs and sees all outputs. It can therefore observe that an account is created right after it instructed an account holder to create an account, and can make the connection between the two. This corresponds to the fact that in a real application an adversary may know that in a long time interval, only one user creates an account, and so the next account that shows up on chain must belong to that user. Of course, our system cannot prevent this - the best we can do is to make sure that the account itself is anonymous. This follows in our model because the ideal adversary - the simulator - will not learn the identity of the holder and will still have to produce account information which are indistinguishable from the real protocol, thus proving that the account information leaks no information about its holder.

Initialize

On (INITIALIZE) from party $P \in \{IP_1, \dots, IP_m, AR_1, \dots, AR_n\}$, output to \mathcal{A} (INITIALIZED, P).

If all parties have been initialized, store (READY).

Issue

On (ISSUE, IP, AL) from an honest account holder AH (or the adversary in the name of corrupted account holder AH) and input (ISSUE, AH) from identity provider IP:

- If not (READY), then ignore.
- If there is already a cid with $\text{Cert}[\text{cid}] = (\text{AH}, \text{IP}, \cdot, \cdot)$, then abort; otherwise, if IP is honest (resp. corrupt), then send (ISSUE) (resp. (ISSUE, AH, IP, AL)) to \mathcal{A} .
- Upon receiving (cid, ISSUE) from \mathcal{A} , if $\text{cid} = \perp$ (in the case of corrupt IP) or if there already exists cid s.t. $\text{Cert}[\text{cid}] \neq \perp$, then abort. Otherwise, set $\text{Cert}[\text{cid}] \leftarrow (\text{AH}, \text{IP}, \text{AL})$ and output (ISSUED, cid) to AH.

Account Creation

Upon inputs (CreateACC, cid, P, (sk_{ACC}, pk_{ACC})) from honest account holders AH (or the adversary in the name of corrupted account holder AH), if not (READY), then ignore. Else, proceed as follows:

- If $\text{Cert}[\text{cid}] = \perp$ then abort, else retrieve (AH', IP, AL) \leftarrow $\text{Cert}[\text{cid}]$.
- Check if $\text{AH}' = \text{AH}$ and $\text{Count}[\text{cid}] < \text{Max}_{\text{ACC}}$ and that AL satisfies the policy P.
- Verify that the key pair (sk_{ACC}, pk_{ACC}) satisfies the relation \mathcal{R}_{ACC} and abort otherwise.
- Output (CreateACC, P, pk_{ACC}, IP) to \mathcal{A} .
- Receiving a response (CreateACC, aid) from \mathcal{A} , if $\text{aid} = \perp$ or $\text{ACC}[\text{aid}] \neq \perp$, then abort, else do the followings:
 - set $\text{ACC}[\text{aid}] \leftarrow (\text{cid}, \text{P}, \text{sk}_{\text{ACC}})$.
 - set $\text{Count}[\text{cid}] \leftarrow \text{Count}[\text{cid}] + 1$.
 - add the tuple (aid, P, pk_{ACC}, IP) to the account buffer.
- Return (Created, aid) to AH.

Account Buffer Release

Upon input (RELEASE, Π) from the adversary \mathcal{A} , remove all tuples from the account buffer and add the permuted tuples (aid, P, pk_{ACC}, IP) of accounts to the account list L .

Accounts Retrieve

On (RETRIEVE) from an account holder or party $P \in \{IP_1, \dots, IP_m, AR_1, \dots, AR_n\}$, output a list including all existing tuples (aid, P, pk_{ACC}, IP) in the account list L .

Revoke

Upon input (REVOKE, aid) from a (possibly corrupt) identity provider IP and a set of (possibly corrupt) anonymity revokers $\{\text{AR}_i\}_{i \in I \subseteq [n]}$, proceed as follows:

- If $\text{ACC}[\text{aid}] = \perp$ then return \perp . Otherwise, retrieve $(\text{cid}, \text{P}, \text{sk}_{\text{ACC}}) \leftarrow \text{ACC}[\text{aid}]$.
- If IP is the same as the identity provider in $\text{Cert}[\text{cid}]$ and $|I| > d$, then return (aid, AH) to the IP and $\{\text{AR}_i\}_{i \in I}$. Otherwise, return \perp . Moreover, if the identity provider or at least one anonymity revoker is corrupt, output (aid, AH) to \mathcal{A} as well.

Trace

Upon input $(\text{TRACE}, \text{AH})$ from a (possibly corrupt) identity provider IP and a set of (possibly corrupt) anonymity revokers $\{\text{AR}_i\}_{i \in I}$, proceed as follows:

- If there is no record $(\text{AH}, \text{IP}, \cdot, \cdot)$ in Cert , then return \perp . Otherwise, retrieve $(\text{AH}, \text{IP}, \cdot) \leftarrow \text{Cert}[\text{cid}]$.
- If $|I| > d$, return to IP and $\{\text{AR}_i\}_{i \in I}$ the list of all aids such that $\text{ACC}[\text{aid}] = (\text{cid}, \cdot, \cdot)$. Moreover, if the identity provider or at least one anonymity revoker is corrupt, return the list to \mathcal{A} as well.

4.3 Issuing Credentials – the Functionality

We describe here $\mathcal{F}_{\text{issue}}$, an ideal functionality capturing the desired properties of the issue protocol, which allows an identity provider to issue credentials to account holders. Note that the functionality can be seen as an augmented blind signature functionality: the account holder receives a signature (under the secret key of the identity provider) on a secret message m (as in blind signatures) but also on some public auxiliary information aux and on a secret key chosen by the account holder. The identity provider is not supposed to learn m (as in blind signatures), but in addition the identity provider learns a ciphertext which is guaranteed to contain an encryption of m and the public key corresponding to the secret key which is being signed.

Functionality Issue $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}$

The functionality is parametrized by an NP relation \mathcal{R} corresponding to the account holders key pair, a signature scheme $\text{SIG} = (\text{Setup}, \text{KeyGen}, \text{Sign}, \text{VerifySig})$ and a (n, d) -threshold encryption scheme $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{TDec})$. We assume that $\{\text{IP}_1, \dots, \text{IP}_m\}$ is the set of identifiers for identity providers.

Setup

- Upon input (SETUP) from any party, and only once, run $\mathbf{p} \leftarrow_{\$} \text{Setup}(1^\lambda)$ and return $(\text{SETUPREADY}, \mathbf{p})$ to all parties.

Initialize

- Upon input $(\text{INITIALIZE}, (\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}}))$ from identity provider IP, ignore if the party is already initialized or if SETUPREADY has not been returned yet. Otherwise, if $(\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}})$ is a valid key pair according to the relation defined by $\text{KeyGen}(\mathbf{p})$, store $(\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}})$ for this party and output $(\text{INITIALIZED}, \text{pk}_{\text{IP}}, \text{IP})$ to \mathcal{A} .

- If all parties have been initialized, store (READY).

Issue

On input $(\text{ISSUE}, (ct, m, r, \text{pk}_{\text{AR}}), \text{aux}, (\text{sk}_{\text{AH}}, \text{pk}_{\text{AH}}), \text{IP})$ from account holder AH and input $(\text{ISSUE}, \text{AH}, \text{pk}_{\text{AR}})$ from identity provider IP:

- If not (READY), then ignore.
- If $(\text{sk}_{\text{AH}}, \text{pk}_{\text{AH}}) \notin \mathcal{R}$ or $ct \neq \text{TEnc}_{\text{pk}_{\text{AR}}}^{\text{n,d}}(m; r)$ then abort.
- Otherwise compute $\sigma \leftarrow \text{Sign}((\text{sk}_{\text{AH}}, m, \text{aux}), \text{sk}_{\text{IP}})$.
- Output σ to AH and $(\text{pk}_{\text{AH}}, \text{aux}, ct)$ to IP.

5 Formal Protocols Specifications

5.1 Identity Layer Protocol

The protocol $\Pi_{\text{id-layer}}$ is run by parties interacting with ideal functionalities $\mathcal{F}_{\text{reg}}, \mathcal{F}_{\text{issue}}, \mathcal{F}_{\text{nizk}}, \mathcal{F}_{\text{crs}}, \mathcal{F}_{\text{ledger}}$ and $\mathcal{F}_{\text{mpc-prf}}$. Let \mathcal{R} and \mathcal{R}_{ACC} be NP relations corresponding to the account holders' key-pair and accounts' key-pair, respectively. Let $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{TDec})$ denote a threshold encryption scheme, PRF a pseudorandom function and $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{VerifySig})$ a signature scheme. Protocol $\Pi_{\text{id-layer}}$ proceeds as follows.

Protocol Identity layer $\Pi_{\text{id-layer}}$

Parameters for ideal functionalities.

- We use a $\mathcal{F}_{\text{ledger}}$ that implements the following VALIDATE predicate: the predicate accepts if the NIZK proof π is valid and if $\text{RegID}_{\text{ACC}}$ has not been seen before.
- We use a \mathcal{F}_{crs} functionality that outputs the public parameters for the signature scheme and the threshold encryption scheme.

The protocol description for an account holder AH.

- On input $(\text{ISSUE}, \text{IP}, \text{AL})$, retrieve the public key PK_{IP} of identity provider IP and the vector PK_{AR} of all public keys of the anonymity revokers via \mathcal{F}_{reg} and proceed as follows:
 - Generate a key pair $(\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}})$ satisfying \mathcal{R} .
 - Choose a random key K for PRF and compute the encryption $E_{\text{RegID}} = \text{TEnc}_{\text{pk}_{\text{AR}}}^{\text{n,d}}(K; r)$ with randomness r .
 - Call $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{BS}}$ on input $(\text{ISSUE}, (E_{\text{RegID}}, K, r, \text{PK}_{\text{AR}}), \text{AL}, (\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}}), \text{IP})$.
After receiving the response σ from $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{BS}}$, set $\text{cid} = (\text{IP}, \text{AL}, \text{IDcred}_{\text{SEC}}, \sigma, K)$.
- On input $(\text{CreateACC}, \text{cid}, P)$, proceed as follows

- If there is no record $\text{cid} = (\text{IP}, \text{AL}, \text{IDcred}_{\text{SEC}}, \sigma, \text{K})$, then abort.
- Generate an account key pair $(\text{pk}_{\text{ACC}}, \text{sk}_{\text{ACC}})$ satisfying \mathcal{R}_{ACC} .
- Compute $\text{E}_{\text{ID}} = \text{TEnc}_{\text{pk}_{\text{AR}}}^{\text{n,d}}(\text{IDcred}_{\text{PUB}}; r')$.
- Compute $\text{RegID}_{\text{ACC}} = \text{PRF}_{\text{K}}(x)$, where this is x 'th account that is created using cid .
- Produce a NIZK π by calling $\mathcal{F}_{\text{nizk}}$ for statement

$$st = (\text{P}, \text{E}_{\text{ID}}, \text{RegID}_{\text{ACC}}, \text{IP}, \text{pk}_{\text{ACC}})$$

using secret witness

$$w = (\sigma, x, r', \text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}, \text{sk}_{\text{ACC}}, \text{IDcred}_{\text{PUB}})$$

for the relation $\mathcal{R}(st, w)$ that outputs \top if:

1. The signature σ is valid for $(\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL})$ under pk_{IP} .
 2. AL satisfies the policy i.e., $\text{P}(\text{AL}) = \top$.
 3. $\text{RegID}_{\text{ACC}} = \text{PRF}_{\text{K}}(x)$ for some $0 < x \leq \text{Max}_{\text{ACC}}$.
 4. $\text{E}_{\text{ID}} = \text{TEnc}_{\text{pk}_{\text{AR}}}^{\text{n,d}}(\text{IDcred}_{\text{PUB}}; r')$.
 5. $(\text{pk}_{\text{ACC}}, \text{sk}_{\text{ACC}})$ is a valid key pair according to \mathcal{R}_{ACC} .
 6. $(\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}})$ is a valid key pair according to \mathcal{R} .
- Let $\text{ACI} = (\text{RegID}_{\text{ACC}}, \text{E}_{\text{ID}}, \text{IP}, \text{pk}_{\text{ACC}}, \text{P}, st, \pi)$ and $\text{SI}_{\text{ACC}} = \text{sk}_{\text{ACC}}$. Send the input $(\text{APPEND}, \text{ACI})$ to $\mathcal{F}_{\text{ledger}}$.
 - Store tuple $(\text{ACI}, \text{SI}_{\text{ACC}})$ internally and return $(\text{Created}, \text{RegID}_{\text{ACC}})$.
- On input (RETRIEVE) , call $\mathcal{F}_{\text{ledger}}$ on input RETRIEVE . After receiving $(\text{RETRIEVE}, L)$ from $\mathcal{F}_{\text{ledger}}$, output L .

The protocol description for identity providers and anonymity revokers.

- On input INITIALIZE from $P \in \{\text{IP}_1, \dots, \text{IP}_m\}$, obtain crs from \mathcal{F}_{crs} , generate key pair $(\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}}) \leftarrow_{\$} \text{KeyGen}(1^\lambda)$ and input $(\text{INITIALIZE}, (\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}}))$ to $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}$.
- On input INITIALIZE from $P \in \{\text{AR}_1, \dots, \text{AR}_n\}$, obtain crs from \mathcal{F}_{crs} , generate key pair $(\text{sk}_{\text{AR}}, \text{pk}_{\text{AR}}) \leftarrow_{\$} \text{TKeyGen}(1^\lambda)$ and input $(\text{REGISTER}, \text{sk}_{\text{AR}}, \text{pk}_{\text{AR}})$ to \mathcal{F}_{reg} .
- On input $(\text{ISSUE}, \text{AH})$ from identity provider IP , call $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}$ with input $(\text{ISSUE}, \text{AH}, \text{pk}_{\text{AR}})$. After receiving the response $(\text{IDcred}_{\text{PUB}}, \text{AL}, \text{E}_{\text{RegID}})$ from $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}$, set $\text{IPIAH} = (\text{AH}, \text{IDcred}_{\text{PUB}}, \text{AL}, \text{E}_{\text{RegID}})$.
- On input (RETRIEVE) , call $\mathcal{F}_{\text{ledger}}$ on input RETRIEVE . After receiving $(\text{RETRIEVE}, L)$ from $\mathcal{F}_{\text{ledger}}$, output L .
- On input $(\text{REVOKE}, \text{RegID}_{\text{ACC}})$ from an identity provider IP and a set of anonymity revokers $\{\text{AR}_i\}_{i \in I \subseteq [n]}$, proceed as follows:

- Anonymity revokers call $\mathcal{F}_{\text{ledger}}$ on input **RETRIEVE**. After receiving $(\text{RETRIEVE}, L)$ from $\mathcal{F}_{\text{ledger}}$, they first look up **ACI** in L that contains $\text{RegID}_{\text{ACC}}$. Next, each AR_i decrypts the E_{ID} of the **ACI** by computing $\mu_i = \text{ShareDec}_{pk_I, sk_i}^{n,d}(E_{\text{ID}})$. Finally, all anonymity revokers combine their shares and compute $\text{IDcred}_{\text{PUB}} = \text{TCombine}_{pk_I}^{n,d}(E_{\text{ID}}, \{\mu_i\}_{i \in I})$ and return $\text{IDcred}_{\text{PUB}}$ to the IP.
- The **ACI** contains the public name **IP** of the identity provider who registered $\text{IDcred}_{\text{PUB}}$. If **IP** is different from the requester, then ignore. Else, the IP locates the $\text{IPIAH} = (\text{AH}, \text{IDcred}_{\text{PUB}}, \text{AL}, E_{\text{RegID}})$ record, containing the $\text{IDcred}_{\text{PUB}}$ that was decrypted and outputs **AH**.
- On input $(\text{TRACE}, \text{AH})$ from an identity provider **IP** and a set of anonymity revokers $\{\text{AR}_i\}_{i \in I \subseteq [n]}$, proceed as follows:
 - IP searches the database to locate the $\text{IPIAH} = (\text{AH}, \text{IDcred}_{\text{PUB}}, \text{AL}, E_{\text{RegID}})$ containing the relevant **AH** and sends E_{RegID} via \mathcal{F}_{smt} to the set of anonymity revokers.
 - Each AR_i computes $K_i = \text{ShareDec}_{pk_I, sk_i}^{n,d}(E_{\text{RegID}})$. Then, they call $\mathcal{F}_{\text{mpc-prf}}$ on input $(\text{COMPUTE}, K_i)$ and receive all values $\text{PRF}_K(x)$ for $x = 1, \dots, \text{Max}_{\text{ACC}}$. These values are all the possible values for $\text{RegID}_{\text{ACC}}$ that the **AH** could have used to form valid accounts.

5.2 Proof of Security for Identity Layer

Tolerated Corruptions. We prove security in two separate cases: with arbitrarily many malicious **AH**s and up to threshold semi-honest **AR**s, or with semi honest **IP** and up to threshold semi-honest **AR**s. Note that for technical reasons we cannot let the **IP** be corrupt (even if only semi-honest) at the same time with a malicious **AH**, since in this case the (monolithic) adversary would learn the secret key of the corrupt **IP** and would be able to forge invalid credentials for the corrupt **AH**.

Assumptions on the environment. We consider executions with restricted adversaries and environments, that only input attribute lists **AL** in the **ISSUE** command which are valid with respect to the account holder. This restriction captures the fact that an honest **IP** in the real world is trusted to check (by non-cryptographic means) that an account holder **AH** actually satisfies the claimed attribute list **AL**.

Theorem 5.1. *Suppose that $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{TDec})$ is a (n, d) -threshold encryption scheme, PRF is a weakly robust pseudorandom function, \mathcal{R} a hard relation, and $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{VerifySig})$ is a EUF-CMA signature scheme, then $\Pi_{\text{id-layer}}$, for all restricted environment (as defined above), securely implements $\mathcal{F}_{\text{id-layer}}$ in the $\{\mathcal{F}_{\text{crs}}, \mathcal{F}_{\text{reg}}, \mathcal{F}_{\text{nizk}}, \mathcal{F}_{\text{smt}}, \mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}, \mathcal{F}_{\text{ledger}}, \mathcal{F}_{\text{mpc-prf}}\}$ -hybrid model in the presence of an actively corrupted **AH**, and d semi-honest anonymity revokers $\text{AR}_1, \dots, \text{AR}_d$, or semi-honest **IP** and d semi-honest anonymity revokers $\text{AR}_1, \dots, \text{AR}_d$.*

The simulator **Sim** is internally emulating the functionalities $\mathcal{F}_{\text{crs}}, \mathcal{F}_{\text{reg}}, \mathcal{F}_{\text{nizk}}, \mathcal{F}_{\text{smt}}, \mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}, \mathcal{F}_{\text{ledger}}, \mathcal{F}_{\text{mpc-prf}}$. At the start **Sim** initializes empty lists $\text{list}_{\text{issue}}, \text{list}_{\text{acc}}, \text{list}_{\text{ledger}}, \text{list}_{\text{h-aid}}, \text{list}_{\text{h-pk}}, \text{list}_{\text{h-sig}}$. Here is the description of the simulator:

- Command **INITIALIZE**: Sim emulates \mathcal{F}_{CRS} and generates public parameters for the signature scheme and threshold encryption scheme. Every time an honest or semi-honest IP or AR invokes the initialize command, the simulator generates a key-pair for them.
- Command **ISSUE**. *Passively corrupted IP, honest AH*. The simulator receives $(\text{ISSUE}, \text{AH}, \text{IP}, \text{AL})$ in case of corrupt IP from the ideal functionality $\mathcal{F}_{\text{id-layer}}$, and has to produce a view indistinguishable from the protocol which is consistent with this. The simulator does so by emulating the output $\mathcal{F}_{\text{issue}}$ for IP namely $(\text{IDcred}_{\text{PUB}}, \text{AL}, \text{E}_{\text{RegID}})$ to IP where the simulator encrypts a dummy value for E_{RegID} and $\text{IDcred}_{\text{PUB}}$ is generated according to \mathcal{R} . The simulator adds E_{RegID} to $\text{list}_{\text{h-ct}}$ and $\text{IDcred}_{\text{PUB}}$ to $\text{list}_{\text{h-pk}}$. Sim adds to $\text{list}_{\text{issue}}$ the entry $\langle (\text{AH}, \text{AL}, \text{IP}); (\text{E}_{\text{RegID}}, \text{IDcred}_{\text{PUB}}, \text{AL}) \rangle$.

Malicious AH, honest IP. When a corrupt AH invokes $\mathcal{F}_{\text{issue}}^{\mathcal{R}, \text{TE}, \text{SIG}}$ on command $(\text{ISSUE}, (\text{E}_{\text{RegID}}, \text{K}, r, \text{PK}_{\text{AR}}), \text{AL}, (\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}}), \text{IP})$, Sim aborts if E_{RegID} is not an encryption of K or if $(\text{IDcred}_{\text{SEC}}, \text{IDcred}_{\text{PUB}})$ is not a valid key-pair; Sim outputs **fail** if $\text{E}_{\text{RegID}} \in \text{list}_{\text{h-ct}}$ or if $\text{IDcred}_{\text{PUB}} \in \text{list}_{\text{h-pk}}$.

Otherwise, Sim calls command **ISSUE** of the functionality on input (IP, AL) and returns $\text{cid} = (\text{IP}, \text{AL}, \text{IDcred}_{\text{SEC}}, \sigma, \text{K})$ to AH where σ is a signature computed by Sim using **Sign** of **SIG** (since the simulator is internally emulating the honest IP w.r.t. the corrupt account holder AH). Sim adds $(\sigma; (\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}); \text{IP})$ to $\text{list}_{\text{h-sig}}$.

- Command **CreateACC**. *Malicious AH*. When a corrupt AH invokes $\mathcal{F}_{\text{ledger}}$ on input $\text{ACI} = (st, \pi)$ (where $st = (\text{P}, \text{E}_{\text{ID}}, \text{RegID}_{\text{ACC}}, \text{IP}, \text{pk}_{\text{ACC}})$), the simulator Sim uses $\mathcal{F}_{\text{nizk}}$ to extract the witness $w = (\sigma, x, r', \text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}, \text{sk}_{\text{ACC}}, \text{IDcred}_{\text{PUB}})$ (or abort if the proof doesn't verify). The simulator outputs **fail** if one of the following condition holds: $(\sigma; (\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}); \text{IP}) \notin \text{list}_{\text{h-sig}}$, if $\text{IDcred}_{\text{PUB}} \in \text{list}_{\text{h-pk}}$, if $\text{E}_{\text{ID}} \in \text{list}_{\text{h-ct}}$, if $\text{RegID}_{\text{ACC}} \in \text{list}_{\text{h-aid}}$, or if $x > \text{Max}_{\text{ACC}}$.

Otherwise the simulator inputs the **CreateACC**($\text{cid}, \text{P}, (\text{sk}_{\text{ACC}}, \text{pk}_{\text{ACC}})$) command to the ideal functionality and, when asked, inputs $\text{aid} = \text{RegID}_{\text{ACC}}$ to the ideal functionality.

Honest AH. For an honest account holder, the simulator upon receiving $(\text{CreateACC}, \text{P}, \text{pk}_{\text{ACC}}, \text{IP})$ from the ideal functionality, picks a random aid in the domain of the PRF and forwards it to the functionality (also adds it into $\text{list}_{\text{h-aid}}$). Then, the simulator prepares $st = (\text{P}, \text{E}_{\text{ID}}, \text{RegID}_{\text{ACC}} = \text{aid}, \text{IP}, \text{pk}_{\text{ACC}})$ where E_{ID} is an encryption of a dummy value (and is added to $\text{list}_{\text{h-ct}}$), simulates a proof π via $\mathcal{F}_{\text{nizk}}$ and appends (st, π) to the buffer of the ledger.

Add entry $(\text{RegID}_{\text{ACC}}, \text{E}_{\text{ID}}, \text{IP}, \text{pk}_{\text{ACC}}, \text{P}, \pi)$ in list_{acc} .

- Command **RELEASE**: the simulator simulates these commands directly simulating the calls to $\mathcal{F}_{\text{ledger}}$ e.g., when the adversary invokes $(\text{RELEASE}, \Pi)$ adds the permuted buffer to the list $\text{list}_{\text{ledger}}$ and then resets the buffer.
- Command **RETRIEVE** Sim emulates the retrieve command in $\mathcal{F}_{\text{ledger}}$ and gives as output $\text{list}_{\text{ledger}}$.
- Command **REVOKE**. *Semi-honest IP and up to d AR, honest AH*. When the IP and a qualified set of AR (of which up to d are corrupt) invoke **REVOKE**, the simulator Sim obtains $\text{RegID}_{\text{ACC}} =$

aid and AH from the input/output of the functionality $\mathcal{F}_{\text{id-layer}}$. Now Sim, using $\text{RegID}_{\text{ACC}}$, searches list_{acc} and retrieves the corresponding E_{ID} . Similarly, using (AH, IP), searches $\text{list}_{\text{issue}}$ and retrieves the corresponding $\text{IDcred}_{\text{PUB}}$. Then the simulator Sim equivocates the decryption of E_{ID} to $\text{IDcred}_{\text{PUB}}$ using SimShare (defined in the simulatability property of the threshold encryption scheme).

Malicious AH and up to d AR. The simulator receives $(\text{RegID}_{\text{ACC}}, \text{AH})$ from the ideal functionality, looks up the ciphertext E_{ID} corresponding to $\text{RegID}_{\text{ACC}}$ and runs the threshold decryption protocol as honest parties would do.

- **Command TRACE.** *Semi-honest IP and up to d AR, honest AH.* When the IP and a qualified set of AR (of which up to d are corrupt) invoke **TRACE**, the simulator Sim receives AH and a list $\text{list}_{\text{RegID}_{\text{ACC}}}$ of aids from the input/output of the functionality $\mathcal{F}_{\text{id-layer}}$. Now Sim recovers the ciphertext E_{RegID} . Finally Sim programs the output of $\mathcal{F}_{\text{mpc-prf}}$ to be consistent with $\text{list}_{\text{RegID}_{\text{ACC}}}$.

Malicious AH and up to d AR. The simulator receives AH and a list of accounts $\{\text{RegID}_{\text{ACC}}\}$ from the ideal functionality. The simulator looks up the ciphertext E_{RegID} corresponding to AH and emulates $\mathcal{F}_{\text{mpc-prf}}$ to output the list $\{\text{RegID}_{\text{ACC}}\}$.

We now argue that the view of the environment in the real world and in the ideal world with the simulator described above are indistinguishable.

Analysis of the Simulated Game. We first argue that the probability of the simulator outputting **fail** is negligible. In the **ISSUE** command, since the emulation of the $\mathcal{F}_{\text{issue}}$ functionality is unconditionally secure, the simulator outputs **fail** only in the following cases: 1. the adversary submits the opening for a ciphertext E_{RegID} generated by the simulator (whose randomness is never used anywhere else, and for which the simulator only uses less than d shares of the secret key). An adversary that makes the simulation fail this way can be turned into an attack on the semantic security of the threshold encryption scheme; 2. the adversary submits the secret key for a public key $\text{IDcred}_{\text{PUB}}$ generated by the simulator (whose secret key is never used by the simulator). An adversary that makes the simulation fail this way can be turned into an attack onto the one-wayness of the key generation algorithm (which contradicts the assumption that \mathcal{R} is a hard-relation). In the **CreateACC** command, since the emulation of the $\mathcal{F}_{\text{nizk}}$ functionality is unconditionally secure, and since the adversary cannot make the ledger accept replayed accounts, the simulator outputs **fail** only in the following cases: 3. the adversary outputs a message/signature pair $(\sigma; (\text{IDcred}_{\text{SEC}}, \text{K}, \text{AL}); \text{IP})$ which passes the verification algorithm for the public key of IP and was not generated by the simulator. An adversary that make the simulation fail this way can be turned into an attack on the unforgeability of the signature scheme; 4. the adversary inputs to the $\mathcal{F}_{\text{nizk}}$ functionality the secret key for a public key $\text{IDcred}_{\text{PUB}}$ generated by the simulator (see bullet 2. above); 5. the adversary submits the opening for a ciphertext E_{ID} generated by the simulator (see bullet 1. above); 6. the adversary inputs to the $\mathcal{F}_{\text{nizk}}$ functionality an account id $\text{RegID}_{\text{ACC}}$, together with the key K and input x such that $\text{RegID}_{\text{ACC}} = \text{PRF}(\text{K}, x)$ and that same $\text{RegID}_{\text{ACC}}$ also had been chosen at random by the simulator when simulating an honest party. An adversary that can make the simulation fail this way can be turned into an attack on the weak robust property of the PRF. 7. the adversary creates an account with $x \geq \text{Max}_{\text{ACC}}$. Since the $\mathcal{F}_{\text{nizk}}$ is emulated with unconditional security, the probability that the simulator outputs **fail** in this way is 0.

We now slowly turn the simulated game into the real protocol via a series of hybrids.

Hybrid 0. This is the simulated protocol with the exception that we never output `fail`. As argued above, the probability that the simulator outputs `fail` is negligible, therefore the distribution generated by the simulated game and this hybrid is computationally close.

Hybrid 1. In the first hybrid we change the way in which the simulator samples the account id `aid` for for honest `AH`. Now, instead of picking random values, the simulator picks keys K for each pair of honest account holders and certificate id (AH, cid) and computes the `aid` as $PRF_K(x)$ where x counts how many accounts were opened by that account holder for that `cid`. An adversary that can distinguish between this hybrid and the simulated protocol can be used to break the pseudorandomness property of `PRF`.

Hybrid 2. In the second hybrid we change the way in which the simulator generates the ciphertexts E_{ID} for the honest parties during the simulation of `CreateACC`. Now the simulator encrypts the value `IDcredPUB` corresponding to the certificate `cid` of the account holder `AH`. An adversary that can distinguish between this hybrid and the previous one can be used to break the semantic security of the threshold encryption scheme.

Hybrid 3. In the third hybrid we change the way in which the simulator generates the ciphertexts E_{RegID} for the honest parties during the simulation of `ISSUE`. Now the simulator encrypts the value K corresponding to the certificate `cid` of the account holder `AH`, consistently with the value of K which had been sampled for that `cid` in hybrid 1. An adversary that can distinguish between this hybrid and the previous one can be used to break the semantic security of the threshold encryption scheme.

Hybrid 4. In the fourth hybrid we change the behaviour of the simulator during `REVOKE` for honest `AH`. Now we let the simulator run the threshold decryption protocol honestly. Note that, since we changed the encryptions E_{RegID} in a previous hybrid to encrypt the right value, the output of the decryption in this and the previous hybrid is identical. Hence, an adversary that can distinguish between this and the previous hybrid can be used to break the partial decryption simulatability (PDS) of the threshold encryption scheme.

Hybrid 5. In this hybrid we change the behaviour of the simulator during `TRACE` for honest `AH`. Now we let the simulator run the code of $\mathcal{F}_{\text{mpc-prf}}$ honestly instead of programming the output of the functionality. Note that, since we changed the encryptions E_{ID} in a previous hybrid to encrypt the right value, the output of the decryption in this and the previous hybrid is identical. Since the simulation of $\mathcal{F}_{\text{mpc-prf}}$ is perfect, this hybrid is identically distributed to the previous one.

Hybrid 6. In the final hybrid we change the behaviour of the simulator when simulating the proofs π for honest account holders `AH`. Up until now the simulator had simply picked random strings π , stored them together with the statement x , and then emulated the answers to `VERIFY` queries of the adversary for such proofs by outputting 1 on behalf of the $\mathcal{F}_{\text{nizk}}$ functionality. Now instead the simulator runs the code of the functionality $\mathcal{F}_{\text{nizk}}$, that is it only outputs that a proof is valid if it can come up with a witness for it. Note that, since we changed all elements in the statement (the encryption, the account id, etc.) in previous hybrids to match the behaviour of a

honest account holder, the simulator knows the necessary witnesses, thus the previously simulated proofs still verify, and this hybrid is identically distributed to the previous one.

Since hybrid 6 is identical to the real protocol (that is, the real protocol run in the hybrid world where all the helping functionalities are available), this concludes the proof.

5.3 Credential Issue Protocol

The issue protocol Π_{issue} uses as its main ingredient a two-round blind signature scheme (as defined in Section 2.3), augmented with a NIZK that proves that the input to the blind-signature protocol is consistent with the ciphertext and the public-key that the account holder sends to the identity provider.

Protocol Issue Π_{issue}

The protocol operates in the $\{\mathcal{F}_{\text{crs}}, \mathcal{F}_{\text{reg}}, \mathcal{F}_{\text{nizk}}, \mathcal{F}_{\text{smt}}\}$ -hybrid model. Let $\text{BS} = (\text{Setup}, \text{KeyGen}, \text{Sign}_1, \text{Sign}_2, \text{Unblind}, \text{VerifySig})$ be a blind signature scheme and $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{TDec})$ be a (n, d) -threshold encryption scheme, and \mathcal{R} an NP relation corresponding to the account holders' key pair.

- Upon input (SETUP), use $\mathcal{F}_{\text{crs}}^{\text{Setup}}$ to generate $\mathbf{p} \leftarrow_{\$} \text{Setup}(1^\lambda)$ and publicize it to all parties.
- Upon input (INITIALIZE, $(\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}})$), the identity provider IP checks if the key pair has a correct distribution with respect to the $\text{KeyGen}(\mathbf{p})$. If yes, stores $(\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}})$ and sends (REGISTER, $\text{sk}_{\text{IP}}, \text{pk}_{\text{IP}}$) to \mathcal{F}_{reg} .
- Upon an input (ISSUE, $(ct, m, r, \text{pk}_{\text{AR}}), \text{aux}, (\text{sk}_{\text{AH}}, \text{pk}_{\text{AH}}), \text{IP}$) to the account holder AH and an input (ISSUE, AH, pk_{AR}) to an identity provider IP, proceed as follows:
 1. AH retrieves pk_{IP} from \mathcal{F}_{reg} , computes $\text{sign}_1 = \text{Sign}_1(\text{pk}_{\text{IP}}, (\text{sk}_{\text{AH}}, m, \text{aux}), \text{pp}; r')$ and sends (PROVE, st, w) to $\mathcal{F}_{\text{nizk}}$ for statement $st = (\text{pk}_{\text{AH}}, \text{sign}_1, ct, \text{aux}, \text{pp})$ using secret witness $w = (\text{sk}_{\text{AH}}, m, r', r)$ for the relation $\mathcal{R}_1(st, w)$ that outputs \top if:
 - (a) $(\text{sk}_{\text{AH}}, \text{pk}_{\text{AH}}) \in \mathcal{R}$.
 - (b) $\text{sign}_1 = \text{Sign}_1(\text{pk}_{\text{IP}}, (\text{sk}_{\text{AH}}, m, \text{aux}), \text{pp}; r')$.
 - (c) $ct = \text{TEnc}_{\text{PK}_{\text{AR}}}^{n,d}(m; r)$
 2. Upon receiving the proof π , AH sends (SEND, IP, (st, π)) to \mathcal{F}_{smt} .
 3. Upon receiving (SENT, AH, (st, π)) from \mathcal{F}_{smt} , IP inputs (VERIFY, st, π) (for relation \mathcal{R}_1) to $\mathcal{F}_{\text{nizk}}$. If they pass, IP computes and sends (through \mathcal{F}_{smt}) $\text{sign}_2 \leftarrow \text{Sign}_2(\text{sk}_{\text{IP}}, \text{sign}_1)$.
AH runs $\text{Unblind}(\text{sign}_2, r')$ and obtains a signature σ on $(\text{sk}_{\text{AH}}, m, \text{aux})$.

5.4 Proof of Security for Issue Protocol

Theorem 5.2. *Assume that $\text{BS} = (\text{Setup}, \text{KeyGen}, \text{Sign}_1, \text{Sign}_2, \text{Unblind}, \text{VerifySig})$ is a blind signature scheme. Then, Π_{issue} securely implements $\mathcal{F}_{\text{issue}}$ in the $\{\mathcal{F}_{\text{crs}}, \mathcal{F}_{\text{reg}}, \mathcal{F}_{\text{nizk}}, \mathcal{F}_{\text{smt}}\}$ -hybrid model in the presence of an actively corrupted AH or a passively corrupted IP.*

Proof. We consider corruptions of AH and IP separately.

(Actively) Corrupt AH. The simulator for a corrupt AH receives (st, π) from the \mathcal{F}_{smt} functionality and extracts the witness w from $\mathcal{F}_{\text{nizk}}$ (or abort if the proof doesn't verify). Now the simulator can input the **ISSUE** command to the ideal functionality which includes the message and randomness used to generate ct , the auxiliary information, and the account public and secret keys. The simulator therefore receives a signature σ from the ideal functionality. Using the simulator guaranteed from the simulatability property (Definition A.3), the simulator computes the appropriate sign_2 message. Note that we can use this simulator since (due to the soundness of the NIZK) we know that sign_1 was computed according to the protocol and our simulator knows the message m and the randomness r' which is used by the adversary to do so.

We claim that this simulation is indistinguishable from the real protocol. To see why, notice that the output of the IP is identical in the real and ideal world (remember that in the $\mathcal{F}_{\text{nizk}}$ hybrid model the proofs are unconditionally sound). Moreover, the view of the corrupted AH consists only of the message sign_2 . Therefore, an adversary that can distinguish between the real and simulated protocol can directly be turned into an adversary that breaks the simulatability property (Definition A.3) of the blind signature scheme.

(Passively) Corrupt IP. The case of a corrupted IP is relatively simple since we only consider passively corrupt IP: Here the simulator gets the input and output of the IP from the ideal functionality (which consists of the account public key, the ciphertext, and the auxiliary information), and has to produce a view indistinguishable from the protocol which is consistent with this output. The simulator does so by computing the first message of the blind signature sign_1 by running Sign_1 with some dummy input. Note that all other parts of the statement st are at this point known to the simulator, which therefore only needs to simulate the proof π , but this is trivial to do in the $\mathcal{F}_{\text{nizk}}$ -hybrid model by simply appending an arbitrary proof π^* to the view, and then answering verify queries so that π^* is a valid proof for the statement st .

We claim that this simulation is indistinguishable from the real protocol. To see why, notice that the only difference in the view of IP between the real and simulated protocol is in the distribution of sign_1 and therefore an adversary that can distinguish between the real and simulated protocol can directly be turned into an adversary that breaks the blindness property (Definition A.2) of the blind signature scheme.

Finally, note that at first glance it might appear strange that the proof does not mention at all the security of the threshold encryption scheme. However, this is because our ideal functionality just guarantees that the ciphertext is correctly computed according to the encryption algorithm. Thus, our protocol securely implements the functionality regardless of which security (if any) is guaranteed by the encryption scheme!

□

6 Putting Everything Together

We presented all components of the system in a modular way. We now describe how to instantiate each of the components needed in the ID-layer.

UC-NIZK. We use two different types of non-interactive zero knowledge proofs in our implementation. One is based on Σ -protocols made non-interactive with the Fiat-Shamir (FS) transform [FS87], and the other is preprocessing-based zkSNARKs [GGPR13, PHGR13] in the crs model. Unfortunately, known instantiations of both types of NIZKs do not satisfy UC-security.

In order to lift SNARK to be UC-secure we use the transformation of Kosba et al. [KZM⁺15], which is recalled in Appendix F.2. At a high level, the transformation works having the prover prove an augmented relation $\mathcal{R}_{\mathcal{L}'}$ that is given in Appendix F.2. A pair of one-time signing/verification keys are generated for each proof. The prover is additionally required to show that a ciphertext encrypts the witness of the underlying relation $\mathcal{R}_{\mathcal{L}}$, or the PRF was correctly evaluated on the signature key under a committed key. Then the prover is required to sign the statement together with the proof of \mathcal{L}' . Since our goal is to use SNARKs on small circuits for the purposes of prover efficiency, we treat the augmented relation $\mathcal{R}_{\mathcal{L}'}$ as a composite statement [AGM18, CFQ19] and use a combination of SNARKs and sigma protocols to prove the augmented relation of the transformation. We use the CL scheme [CL15] for encryption, and $f_k : x \rightarrow H(x)^k$, for $k \in \mathbb{Z}_q$ as the PRF where H maps bit strings to group elements. This PRF can be shown to be secure under the DDH assumption where H is modeled as a random oracle. We can use a sigma protocol to prove correct evaluation of the PRF given public input, public output, and committed key g^k . A standard sigma protocol proof of equality of discrete logarithms can be used to prove equality of CL encrypted and Pedersen committed messages. The composition theorem from [AGM18] can be invoked to argue security of the NIZK for the composite statement formed as the AND of the statements of the lifting transformation.

In Appendix F.1, we show that a simulation-sound NIZK (such as Fiat-Shamir as shown in [FKMV12]) and a perfectly correct CPA-secure encryption scheme are sufficient to instantiate a simulation-extractable NIZK by transforming the relation to include a ciphertext encrypting the witness.⁴

While this lifting technique for transforming a (sound) NIZK to a knowledge-sound NIZK is folklore, the use of CL encryption scheme [CL15] for this goal is novel up to our knowledge. Our choice of the CL encryption scheme in the transformation means that we can at the same time encrypt messages in the same plaintext space as the commitment schemes (thus allowing for efficient proofs of equality of discrete logarithms), and guarantee efficient decryption by the extractor. This is as opposed to using e.g., Paillier (where we could have efficient decryption but would need range proofs to prove equality of exponents in different groups) or ElGamal “in the exponent” (where the group order could be the same but efficient decryption can only be achieved by encrypting the witness in short chunks).

Implementation of Π_{issue} . We instantiate the blind signature scheme BS by the Pointcheval-Sanders (PS) signature scheme [PS16]. We recall the PS scheme in Appendix A.3, and prove that the PS signature satisfies the definitions of Blindness A.2 (here we prove a stronger variant than what given in the original paper) and Simulatability A.3 (which we define, as it is needed for proving UC security of the overall construction). We adopt the threshold encryption TE scheme (Definition A.5) described in [DHMR07] which follows the share and encrypt paradigm. We use the CL encryption scheme [CL15] to encrypt. Once again, our choice of CL encryption scheme means that we can at the same time encrypt messages in the same plaintext space as the commitment schemes (for efficient equality proofs) and guarantee efficient decryption when needed in the TRACE command by the ARs.

We now describe the Σ -protocols we use to prove relation \mathcal{R}_1 in Π_{issue} . We let \mathcal{R} be the discrete log relation where $\text{pk} = g^{\text{sk}}$. Then, we can prove that public keys and secret keys satisfy \mathcal{R} using the standard Σ -protocol **dlog** from Appendix B.3. The message output by Sign_1 in the PS blind signature is essentially a Pedersen commitment for vectors (see Appendix A.3). So we prove that

⁴The notions of black-box simulation extractable NIZK and UC-secure NIZK are interchangeable [Gro06, CLOS02].

Sign_1 was executed correctly using the AggregateDL protocol described in Appendix B.6 (note that due to the homomorphic nature of Pedersen commitment we don’t need to prove that the values in the AL which are leaked to the IP are correct, since both parties can add those to the commitment “in public”). Finally, we use the protocol in Appendix C.1 for proving that the ciphertext encrypts the right value, and use standard “AND” composition of Σ -protocols to assert that the values appearing in different proofs are consistent.

Implementation of $\Pi_{\text{id-layer}}$. We instantiate the weakly robust PRF scheme (Definition 2.2) with Dodis-Yampolskiy PRF [DY05]. In this case we use ElGamal as the base encryption scheme for the “share-and-encrypt” ad-hoc threshold encryption scheme TE (Definition 2.4). This is because we are encrypting the public key as a group element, which can also be seen as an encryption of the secret key for “ElGamal in the exponent”. Note that this allows to both easily prove knowledge of the secret key, and to make sure that the ARs will only learn the AH public key when decrypting. Due to the algebraic nature of the DY PRF, we can efficiently evaluate it inside an MPC protocol as required to implement $\mathcal{F}_{\text{mpc-prf}}$ using techniques described in [ST19, DOK⁺20].

When creating a new account, we use both SNARKs and Sigma protocols for proving a single composite statement consisting of a circuit-part and an algebraic part using the technique of [AGM18] to obtain SNARK on algebraically committed input. This commitment is used to tie the witness of the Sigma protocol to the witness used in the SNARK. We describe briefly how we use a combination of Σ -protocols and SNARKs in order to prove relation \mathcal{R} in $\Pi_{\text{id-layer}}$. We use the protocol in Appendix C.2 for proving knowledge of a signature. We use SNARKs on committed input to prove that account holder’s attribute list satisfies a certain policy, and for proving that $x \leq \text{Max}_{\text{ACC}}$ for committed x and public Max_{ACC} . The protocols for proving that $\text{RegID}_{\text{ACC}} = \text{PRF}_{\mathcal{K}}(x)$ and $\text{E}_{\text{ID}} = \text{TEnc}_{\text{PK}_{\text{AR}}}^{\text{n,d}}(\text{IDcred}_{\text{PUB}})$ are described in Appendix C.2. Note that all the Σ -protocol proofs are made non-interactive using Fiat-Shamir as discussed in Appendix B.2.

Implementation of a transaction layer. In Appendix E, we include a high-level description of a method for transferring money on the ledger in the accounts-based model.

Acknowledgements. The authors would like to thank all members of the Concordium Blockchain Research Center and the Concordium AG for useful feedback, and in particular: Matthias Hall-Andersen, Jesper Buus Nielsen, Torben Pedersen, Daniel Tschudi.

References

- [ACC⁺20] Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhayaoui, and Björn Tackmann. Privacy-preserving auditable token payments in a permissioned blockchain system. In Proceedings of the 2nd ACM Conference on Advances in Financial Technologies, pages 255–267, 2020.
- [AGM18] Shashank Agrawal, Chaya Ganesh, and Payman Mohassel. Non-interactive zero-knowledge proofs for composite statements. In Hovav Shacham and Alexandra Boldyreva, editors, CRYPTO 2018, Part III, volume 10993 of LNCS, pages 643–673. Springer, Heidelberg, August 2018.
- [ASM06] Man Ho Au, Willy Susilo, and Yi Mu. Constant-size dynamic k-TAA. In Roberto De Prisco and Moti Yung, editors, SCN 06, volume 4116 of LNCS, pages 111–125. Springer, Heidelberg, September 2006.

- [ASMC12] Man Ho Au, Willy Susilo, Yi Mu, and Sherman SM Chow. Constant-size dynamic k-times anonymous authentication. IEEE Systems Journal, 7(2):249–261, 2012.
- [Bag19] Karim Baghery. On the efficiency of privacy-preserving smart contract systems. In Johannes Buchmann, Abderrahmane Nitaj, and Tajje eddine Rachidi, editors, AFRICACRYPT 19, volume 11627 of LNCS, pages 118–136. Springer, Heidelberg, July 2019.
- [BCG⁺14] 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 Computer Society Press, May 2014.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In 42nd FOCS, pages 136–145. IEEE Computer Society Press, October 2001.
- [CCL⁺20] Guilhem Castagnos, Dario Catalano, Fabien Laguillaumie, Federico Savasta, and Ida Tucker. Bandwidth-efficient threshold EC-DSA. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, PKC 2020, Part II, volume 12111 of LNCS, pages 266–296. Springer, Heidelberg, May 2020.
- [CFQ19] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, ACM CCS 2019, pages 2075–2092. ACM Press, November 2019.
- [Cha83] David Chaum. Blind signature system. In David Chaum, editor, CRYPTO’83, page 153. Plenum Press, New York, USA, 1983.
- [CHK⁺06] Jan Camenisch, Susan Hohenberger, Markulf Kohlweiss, Anna Lysyanskaya, and Mira Meyerovich. How to win the clonewars: Efficient periodic n-times anonymous authentication. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, ACM CCS 2006, pages 201–210. ACM Press, October / November 2006.
- [CHL05] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya. Compact e-cash. In Ronald Cramer, editor, EUROCRYPT 2005, volume 3494 of LNCS, pages 302–321. Springer, Heidelberg, May 2005.
- [Cho09] Sherman S. M. Chow. Real traceable signatures. In Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini, editors, SAC 2009, volume 5867 of LNCS, pages 92–107. Springer, Heidelberg, August 2009.
- [CL15] Guilhem Castagnos and Fabien Laguillaumie. Linearly homomorphic encryption from DDH. In Kaisa Nyberg, editor, CT-RSA 2015, volume 9048 of LNCS, pages 487–505. Springer, Heidelberg, April 2015.
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In 34th ACM STOC, pages 494–503. ACM Press, May 2002.

- [CLT18] Guilhem Castagnos, Fabien Laguillaumie, and Ida Tucker. Practical fully secure unrestricted inner product functional encryption modulo p . In Thomas Peyrin and Steven Galbraith, editors, ASIACRYPT 2018, Part II, volume 11273 of LNCS, pages 733–764. Springer, Heidelberg, December 2018.
- [CMS96] Jan Camenisch, Ueli M. Maurer, and Markus Stadler. Digital payment systems with passive anonymity-revoking trustees. In Elisa Bertino, Helmut Kurth, Giancarlo Martella, and Emilio Montolivo, editors, ESORICS’96, volume 1146 of LNCS, pages 33–43. Springer, Heidelberg, September 1996.
- [CS97] Jan Camenisch and Markus Stadler. Efficient group signature schemes for large groups (extended abstract). In Burton S. Kaliski Jr., editor, CRYPTO’97, volume 1294 of LNCS, pages 410–424. Springer, Heidelberg, August 1997.
- [CZJ⁺17] Ethan Cecchetti, Fan Zhang, Yan Ji, Ahmed E. Kosba, Ari Juels, and Elaine Shi. Solidus: Confidential distributed ledger transactions via PVORM. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, ACM CCS 2017, pages 701–717. ACM Press, October / November 2017.
- [DDP06] Ivan Damgård, Kasper Dupont, and Michael Østergaard Pedersen. Unclonable group identification. In Serge Vaudenay, editor, EUROCRYPT 2006, volume 4004 of LNCS, pages 555–572. Springer, Heidelberg, May / June 2006.
- [DHMR07] Vanesa Daza, Javier Herranz, Paz Morillo, and Carla Ràfols. CCA2-secure threshold broadcast encryption with shorter ciphertexts. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, ProvSec 2007, volume 4784 of LNCS, pages 35–50. Springer, Heidelberg, November 2007.
- [DOK⁺20] Anders P. K. Dalskov, Claudio Orlandi, Marcel Keller, Kris Shrishak, and Haya Shulman. Securing DNSSEC keys via threshold ECDSA from generic MPC. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, ESORICS 2020, Part II, volume 12309 of LNCS, pages 654–673. Springer, Heidelberg, September 2020.
- [DP92] Alfredo De Santis and Giuseppe Persiano. Zero-knowledge proofs of knowledge without interaction (extended abstract). In 33rd FOCS, pages 427–436. IEEE Computer Society Press, October 1992.
- [DY05] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In Serge Vaudenay, editor, PKC 2005, volume 3386 of LNCS, pages 416–431. Springer, Heidelberg, January 2005.
- [FKMV12] Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the Fiat-Shamir transform. In Steven D. Galbraith and Mridul Nandi, editors, INDOCRYPT 2012, volume 7668 of LNCS, pages 60–79. Springer, Heidelberg, December 2012.
- [FMMO19] Prastudy Fauzi, Sarah Meiklejohn, Rebekah Mercer, and Claudio Orlandi. Quisquis: A new design for anonymous cryptocurrencies. In Steven D. Galbraith and Shiho Moriai, editors, ASIACRYPT 2019, Part I, volume 11921 of LNCS, pages 649–678. Springer, Heidelberg, December 2019.

- [FOR17] Pooya Farshim, Claudio Orlandi, and Răzvan Roşie. Security of symmetric primitives under incorrect usage of keys. *IACR Trans. Symm. Cryptol.*, 2017(1):449–473, 2017.
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
- [GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3):11:1–11:35, 2012.
- [Gro06] Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 444–459. Springer, Heidelberg, December 2006.
- [KTY04] Aggelos Kiayias, Yiannis Tsiounis, and Moti Yung. Traceable signatures. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 571–589. Springer, Heidelberg, May 2004.
- [KZM⁺15] Ahmed Kosba, Zhichao Zhao, Andrew Miller, Yi Qian, Hubert Chan, Charalampos Papamanthou, Rafael Pass, abhi shelat, and Elaine Shi. How to use SNARKs in universally composable protocols. Cryptology ePrint Archive, Report 2015/1093, 2015. <https://eprint.iacr.org/2015/1093>.
- [KZZ16] Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 705–734. Springer, Heidelberg, May 2016.
- [Mau09] Ueli M. Maurer. Unifying zero-knowledge proofs of knowledge. In Bart Preneel, editor, *AFRICACRYPT 09*, volume 5580 of *LNCS*, pages 272–286. Springer, Heidelberg, June 2009.
- [MPJ⁺13] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. A fistful of bitcoins: characterizing payments among men with no names. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 127–140, 2013.
- [NS05] Lan Nguyen and Reihaneh Safavi-Naini. Dynamic k-times anonymous authentication. In John Ioannidis, Angelos Keromytis, and Moti Yung, editors, *ACNS 05*, volume 3531 of *LNCS*, pages 318–333. Springer, Heidelberg, June 2005.

- [NVV18] Neha Narula, Willy Vasquez, and Madars Virza. zkLedger: Privacy-preserving auditing for distributed ledgers. Cryptology ePrint Archive, Report 2018/241, 2018. <https://eprint.iacr.org/2018/241>.
- [PBF⁺19] Andrew Poelstra, Adam Back, Mark Friedenbach, Gregory Maxwell, and Pieter Wuille. Confidential assets. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, FC 2018 Workshops, volume 10958 of LNCS, pages 43–63. Springer, Heidelberg, March 2019.
- [Ped92] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, CRYPTO’91, volume 576 of LNCS, pages 129–140. Springer, Heidelberg, August 1992.
- [PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In 2013 IEEE Symposium on Security and Privacy, pages 238–252. IEEE Computer Society Press, May 2013.
- [PS16] David Pointcheval and Olivier Sanders. Short randomizable signatures. In Kazue Sako, editor, CT-RSA 2016, volume 9610 of LNCS, pages 111–126. Springer, Heidelberg, February / March 2016.
- [RSY18] Leonid Reyzin, Adam Smith, and Sophia Yakubov. Turning HATE into LOVE: Homomorphic ad hoc threshold encryption for scalable MPC. Cryptology ePrint Archive, Report 2018/997, 2018. <https://eprint.iacr.org/2018/997>.
- [Sch90] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, CRYPTO’89, volume 435 of LNCS, pages 239–252. Springer, Heidelberg, August 1990.
- [ST19] Nigel P. Smart and Younes Talibi Alaoui. Distributing any elliptic curve based protocol. In Martin Albrecht, editor, 17th IMA International Conference on Cryptography and Coding, volume 11929 of LNCS, pages 342–366. Springer, Heidelberg, December 2019.
- [SU12] Dominique Schröder and Dominique Unruh. Security of blind signatures revisited. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, PKC 2012, volume 7293 of LNCS, pages 662–679. Springer, Heidelberg, May 2012.
- [TFS04] Isamu Teranishi, Jun Furukawa, and Kazue Sako. k-Times anonymous authentication (extended abstract). In Pil Joong Lee, editor, ASIACRYPT 2004, volume 3329 of LNCS, pages 308–322. Springer, Heidelberg, December 2004.
- [TS06] Isamu Teranishi and Kazue Sako. k-times anonymous authentication with a constant proving cost. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, PKC 2006, volume 3958 of LNCS, pages 525–542. Springer, Heidelberg, April 2006.
- [Zca] Zcash regulatory and compliance brief. <https://z.cash/wp-content/uploads/2020/07/Zcash-Regulatory-Brief-062020.pdf>. 2020-06-01.

A Additional Definitions

A.1 Bilinear groups

Throughout the paper we let \mathcal{G} be a bilinear group generator that on input security parameter λ returns $(q, \mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T, e, g, \tilde{g}) \leftarrow \mathcal{G}(1^\lambda)$ with the following properties:

- $\mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T$ are groups of prime order q .
- $e : \mathbb{G} \times \tilde{\mathbb{G}} \rightarrow \mathbb{G}_T$ is a non-degenerate efficiently computable bilinear map such that $\forall U \in \mathbb{G}, \forall V \in \tilde{\mathbb{G}}, \forall a, b \in \mathbb{Z} : e(U^a, V^b) = e(U, V)^{ab}$.
- g and \tilde{g} are generators for \mathbb{G} and $\tilde{\mathbb{G}}$ respectively.

A.2 Commitment Schemes

Definition A.1 (Commitment Schemes). *A commitment scheme* $\text{COM} = (\text{Setup}, \text{Commit}, \text{OpenVrf})$ consists of the following algorithms:

- $\text{Setup}(1^\lambda)$: given the security parameter 1^λ , it outputs a commitment key ck .
- $\text{Commit}_{\text{ck}}(m)$: a probabilistic algorithm that given a message m , outputs a pair $(cm, r_{cm}) \leftarrow \text{Commit}_{\text{ck}}(m)$ of commitment cm and an opening r_{cm} . We sometimes write $\text{Commit}_{\text{ck}}(m; r_{cm})$ when we want to be able to fix the value of the randomness r_{cm} to a specific value.
- $\text{OpenVrf}_{\text{ck}}(cm, m, r_{cm})$: a deterministic algorithm that outputs a bit indicating acceptance or rejection.

Hiding property. We say that COM is hiding if for all non-uniform PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that

$$\left| \frac{1}{2} - \Pr \left[b' = b \mid \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda); (m_0, m_1) \leftarrow \mathcal{A}(\text{ck}); \\ b \leftarrow \{0, 1\}; (cm, r_{cm}) \leftarrow \text{Commit}_{\text{ck}}(m_b); b' \leftarrow \mathcal{A}(cm); \end{array} \right] \right| \leq \text{negl}(\lambda)$$

where the probability is over the randomness of \mathcal{A} , Setup , Commit and the choice of bit b . We say that COM is perfectly hiding if $\text{negl}(\lambda) = 0$.

Binding property. A commitment scheme COM is binding if for all non-uniform PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that

$$\Pr \left[\begin{array}{l} \text{OpenVrf}_{\text{ck}}(cm, m_0, r_0) = 1 \wedge \\ \text{OpenVrf}_{\text{ck}}(cm, m_1, r_1) = 1 \wedge m_0 \neq m_1 \end{array} \mid \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda); \\ (cm, m_0, m_1, r_0, r_1) \leftarrow \mathcal{A}(\text{ck}); \end{array} \right] \leq \text{negl}(\lambda)$$

where the probability is over the randomness of \mathcal{A} and Setup . We say that COM is perfectly binding if $\text{negl}(\lambda) = 0$.

$\text{Exp}_A^{\text{Blind}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $\begin{aligned} & \mathbf{p} \leftarrow \text{Setup}(1^\lambda); (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\mathbf{p}); \\ & (m_0, m_1) \leftarrow \mathcal{A}(\text{sk}, \mathbf{p}); \\ & b \leftarrow \{0, 1\}; \\ & (\text{sign}_1, \omega) \leftarrow \text{Sign}_1(\text{pk}, m_b); \\ & (\overline{\text{sign}}_1, \overline{\omega}) \leftarrow \text{Sign}_1(\text{pk}, m_{1-b}); \\ & \text{sign}_2, \overline{\text{sign}}_2 \leftarrow \mathcal{A}(\text{sign}_1, \overline{\text{sign}}_1); \\ & \text{let } \sigma_b, \sigma_{1-b} \text{ denote the (possibly undefined) local outputs of} \\ & \text{Unblind}(\text{sign}_2, \omega) \text{ and } \text{Unblind}(\overline{\text{sign}}_2, \overline{\omega}) \text{ respectively.} \\ & \text{if } \text{VerifySig}(\text{pk}, m_0, \sigma_0) = 0 \text{ or } \text{VerifySig}(\text{pk}, m_1, \sigma_1) = 0, \text{ then } (\sigma_0, \sigma_1) = (\perp, \perp); \\ & b^* \leftarrow \mathcal{A}(\sigma_0, \sigma_1); \\ & \text{if } b = b^* \wedge m_0 = m_1 , \text{ then return 1; else return 0;} \end{aligned}$
--

Figure 1: Experiment in the definition of blindness property

A.3 (Blind) Signature Scheme

We define the properties we require from blind signature schemes in this work – Blindness and Simulatability. Blindness captures the requirement that we run two blind signing protocols on two messages of the adversary’s choice, and the adversary should not be able to say which input-signature pair corresponds to which execution. This is stronger than the definition in [PS16]; however, we can show that the PS blind signature scheme can be made to satisfy this stronger, standard definition. This follows from the rerandomizability of PS signatures.

Definition A.2 (Blindness). *An interactive signature scheme $\text{BS} = (\text{Setup}, \text{KeyGen}, \text{Sign}_1, \text{Sign}_2, \text{Unblind}, \text{VerifySig})$ is called blind if for any PPT adversary \mathcal{A} , $\Pr[\text{Exp}_A^{\text{Blind}}(\lambda) = 1] \approx_\lambda 1/2$, where the experiment $\text{Exp}_A^{\text{Blind}}(\lambda)$ is defined in Fig. 1.*

Definition A.3 (Simulatability). *An interactive signature scheme $\text{BS} = (\text{Setup}, \text{KeyGen}, \text{Sign}_1, \text{Sign}_2, \text{Unblind}, \text{VerifySig})$ is called simulatable if there exist a PPT algorithm Sim s.t. for any PPT adversary \mathcal{A} , $\Pr[\text{Exp}_A^{\text{Sim}}(\lambda) = 1] \approx_\lambda 1/2$, where the experiment $\text{Exp}_A^{\text{Sim}}(\lambda)$ is defined in Fig. 2.*

A.3.1 Instantiation with PS Signature scheme [PS16]

We specify here the multi-message Pointcheval-Sanders (PS) signature scheme and how it realizes our definition of Blind signature. The PS scheme is a randomizable signature that uses groups with asymmetric pairing and allows to sign messages that are completely unknown to the signer.

- $\text{Setup}(1^\lambda)$: Sample $\mathbf{p} := (q, \mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T, e, g, \tilde{g}) \leftarrow \mathcal{G}(1^\lambda)$.
- $\text{KeyGen}(\mathbf{p})$: Given $(q, \mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T, e, g, \tilde{g})$, the key generation works as follows:
 1. Choose $x, y_i \leftarrow \mathbb{F}_q$ for $i = 1, \dots, \ell^5$.
 2. Set $X = g^x, \tilde{X} = \tilde{g}^x, Y_i = g^{y_i}$ and $\tilde{Y}_i = \tilde{g}^{y_i}$ for $i = 1, \dots, \ell$.

⁵Note that the scheme allows to sign vectors over \mathbb{F}_q of length ℓ , so ℓ should be chosen with this in mind.

$\text{Exp}_{\mathcal{A}}^{\text{Sim}}(\lambda)$ <hr style="border: 0; border-top: 1px solid black; margin: 2px 0;"/> $\begin{aligned} & \mathbf{p} \leftarrow \text{Setup}(1^\lambda); (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\mathbf{p}); \\ & b \leftarrow \{0, 1\}; \\ & (m; r) \leftarrow \mathcal{A}(\text{sk}, \mathbf{p}); \\ & (\text{sign}_1, \omega) \leftarrow \text{Sign}_1(\text{pk}, m; r); \\ & \text{sign}_2^0 \leftarrow \text{Sign}_2(\text{sk}, \text{sign}_1); \\ & \sigma = \text{Unblind}(\text{sign}_2^0, \omega); \\ & \text{sign}_2^1 \leftarrow \text{Sim}(\text{pk}, m, \omega, \sigma); \\ & b^* \leftarrow \mathcal{A}(\sigma, m, \text{sign}_2^1); \\ & \text{if } b = b^*, \text{ then return 1; else return 0;} \end{aligned}$
--

Figure 2: Experiment in the definition of simulatability property

3. The public key is now $(\tilde{X}, \{Y_i\}_{i \in [\ell]}, \{\tilde{Y}_i\}_{i \in [\ell]})$, while the secret key is X .

- $\text{Sign}_1(\text{pk}, m = (m_1, \dots, m_\ell))$: The user chooses $\omega \leftarrow \mathbb{F}_q$, and computes $\text{sign}_1 = g^\omega \prod_{i=1}^{\ell} Y_i^{m_i}$. Note that ω is random and not used for anything else, and therefore sign_1 perfectly hides the rest of the m_i .
- $\text{Sign}_2(\text{sk}, \text{sign}_1)$: The signer chooses $\alpha \leftarrow \mathbb{F}_q$ and sets $a' = g^\alpha \neq 1_{\mathbb{G}}$ and $b' = (X \cdot \text{sign}_1)^\alpha$. They output $\text{sign}_2 = (a', b')$.
- $\text{Unblind}(\text{sign}_2, \omega)$: For $\text{sign}_2 = (a', b')$, the user first computes $\hat{\sigma} = (a', b'/a'^\omega)$. Then, they rerandomize $\hat{\sigma}$ by choosing $\gamma \leftarrow \mathbb{F}_q$ and computing $\sigma = \hat{\sigma}^\gamma = (\hat{\sigma}_1^\gamma, \hat{\sigma}_2^\gamma)$. The user returns σ .
- $\text{VerifySig}(\text{pk}, m = (m_1, \dots, m_\ell), \sigma)$: The algorithm parses σ as (σ_1, σ_2) and outputs 1 if the following checks pass:
 1. $\sigma_1 \neq 1_{\mathbb{G}}$.
 2. $e(\sigma_1, \tilde{X} \prod \tilde{Y}_j^{m_j}) = e(\sigma_2, \tilde{g})$.

PS satisfies Correctness, Unforgeability, Blindness and Simulatability. **Correctness.** If $\sigma = (\sigma_1 = h, \sigma_2 = h^{x + \sum y_j m_j})$, then

$$\begin{aligned} e(\sigma_1, \tilde{X} \prod \tilde{Y}_j^{m_j}) &= e(h, \tilde{X} \prod \tilde{Y}_j^{m_j}) = e(h, \tilde{g})^{x + \sum y_j m_j} \\ &= e(h^{x + \sum y_j m_j}, \tilde{g}) = e(\sigma_2, \tilde{g}). \end{aligned}$$

Unforgeability. As shown in [PS16], the EUF-CMA security of this signature is equivalent to the single-message version, where its security relies on the following assumption⁶.

Definition A.4 (PS Assumption 1 [PS16]). *Let $(q, \mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T, e, g, \tilde{g})$ be a bilinear group setting of type 3. For $(X = g^x, Y = g^y)$ and $(\tilde{X} = \tilde{g}^x, \tilde{Y} = \tilde{g}^y)$, where x and y are random scalars in \mathbb{Z}_q , we define the oracle $\mathcal{O}(m)$ on input $m \in \mathbb{Z}_q$ that chooses a random $h \in \mathbb{G}$ and outputs the*

⁶The assumption is related to the LRSW assumption and is proved in [PS16] to hold in the bilinear generic group model.

$\text{Exp}_{\mathcal{A}}^{\text{SSS}}(\lambda, n, d)$ <hr style="border: 0.5px solid black;"/> $[n] \supseteq \mathcal{C} \leftarrow \mathcal{A}(\lambda, n, d);$ $(\text{sk}_i, \text{pk}_i) \leftarrow_{\$} \text{TKeyGen}(1^\lambda) \text{ for } i \in [n];$ $([n] \supseteq R, m_0, m_1) \leftarrow \mathcal{A}(\{\text{pk}_i\}_{i \in [n]}, \{\text{sk}_i\}_{i \in \mathcal{C}});$ $b \leftarrow_{\$} \{0, 1\}; ct \leftarrow_{\$} \text{TEnc}_{\text{pk}_R}^{n, d}(m_b);$ $b' \leftarrow \mathcal{A}(ct);$ $\text{if } (b' = b \wedge m_0 = m_1 \wedge R \cap \mathcal{C} \leq d);$ $\text{then return 1; else return 0;}$
--

Figure 3: Static Semantic Security Experiment for Threshold Encryption

pair $P = (h, h^{x+my})$. Given $(g, Y, \tilde{g}, \tilde{X}, \tilde{Y})$ and unlimited access to this oracle, no adversary can efficiently generate such a pair, with $h \neq 1_G$, for a new scalar m^* , not asked to \mathcal{O} .

Blindness. Let us define a game H_0 which is defined as game $\text{Exp}_{\mathcal{A}}^{\text{Blind}}$ with the only difference that $b = 0$. We can also define a symmetrical experiment H_1 where the bit $b = 1$.

To claim indistinguishability between these two games we make the following observations: 1) sign_1 perfectly hides m ; 2) the message sign_2 coming from a corrupted signer is rerandomized (through the Unblind algorithm) before the final signature σ is given to adversary \mathcal{A} . From the above two observations follow that the execution in which \mathcal{A} participates first in the computation of σ_0 and then of σ_1 is identically distributed to the one where \mathcal{A} participates first in the computation of σ_1 and then of σ_0 . Therefore H_0 is indistinguishable from H_1 .

Simulatability. We start by defining the simulator Sim . Sim given m, ω and the signature $\sigma = (\sigma_1, \sigma_2)$ chooses $\beta \leftarrow_{\$} \mathbb{F}_q$ and returns $\text{sign}_2 = (a', b')$ where $a' = \sigma_1^\beta$ and $b' = \sigma_2^\beta \cdot \sigma_1^{\omega \cdot \beta}$.

We now argue that the game H_0 where sign_2 is computed honestly (i.e., as output of Sign_2) is identically distributed to a simulated game H_1 where sign_2 is computed by Sim . This indistinguishability follows from the fact that the distribution of the message sign_2 output by Sim and the distribution of the output of Sign_2 are identical. We conclude the proof observing that H_0 corresponds to the experiment $\text{Exp}_{\mathcal{A}}^{\text{Sim}}$ where the bit $b = 0$ and H_1 corresponds to $\text{Exp}_{\mathcal{A}}^{\text{Sim}}$ with $b = 1$. \square

A.4 Threshold Encryption Scheme

Definition A.5. A (n, d) -threshold encryption scheme $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{ShareDec}, \text{TCombine})$ is (n, d) -statically semantic secure (SSS) if for all PPT adversaries \mathcal{A} , $\Pr[\text{Exp}_{\mathcal{A}}^{\text{SSS}}(\lambda, n, d) = 1] \approx_{\lambda} 1/2$, where the experiment $\text{Exp}_{\mathcal{A}}^{\text{SSS}}(\lambda, n, d)$ is defined in Fig. 3.

Definition A.6. We say that a (n, d) -threshold encryption scheme $\text{TE} = (\text{TKeyGen}, \text{TEnc}, \text{ShareDec}, \text{TCombine})$ is (n, d) -partial decryption simulatable (PDS) if there exists an efficient algorithm SimPart such that for all PPT adversaries \mathcal{A} , $\Pr[\text{Exp}_{\mathcal{A}}^{\text{PDS}}(\lambda, n, d) = 1] \approx_{\lambda} 1/2$, where the experiment $\text{Exp}_{\mathcal{A}}^{\text{PDS}}(\lambda, n, d)$ is described in Fig. 4.

A.4.1 Construction based on Share and Encrypt paradigm.

Below, we recall the construction of [DHMR07]. Let $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a public key encryption scheme and $\text{SS} = (\text{Share}, \text{Reconstruct})$ be a secret sharing scheme.

$\text{Exp}_{\mathcal{A}}^{\text{PDS}}(\lambda, n, d)$ <hr style="border: 0.5px solid black;"/> $[n] \supseteq \mathcal{C} \leftarrow \mathcal{A}(\lambda, n, d);$ $(\text{sk}_i, \text{pk}_i) \leftarrow_{\$} \text{TKeyGen}(1^\lambda) \text{ for } i \in [n];$ $([n] \supseteq R, m_0, m_1) \leftarrow \mathcal{A}(\{\text{pk}_i\}_{i \in [n]}, \{\text{sk}_i\}_{i \in \mathcal{C}});$ $b \leftarrow_{\$} \{0, 1\}; ct_0 \leftarrow_{\$} \text{TEnc}_{\vec{\text{pk}}_R}^{n,d}(m_0); ct_1 \leftarrow_{\$} \text{TEnc}_{\vec{\text{pk}}_R}^{n,d}(m_1);$ $\text{if } b = 0 \text{ then } \mu_i = \text{ShareDec}_{\vec{\text{pk}}_R, \text{sk}_i}^{n,d}(ct_0) \text{ for } i \in R \setminus \mathcal{C};$ $\text{else if } b = 1 \text{ then } \left(\mu_i = \text{ShareDec}_{\vec{\text{pk}}_R, \text{sk}_i}^{n,d}(ct_1) \text{ for } i \in R \cap \mathcal{C} \right.$ $\left. \wedge \mu_i = \text{SimPart}(n, d, \vec{\text{pk}}_R, c_0, \{\mu_j\}_{j \in R \cap \mathcal{C}}, m_0) \text{ for } i \in R \setminus \mathcal{C} \right);$ $b' \leftarrow \mathcal{A}(ct_b, \{\mu_j\}_{j \in R \setminus \mathcal{C}});$ $\text{if } (b' = b \wedge m_0 = m_1 \wedge R \cap \mathcal{C} \leq d);$ $\text{then return 1; else return 0;}$

Figure 4: Partial Decryption Simulatability Experiment for Threshold Encryption

- $\text{TKeyGen}(1^\lambda)$: return $\text{KeyGen}(1^\lambda)$.
- $\text{TEnc}_{\vec{\text{pk}}_R}^{n,d}(m)$: return $ct = (ct_1, \dots, ct_n)$, where $ct_i = \text{Enc}_{\text{pk}_i}(m_i)$ for $i \in R$ and $(m_1, \dots, m_n) = \text{Share}^{n,d}(m)$.
- $\text{ShareDec}_{\vec{\text{pk}}_R, \text{sk}_i}^{n,d}(ct)$: let $ct = (ct_1, \dots, ct_n)$. Return decryption share $\mu_i = \text{Dec}_{\text{sk}_i}(ct_i)$.
- $\text{TCombine}_{\vec{\text{pk}}_R}^{n,d}(ct, \{\mu_i\}_{i \in I})$: Return $m = \text{Reconstruct}^{n,d}(\{\mu_i\}_{i \in I})$.

The following theorem is proved in [RSY18] (see Appendix D, Thm. 11 of [RSY18]).

Lemma A.7. *The threshold encryption scheme TE described above is (n, d) -statically secure, as long as SS is a secure share simulatable (n, d) -secret sharing scheme, and PKE is a CPA-secure public key encryption scheme.*

A.5 Secret Sharing Scheme

Informally, a (n, d) -secret sharing of a secret value s is an encoding of s into n shares, such that any $d + 1$ shares together can reconstruct s , whereas fewer shares reveal no information about s .

Definition A.8. *A (n, d) -secret sharing scheme $SS = (\text{Share}, \text{Reconstruct})$ over message space S consists of the following algorithms:*

- $\text{Share}^{n,d}(s; r) \rightarrow ([s]_1, \dots, [s]_n)$ is a randomized algorithm that on any input $s \in S$ and randomness r , outputs n shares $([s]_1, \dots, [s]_n)$.
- $\text{Reconstruct}^{n,d}([s]_{i_1}, \dots, [s]_{i_{d+1}}) \rightarrow s'$ is a deterministic algorithm that takes $d + 1$ shares as input and returns the reconstructed message $s' \in S$.

$\text{Exp}_{\mathcal{A}}^{\text{SS-Sim}}(\lambda, n, d)$ <hr style="border: 0.5px solid black;"/> $(\mathcal{C}, s_0, s_1) \leftarrow \mathcal{A}(n, d);$ <p>if $s_0 \notin S \vee s_1 \notin S \vee s_0 \neq s_1$ then return 0;</p> $([s_0]_1, \dots, [s_0]_n) \leftarrow \text{Share}^{n,d}(s_0);$ $([s_1]_1, \dots, [s_1]_n) \leftarrow \text{Share}^{n,d}(s_1);$ $\{[s'_0]_i\}_{i \in [n] \setminus \mathcal{C}} \leftarrow \text{SimShare}(n, d, \{[s_1]_i\}_{i \in \mathcal{C}}, s_0);$ $b \leftarrow \{0, 1\};$ <p>if $b = 0$ then $x = \{[s_0]_i\}_{i \in [n]}$;</p> <p>else $x = \{[s_1]_i\}_{i \in \mathcal{C}} \cup \{[s'_0]_i\}_{i \in [n] \setminus \mathcal{C}}$;</p> $b' \leftarrow \mathcal{A}(x);$ <p>if $b' = b \wedge \mathcal{C} \leq d$ then return 1;</p> <p>else return 0;</p>
--

Figure 5: Share Simulatability experiment for Secret Sharing

Informally, correctness requires that $s' = s$, and privacy requires that given d or fewer shares of either s_0 or s_1 , no efficient adversary can guess which message was shared.

Share Simulatability. We additionally use another property for a secret sharing as defined in [RSY18] which requires that given any set of d or fewer shares of s_0 and given s_1 , there exists a PPT algorithm SimShare that generates the rest of the shares in such a way that is indistinguishable from a fresh sharing of s_1 .

Definition A.9 (Share Simulatability). *We say that a (n, d) -secret sharing scheme $\text{SS} = (\text{Share}, \text{Reconstruct})$ over message space S is share simulatable if there exists an efficient simulation algorithm SimShare such that for all PPT adversaries \mathcal{A} , $\Pr[\text{Exp}_{\mathcal{A}}^{\text{SS-Sim}}(\lambda, n, d) = 1] \approx_{\lambda} 1/2$, where the experiment $\text{Exp}_{\mathcal{A}}^{\text{SS-Sim}}(\lambda, n, d)$ is described in Fig. 5.*

A.6 Non-Interactive Zero-Knowledge Proofs

Definition A.10 (NIZK). *A non-interactive zero-knowledge proof system (NIZK) for an NP language \mathcal{L} with relation $\mathcal{R}_{\mathcal{L}}$ consists of the following four algorithms:*

- $\text{CRSGen}(1^\lambda, \mathcal{L})$. *On input 1^λ and the description of the language \mathcal{L} , generates a common reference string crs , a trapdoor τ and an extraction key ek .*
- $\text{Prove}(\text{crs}, x, w)$. *On input of a crs , a statement x with witness w , outputs a proof π .*
- $\text{Verify}(\text{crs}, x, \pi)$. *Given a crs , a statement x and a proof π , outputs a bit indicating accept or reject.*
- $\text{SimProve}(\text{crs}, \tau, x)$. *On input a crs , a trapdoor τ and a statement x , outputs a simulated proof π without needing a witness for x .*

Completeness. A NIZK is (perfectly) complete, if an honest prover with a valid witness can always convince an honest verifier. More formally, for any $(x, w) \in \mathcal{R}_{\mathcal{L}}$, we have:

$$\Pr \left[\begin{array}{l} (\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}), \pi \leftarrow \text{Prove}(\text{crs}, x, w) : \\ \text{Verify}(\text{crs}, x, \pi) = 1 \end{array} \right] = 1$$

Soundness. A NIZK scheme for the language \mathcal{L} is called (computationally) sound, if for all PPT adversaries \mathcal{A} , we have:

$$\Pr \left[\begin{array}{l} (\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}), (x, \pi) \leftarrow \mathcal{A}(\text{crs}) : \\ \text{Verify}(\text{crs}, x, \pi) = 1 \wedge x \notin \mathcal{L} \end{array} \right] \approx_\lambda 0$$

Zero-knowledge. A NIZK scheme for the language \mathcal{L} is called (computationally) zero-knowledge if for all PPT adversary \mathcal{A} ,

$$\begin{aligned} & \Pr \left[(\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}) : \mathcal{A}^{\text{Prove}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 \right] \\ & \approx_\lambda \Pr \left[(\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}) : \mathcal{A}^{\text{SimProve}(\text{crs}, \tau, \cdot)}(\text{crs}) = 1 \right] \end{aligned}$$

Simulation Soundness. A NIZK proof for the language \mathcal{L} is called simulation sound if for all PPT adversaries \mathcal{A} ,

$$\Pr \left[\begin{array}{l} (\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}); (x, \pi) \leftarrow \mathcal{A}^{\text{SimProve}(\text{crs}, \tau, \cdot)}(\text{crs}) : \\ (x, \pi) \notin \mathcal{Q} \wedge x \notin \mathcal{L} \wedge \text{Verify}(\text{crs}, x, \pi) = 1 \end{array} \right] \approx_\lambda 0$$

where \mathcal{Q} is the list of simulation queries and responses (x_i, π_i) .

Simulation extractability. This is a strong notion which requires that for any adversary that outputs a proof after seeing many simulated proofs (for either true or false statements), there exists an efficient extractor that can extract the witness from the proof. More formally, we say that a NIZK scheme for the language \mathcal{L} is simulation extractable if there exists an efficient algorithm Ext (called extractor), such that for any PPT adversary \mathcal{A} , we have:

$$\Pr \left[\begin{array}{l} (\text{crs}, \tau, \text{ek}) \leftarrow \text{CRSGen}(1^\lambda, \mathcal{L}); (x, \pi) \leftarrow \mathcal{A}^{\text{SimProve}(\text{crs}, \tau, \cdot)}(\text{crs}, \text{ek}); \\ w \leftarrow \text{Ext}(\text{crs}, \text{ek}, x, \pi) : (x, \pi) \notin \mathcal{Q} \wedge (x, w) \notin \mathcal{R}_{\mathcal{L}} \wedge \text{Verify}(\text{crs}, x, \pi) = 1 \end{array} \right] \approx_\lambda 0$$

where \mathcal{Q} contains the list of statement-proof pairs that \mathcal{A} asks the oracle $\text{SimProve}(\text{crs}, \tau, \cdot)$. Note that this definition considers a strong version of simulation extractability, where the adversary should not even be able to generate a new proof (different from oracle's response) for a statement that has been queried before. We can relax this definition by considering a weaker version where the adversary may be able to generate a new proof for a queried statements. However, for the universally composable NIZKs that we use for realizing $\mathcal{F}_{\text{nizk}}$ (defined in section 4), we require the stronger definition. Thus, when we discuss about the notion of simulation extractability in section F, we mean the stronger version defined above.

B Implementation Details

This section contains a detailed description of the protocols to implement for the ID layer.

B.1 Auxiliary Σ -protocols.

Let $(q, \mathbb{G}, \tilde{\mathbb{G}}, \mathbb{G}_T, e, g, \tilde{g})$ be a bilinear group as defined in A.1. In the Σ protocols described in the following we assume that all group elements and commitment keys and public keys for signature and encryption schemes come from a group of order q , where in some cases it does not matter if we use \mathbb{G} , $\tilde{\mathbb{G}}$ or \mathbb{G}_T , but there are also cases where a particular group has to be used. The protocol applying a Σ -protocol will specify which group is used. We denote by $\text{PK}\{(x, y, \dots) : \text{statements about } x, y\}$, a zero knowledge proof of knowledge of x, y, \dots that satisfies statements. Here, x, y, \dots are private (witness), and other values in statements are public. The sigma protocols we describe below are standard protocols that have appeared in other works [Sch90, CS97] and follow from the *pre-image protocol* of Maurer [Mau09] for proving knowledge of a pre-image of a group homomorphism that unifies and generalizes a large class of protocols in literature.

B.2 Using Fiat-Shamir for non-interactive Proof

All Σ protocols below are described as interactive 2-party protocols. When using them for non-interactive proofs, we use the Fiat-Shamir paradigm with hash function to replace a random oracle. More precisely, this will be done as follows: Suppose the messages in the protocol are (a, c, z) , where c is the random challenge from the verifier, and where x is the public input. Then the prover executes the protocol on his own as follows: they compute a first, and then $\text{H}_{\text{RO}}(x, a)$. They set c to be the first $\lfloor \log_2 q \rfloor$ bits of the hash value. Finally they compute the last message z . The non-interactive proof is (c, z) .

To verify a proof we use the fact that all our Σ protocols have a verification equation that allows you to compute a from x, c and z . So given x and e, z , the verifier computes a , $\text{H}_{\text{RO}}(x, a)$ and compares c to the first $\lfloor \log_2 q \rfloor$ bits of the hash value.

If a proof consists of several, say T Σ protocols, we do them all in parallel as follows: the prover computes all T first messages and concatenates them to get M_1 . Then they compute $\text{H}_{\text{RO}}(X, M_1)$ where X is the concatenation of all public inputs to the protocols. They set c to be the first $\lfloor \log_2 q \rfloor$ bits of the hash value. Finally they compute the T last messages and concatenates them all to get M_2 . The non-interactive proof is (c, M_2) .

To verify X and (c, M_2) , where X contains all the public inputs, the verifier splits M_2 into T individual last messages. As above, they use the verification equations to compute the T first messages. They concatenate them to get M_1 , computes $\text{H}_{\text{RO}}(X, M_1)$ and sets c to the first $\lfloor \log_2 q \rfloor$ bits of the hash value.

B.3 Sigma protocol for proving knowledge of discrete log

Protocol $\text{dlog} : \text{PK}\{(x) : y = g^x\}$

- The prover computes and sends $a = g^\alpha$ for randomly chosen $\alpha \in \mathbb{Z}_q^*$.
- The verifier sends a random challenge c at random from \mathbb{Z}_q .
- The prover sends $z = \alpha + cx \pmod q$
- The verifier checks if $a = g^{-z}y^c$. If yes, the verifier accepts.

We describe here for completeness the non-interactive variant using Fiat-Shamir but remember that this approach will be used in parallel for several protocols for applications that require more than one Sigma protocol.

Protocol NI dlog : PK $\{(x) : y = g^x\}$

- The prover computes $a = g^\alpha$ for randomly chosen $\alpha \in \mathbb{Z}_q^*$, $c = H(y||a)$, $z = \alpha + cx$ and sends (c, z) to the verifier.
- The verifier computes if $a = y^{-c}g^z$ and checks that $c = H(y||a)$. If yes, the verifier accepts.

B.4 Sigma protocol for proving equality of committed value and ElGamal encrypted value

The prover has produced an ElGamal encryption $(e_1, e_2) = (g_1^R, g_1^x h_1^R)$ under public key g_1, h_1 that is an encryption of g_1^x with randomness R . They have also committed to x : $C = g^x h^r$ and wants to show that the same x appears in the encryption and in the commitment.

Protocol com-enc-eq : PK $\{(x, r, R) : e_1 = g_1^R \wedge e_2 = g_1^x h_1^R \wedge C = g^x h^r\}$

1. The prover computes $a_1 = g_1^\alpha, a_2 = g_1^\beta h_1^\alpha, a_3 = g^\beta h^\gamma$ for randomly chosen $\alpha, \beta, \gamma \in \mathbb{Z}_q^*$ and sends (a_1, a_2, a_3) to the verifier.
2. The verifier sends a random challenge c at random from \mathbb{Z}_q .
3. The prover computes $z_1 = \alpha + cR \pmod q, z_2 = \beta + cx \pmod q, z_3 = \gamma + cr \pmod q$ and sends (z_1, z_2, z_3) to the verifier.
4. The verifier checks if $a_1 = g_1^{-z_1} e_1^c, a_2 = g_1^{-z_2} h_1^{-z_1} e_2^c, a_3 = g^{-z_2} h^{-z_3} C^c$. If yes, the verifier accepts.

B.5 Proof of Equality of Aggregated Discrete Logs & Commitments

We now describe a protocol for proving equality of the discrete logarithms (a_1, \dots, a_n) in $y = \prod_{i=1}^n G_i^{a_i}$ and individual algebraic commitments to them. Using standard notation, we denote the protocol by PK $\{(a_1, \dots, a_n, r_1, \dots, r_n) : y = \prod_{i=1}^n G_i^{a_i} \wedge C_1 = g^{a_1} h^{r_1} \wedge \dots \wedge C_n = g^{a_n} h^{r_n}\}$.

Let \mathbb{G}, \mathbb{H} be groups of prime order q . Given $y = \prod G_i^{a_i}$ and $C_i = g^{a_i} h^{r_i}$, where G_i are generators of the group \mathbb{G} , g is a generator of \mathbb{H} and h is a random element in \mathbb{H} . The prover does not know the discrete logarithm of h with respect to g , and the discrete logarithms of G_i s with respect to each other. We want to prove equality of the discrete logarithms in y and the respective values committed to in C_i s.

Protocol comEq : $\text{PK}\{(a_1, \dots, a_n, r_1, \dots, r_n) : y = \prod_{i=1}^n G_i^{a_i} \wedge C_1 = g^{a_1} h^{r_1} \wedge \dots \wedge C_n = g^{a_n} h^{r_n}\}$

Given $y = \prod_{i=1}^n G_i^{a_i}$ and $C_i = g^{a_i} h^{r_i}$

1. The prover computes the following values: $u = \prod_{i=1}^n G_i^{\alpha_i}$ and $v_i = g^{\alpha_i} h^{R_i}$ for randomly chosen $\alpha_i, R_i \in \mathbb{Z}_q$ and sends u, v_i to the verifier.
2. The verifier chooses a challenge c at random from \mathbb{Z}_{2^k} for a fixed k , such that $2^k < q$, and sends it to the prover.
3. For a challenge string c , prover computes and sends the tuple (s_i, t_i)

$$s_i = \alpha_i - c a_i \pmod{q}, \quad t_i = R_i - c r_i \pmod{q}$$

4. Verification: Check if $u = y^c \prod G_i^{s_i}$ and $v_i = (C_i)^c g^{s_i} h^{t_i}$. The verifier accepts if checks succeed for all i .

B.6 Proof of Aggregated Discrete Logs

This is a specialization of the previous protocol and proves knowledge of discrete logarithms (a_1, \dots, a_n) in $y = \prod_{i=1}^n G_i^{a_i}$. Protocol is denoted by $\text{PK}\{(a_1, \dots, a_n) : y = \prod_{i=1}^n G_i^{a_i}\}$. The protocol is derived by simply ignoring in the previous protocol everything related to the commitments C_i .

Protocol AggregateDL : $\text{PK}\{(a_1, \dots, a_n) : y = \prod_{i=1}^n G_i^{a_i}\}$

Given $y = \prod_{i=1}^n G_i^{a_i}$

1. The prover computes the following values: $u = \prod_{i=1}^n G_i^{\alpha_i}$ for randomly chosen $\alpha_i \in \mathbb{Z}_q$ and sends u to the verifier.
2. The verifier chooses a challenge c at random from \mathbb{Z}_{2^k} for a fixed k , such that $2^k < q$, and sends it to the prover.
3. For a challenge string c , prover computes and sends the tuple (s_i)

$$s_i = \alpha_i - c a_i \pmod{q}$$

4. Verification: Check if $u = y^c \prod G_i^{s_i}$. The verifier accepts if checks succeed for all i .

B.7 Proof of Equality for Commitments in Different Groups

This is essentially a specialization of the protocol from section B.5. Assume we have discrete logarithms (a_1, a_2) in $y = G_1^{a_1} G_2^{a_2}$ which can be seen as a commitment to a_1 with randomness a_2 , as well as a commitment to a_1 in a different group. Using standard notation, we denote the protocol by $\text{PK}\{(a_1, a_2, r) : y = G_1^{a_1} G_2^{a_2} \wedge C = g^{a_1} h^r\}$.

Let \mathbb{G}, \mathbb{H} be groups of prime order q . Given $y = G_1^{a_1} G_2^{a_2}$ and $C = g^{a_1} h^r$, where G_i are generators of the group \mathbb{G} , g is a generator of \mathbb{H} and h is a random element in \mathbb{H} . The prover does not know the discrete logarithm of h with respect to g , and the discrete logarithms of G_i s with respect to each other. We want to prove that y and C are commitments to the same value.

Protocol $\text{PK}\{(a_1, a_2, r) : y = G_1^{a_1} G_2^{a_2} \wedge C = g^{a_1} h^r\}$

Given $y = G_1^{a_1} G_2^{a_2}$ and $C = g^{a_1} h^r$

1. The prover computes the following values: $u = G_1^{\alpha_1} G_2^{\alpha_2}$ and $v = g^{\alpha_1} h^R$ for randomly chosen $\alpha_i, R \in \mathbb{Z}_q$ and sends u, v to the verifier.
2. The verifier chooses a challenge c at random from \mathbb{Z}_{2^k} for a fixed k , such that $2^k < q$, and sends it to the prover.
3. For a challenge string c , prover computes and sends the tuple (s_1, s_2, t) , computed as follows:

$$s_1 = \alpha_1 - ca_1 \pmod{q}, \quad s_2 = \alpha_2 - ca_2 \pmod{q}, \quad t = R - cr \pmod{q}$$

4. Verification: Check if $u = y^c \prod G_1^{s_1} G_2^{s_2}$ and $v = C^c g^{s_1} h^t$. The verifier accepts if all checks succeed.

B.8 Proof of multiplicative relation on committed values

This protocol proves, for committed values a_1, a_2, a_3 in commitments C_1, C_2, C_3 , that $a_1 a_2 = a_3 \pmod{q}$. Protocol is denoted by $\text{PK}\{(a_1, a_2, a_3, r_1, r_2, r_3) : C_i = g^{a_i} h^{r_i} \text{ for } i = 1..3 \wedge a_1 a_2 = a_3 \pmod{q}\}$.

Protocol $\text{comMult} : \text{PK}\{(a_1, a_2, a_3, r_1, r_2, r_3) : C_i = g^{a_i} h^{r_i} \text{ for } i = 1..3 \wedge a_1 a_2 = a_3 \pmod{q}\}$

Given $C_i = g^{a_i} h^{r_i}$ for $i = 1, 2, 3$, we will prove the relation $a_1 a_2 = a_3 \pmod{q}$ by proving that $C_3 = C_1^{a_2} h^r$ in parallel with proving that $C_i = g^{a_i} h^{r_i}$ for $i = 1, 2, 3$. The condition is satisfied if we set $r = r_3 - r_1 a_2 \pmod{q}$, so the prover can do this.

If the proof succeeds, it follows that C_3 can be opened to both a_3 and $a_1 a_2$, so we get what we want unless the binding condition is broken.

1. The prover computes the following values: $v_i = g^{\alpha_i} h^{R_i}$ for $i = 1, 2, 3$, $v = C_1^{a_2} h^R$ for randomly chosen $\alpha_i, R_i, \alpha, R \in \mathbb{Z}_q$ and sends $\{v_i\}, v$ to the verifier.

Note that v will later be used in a verification equation to test $C_3 = C_1^{a_2} h^r$. This is why we need to use α_2 in the exponent of C_1 in the expression for v .

2. The verifier chooses a challenge c at random from \mathbb{Z}_{2^k} for a fixed k , such that $2^k < q$, and sends it to the prover.
3. For a challenge string c , prover computes and sends the tuple (s_i, t_i) for $i = 1, 2, 3$ and t :

$$s_i = \alpha_i - ca_i \pmod{q}, t_i = R_i - cr_i \pmod{q}, t = R - cr \pmod{q}$$

4. Verification: Check if $v_i = (C_i)^c g^{s_i} h^{t_i}$ for $i = 1, 2, 3$ and $v = C_3^c C_1^{s_2} h^t$. The verifier accepts if all checks succeed.

C Protocols for ID layer

C.1 Protocol for Π_{issue}

Protocol For proving that $ct = \text{TEnc}_{\text{PK}_{\text{AR}}}^{n,d}(m; r)$

Recall that we follow the share-and-encrypt paradigm for the threshold encryption by using Shamir Secret sharing and the CL encryption scheme. This means that the encryption of m is supposed to be consisting of n ciphertexts $\text{CL.Enc}(\text{pk}_{\text{AR}_i}, \text{sh}(m)_i)$ for $i = 1, \dots, n$ where $\text{sh}(m)_i$ is the i 'th share of m .

Let M' be the commitment of m on group $\tilde{\mathbb{G}}$. The prover AH proves to the verifier IP that M' contains the same value m as does ct :

1. AH makes a commitment B_0 to m under the default commitment key ck (which is in the CRS), so $B_0 \in \mathbb{H}$. They use the protocol from Section B.7 to show that M' and B_0 both contain m .
2. AH establishes a commitment M_i to each share $\text{sh}(m)_i$ of m as follows: when AH secret-shares m , they use a polynomial g to get shares $\text{sh}(m)_i$. Let the coefficients of g be $b_0 = m, b_1, \dots, b_d$. Note that we have $\text{sh}(m)_i = g(i) = \sum_{j=0}^d b_j \cdot i^j$. Now, the AH makes and sends commitments B_j to b_j , where we notice that B_0 has already been constructed above. Using the homomorphic property of commitments one can compute $K_i = \prod_{j=0}^d B_j^{i^j}$ which is a commitment to $\text{sh}(m)_i$. Since AH, the prover, has created the B_j s, they also know how to open the K_i s.
3. AH finally runs a protocol similar to `com-enc-eq` from section B.4 but for the CL encryption scheme to show the IP that each $\text{CL.Enc}(\text{pk}_{\text{AR}_i}, \text{sh}(m)_i)$ contains the same value as K_i for $i = 1 \dots n$.

C.2 Protocol for $\Pi_{\text{id-layer}}$

Proving that you know a signature. The user has a signature $\sigma = (a, b)$ on a message (m_0, \dots, m_ℓ) and wants to convince a verifier of this fact. This is done as follows:

Protocol For proving knowledge of a signature

1. Choose $r, r' \leftarrow \mathbb{F}_q$, and compute a blinded version of the signature as follows:

$$\hat{a} = a^r, \hat{b} = (b \cdot a^{r'})^r$$

and sends the blinded signature (\hat{a}, \hat{b}) to the verifier.

2. Both parties compute locally the following values:

$$v_1 = e(\hat{a}, \tilde{g}), v_2 = e(\hat{b}, \tilde{g}), v_3 = e(\hat{a}, \tilde{X}), u_i = e(\hat{a}, \tilde{Y}_i)$$

3. The verifier checks that $\hat{a} \neq 1_{\mathbb{G}}$. If not, the verifier rejects. The user gives a ZK proof of knowledge:

$$PK((m_0, \dots, m_\ell, r') : v_2 = v_3 \cdot v_1^{r'} \prod_{i=1}^{\ell} u_i^{m_i})$$

4. The verifier accepts if the proof is valid.

Adaptation for Implementing $\Pi_{\text{id-layer}}$. In the ID layer, the AH will act as user/prover and anyone else can act as verifier. However, the above protocol needs to be further adapted so we can show specific properties of attributes: recall that we set $(m_0, \dots, m_\ell) = (r, \text{IDcred}_{\text{SEC}}, K, a_1, \dots, a_v)$, padding the right-hand side with 0's if needed to get ℓ entries. The AH can supply also commitments to all fields and attributes: $i = 1 \dots \ell : C_i = \text{Commit}_{\text{ck}}(m_i, r_i)$. So we execute the protocol exactly as specified above, but the proof of knowledge in step 3 is extended as follows:

$$PK((m_0, \dots, m_\ell, r') : v_2 = v_3 \cdot v_1^{r'} \prod_{i=1}^{\ell} u_i^{m_i}, C_i = \text{Commit}_{\text{ck}}(m_i, r_i) \text{ for } i = 1 \dots \ell).$$

For this proof of knowledge, we use the **comEq** protocol from section B.5-

The AH can now use the C_i to prove desired properties of the fields and attributes. This is fleshed out later in this appendix.

Protocol For proving that $\text{RegID}_{\text{ACC}} = \text{PRF}_K(x)$ for some $x \leq \text{Max}_{\text{ACC}}$

Recall that $\text{PRF}_K(x) = \bar{g}^{1/(x+K)}$. This value can be seen as $\text{Commit}_{\text{ck}}(1/(K+x))$ where the randomness is 0. Let $C_1 = \text{Commit}_{\text{ck}}(K)$ and $C_3 = \text{Commit}_{\text{ck}}(a_1) = \text{Commit}_{\text{ck}}(\text{Max}_{\text{ACC}})$.

1. AH makes commitments $F_x = \text{Commit}_{\text{ck}}(x)$ and $F_1 = g$ which can be written as $\text{Commit}_{\text{ck}}(1)$ where the randomness is 0. Note that $C_1 F_x = \text{Commit}_{\text{ck}}(K+x)$.
2. AH runs Protocol **comMult** from Section B.8 with input commitments $C_1 F_x, \text{RegID}_{\text{ACC}}$ and F_1 . This shows that $\text{RegID}_{\text{ACC}} = \text{PRF}_K(x)$.
3. AH finally uses a zkSNARK on committed inputs based on F_x and C_3 to show that

$$x \leq \text{Max}_{\text{ACC}}.$$

Protocol For proving that $E_{\text{ID}} = \text{TEnc}_{\text{PK}_{\text{AR}}}^{\text{n,d}}(\text{IDcred}_{\text{PUB}}; r')$

Let C_0 be the commitment to $\text{IDcred}_{\text{SEC}}$. Recall that we follow share-and-encrypt paradigm for the threshold encryption (see A.4.1) which means in order to encrypt $\text{IDcred}_{\text{PUB}}$ correctly, AH first secret shares $\text{IDcred}_{\text{SEC}}$ using a polynomial f to get shares $\text{sh}(\text{IDcred}_{\text{SEC}})_i$ and then computes the i 'th component of E_{ID} as Elgamal encryption $\text{Enc}_{\text{pk}_{\text{AR}_i}}(\bar{g}^{\text{sh}(\text{IDcred}_{\text{SEC}})_i}; r'_i)$. Let the coefficients of f be $a_0 = \text{IDcred}_{\text{SEC}}, a_1, \dots, a_d$. Note that we have $\text{sh}(\text{IDcred}_{\text{SEC}})_i = f(i) = \sum_{j=0}^d a_j \cdot i^j$. Now,

1. AH makes commitments A_j to a_j , where we set $A_0 = C_0$. Using the homomorphic property of commitments one can compute $S_i = \prod_{j=0}^d A_j^{i^j}$ which is a commitment to $\text{sh}(\text{IDcred}_{\text{SEC}})_i$. Since AH, the prover, has created the A_i s, they also know how to open the S_i s.
2. For $i = 1 \dots n$, the AH runs the protocol `com-enc-eq` from section B.4 using the two components in $E_{\text{pk}_{\text{AR}_i}}(\bar{g}^{\text{sh}(\text{IDcred}_{\text{SEC}})_i}; r'_i)$ as e_1, e_2 and S_i as C .

D CL Framework

In [CL15], Castagnos and Laguillaumie introduced the so-called CL framework, which allows to construct cyclic groups \mathbb{G} where the decisional Diffie-Hellman (DDH) assumption is believed to hold and furthermore, there exists a subgroup \mathbb{H} of \mathbb{G} such that the discrete logarithm problem in \mathbb{H} is easy. As the main application of this framework, they construct the first practical linearly homomorphic encryption scheme where the plaintext space is \mathbb{Z}_q with prime order q , where the size of q can be made independent of the security parameter. We make use of the CL encryption schemes in some of our constructions. In particular, the CL framework allows to encrypt values “in the exponent”, so that proving statements about the message is easy using efficient Σ -protocols in [CCL⁺20], while at the same time allowing for efficient decryption. (As opposed to standard ElGamal encryption, where it is necessary to limit the size of the message in the exponent to allow for decryption). This will be useful both to instantiate our threshold encryption scheme (used to allow anonymity revokers to “trace” a suspicious account holder), and in our construction of extractable NIZK proofs in the crs model (see F).

In the following, we recall the definition of CL framework from [CCL⁺20].

D.1 Definition

Let λ be a positive integer and q be a μ -bit prime for $\mu \geq \lambda$. The framework Gen_{CL} consists of two algorithms $\text{Gen}_{\text{CL}} = (\text{Gen}, \text{Solve})$ defined as follows:

- $\text{Gen}(1^\lambda, q) \rightarrow \mathfrak{p} := (\widehat{\mathbb{G}}, \mathbb{G}, \mathbb{H}, \mathbb{G}^q, \tilde{s}, g, h, g_q)$: the algorithm takes the security parameter λ and a prime q as inputs and outputs a tuple $\mathfrak{p} := (\widehat{\mathbb{G}}, \mathbb{G}, \mathbb{H}, \mathbb{G}^q, \tilde{s}, g, h, g_q)$, where

- $(\widehat{\mathbb{G}}, \cdot)$ is a finite abelian group with order $\widehat{n} := q \cdot \widehat{s}$ where the bitsize of \widehat{s} is a function of λ and $\gcd(q, \widehat{s}) = 1$. It is required that valid encodings of elements in $\widehat{\mathbb{G}}$ can efficiently be recognized.
 - (\mathbb{G}, \cdot) is a cyclic subgroup of $\widehat{\mathbb{G}}$ of order $n := q \cdot s$ where s divides \widehat{s} .
 - (\mathbb{H}, \cdot) is the unique cyclic subgroup of $\widehat{\mathbb{G}}$ of order q , generated by h .
 - $\mathbb{G}^q := \{x^q | x \in \mathbb{G}\}$ is the subgroup of \mathbb{G} of order s . Since $\mathbb{H} \subset \mathbb{G}$, it holds that $\mathbb{G} \simeq \mathbb{G}^q \times \mathbb{H}$.
 - \tilde{s} is an upper bound of \widehat{s} .
 - g, h and g_q are respectively the generators of \mathbb{G}, \mathbb{H} and \mathbb{G}^q , where $g := h \cdot g_q$.
- $\text{Solve}(q, \mathfrak{p}, X) \rightarrow x'$: this is a DPT algorithm that solves the discrete logarithm problem in \mathbb{H} :

$$\Pr \left[x = x' : \mathfrak{p} \leftarrow \text{Gen}(1^\lambda, q), x \leftarrow \mathbb{Z}/q\mathbb{Z}, X \leftarrow h^x \right. \\ \left. x' \leftarrow \text{Solve}(q, \mathfrak{p}, X) \right] = 1$$

D.2 Hard Subgroup Membership Problem

We now recall the hard subgroup membership problem within a group with an easy DL subgroup (HSM-CL) from [CLT18]. The HSM-CL assumption states that it is hard to distinguish the elements of \mathbb{G}^q in \mathbb{G} .

Definition D.1 (HSM-CL assumption). *Let λ be a positive integer and $\text{Gen}_{\text{CL}} = (\text{Gen}, \text{Solve})$ be a generator that generates a group with an easy DL subgroup as defined above. Let \mathcal{D} (resp. \mathcal{D}_q) be a distribution over the integers such that the distance between the distribution $\{g^x, x \leftarrow \mathcal{D}\}$ (resp. $\{g_q^x, x \leftarrow \mathcal{D}_q\}$) and uniform distribution in \mathbb{G} (resp. in \mathbb{G}^q) is less than $\delta(\lambda) = \text{negl}(\lambda)$. The advantage of an adversary \mathcal{A} for the HSM-CL problem is now defined as*

$$\text{Adv}_{\mathcal{A}}^{\text{hsm-cl}}(\lambda) = |\Pr[b = b' : \mathfrak{p} \leftarrow \text{Gen}(1^\lambda, q), |q| \geq \lambda, x \leftarrow \mathcal{D}, x' \leftarrow \mathcal{D}_q, \\ b \leftarrow \{0, 1\}, X_0 \leftarrow g^x, X_1 \leftarrow g_q^{x'}, b' \leftarrow \mathcal{A}(q, \mathfrak{p}, X_b, \text{Solve}(\cdot))] - 1/2|$$

We say that the HSM-CL problem is ε -hard for Gen_{CL} if for all PPT adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\text{hsm-cl}}(\lambda) \leq \varepsilon(\lambda)$. And we say HSM-CL holds for Gen_{CL} if it is ε -hard for Gen_{CL} and $\varepsilon(\lambda) = \text{negl}(\lambda)$.

D.3 PKE scheme under HSM-CL assumption

We recall the scheme described in [CLT18] with the following setting: the plaintext space is \mathbb{Z}_q for μ -bit prime q and $\mu \geq \lambda$; and the secret key x and randomness r are drawn from distribution \mathcal{D}_q . As shown in [CLT18] (lemma 4), to instantiate \mathcal{D}_q in practice, one can use the uniform distribution on $\{0, \dots, S\}$ for $S := 2^{\lambda-2} \cdot \tilde{s}$. The scheme is depicted in Fig.

Theorem D.2. [CLT18] *The scheme described in Fig. 6 is semantically secure under chosen plaintext attacks (IND-CPA) under the HSM-CL assumption.*

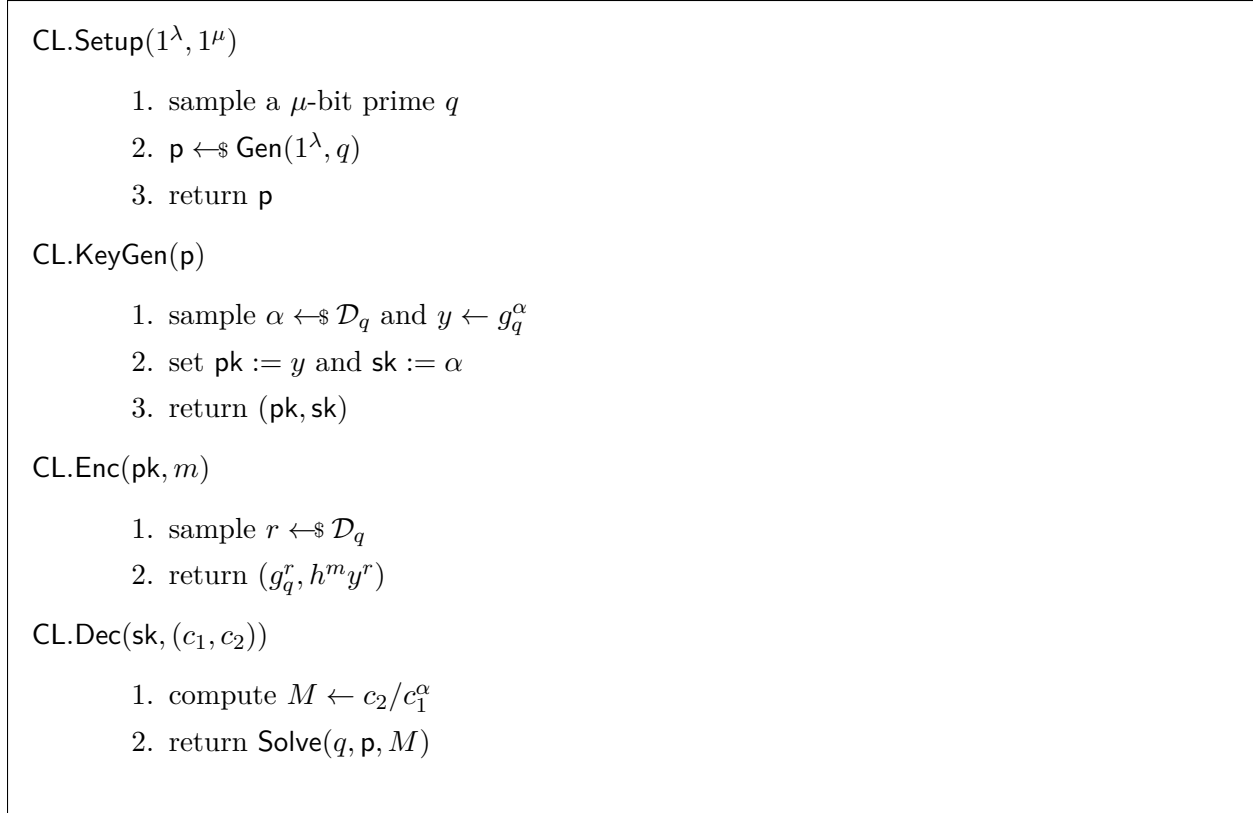


Figure 6: Encryption scheme from HSM-CL [CLT18]

E Transaction Layer

Our paper shows how account holders in our design can create accounts after registering with identity providers. As we have seen, an account includes an account-specific public key where the account holder has the private key. Since our protocol is completely generic, the “key” can in fact be a vector of keys which includes a key for encryption, one for signatures, etc. and therefore accounts can be used for transactions in many different ways.

In this section, we include an informal presentation of one way in which accounts can be used for transferring money on the blockchain, assuming each account has a public signing key and a public encryption key. We stress that this is just as example of one of many possible ways of doing this and, as we prove our ID layer secure in the UC framework one can use our ID layer with any other transaction layer, or even other applications not involving payments.

The system we sketch supports plaintext transactions and encrypted transactions. All transactions must be signed by the account holder from which the transaction originates. When a transaction is published, the ledger will check the signature and allow the transaction to go through if the signature is valid, and possibly if other constraints (described in detail below) are satisfied. Note that payments are linked to the sending and the receiving account and if several payments are made between the same accounts this will be visible on chain. We allow this to have a trade-off between efficiency and privacy: one can make several payments from an account and only suffer

the cost of opening it once. On the other hand, one can also choose to use each account once for complete privacy.

Let ACC be an account with encryption key ek_{ACC} . An account will hold a public amount p and a secret amount s . The secret amount is represented as a set $\mathcal{S} = \{S_i\}_{i=1}^{|\mathcal{S}|}$, where $S_i = \text{Enc}_{\text{ek}_{\text{ACC}}}(s_i)$ and $s = \sum_i s_i$.

E.0.1 Plaintext Transactions.

Plaintext transactions happens by a matched reduction and increment of the public amounts on the sending and receiving accounts.

E.0.2 Encrypted Transactions.

Let $\mathcal{S}_1 = \{S_i\}_{i=1}^n$ be the representation of the secret amount for account ACC_1 owned by AH_1 , where $S_i = \text{Enc}_{\text{ek}_{\text{ACC}_1}}(s_i)$. To do an encrypted transaction of amount a from ACC_1 to some account ACC_2 , AH_1 proceeds as follows:

1. Compute $s' = \sum_{i=1}^n s_i - a$.
2. Create $S' = \text{Enc}_{\text{ek}_{\text{ACC}_1}}(s')$.
3. Create $A = \text{Enc}_{\text{ek}_{\text{ACC}_2}}(a)$.
4. Compute a NIZK proof π that (S', A, S_1, \dots, S_n) contains numbers (s', a, s_1, \dots, s_n) such that

$$a \geq 0, \tag{1}$$

$$s' \geq 0, \tag{2}$$

$$s' = \sum_{i=1}^n s_i - a. \tag{3}$$

The transaction contains (S', A, π) . When the transaction is executed, the ledger checks the proof and if it is valid, then let $\mathcal{S}_1 = \{S'\}$ and $\mathcal{S}_2 = \mathcal{S}_2 \cup \{A\}$, where \mathcal{S}_2 is the representation of the secret amount for ACC_2 .

E.0.3 Compressing an account.

From the above, it follows that after receiving some number of transactions, the receiving account will contain several encrypted amounts, and this may become impractical to handle at some point. Also, note that an account owner does not actually know the amount on his account until he has decrypted the transactions that come into his account. To solve this, the account owner of ACC_2 can execute a compression transaction, which works as follows:

1. Let $\mathcal{S} = \{S_i\}$ be the representation of existing secret amount for ACC_2 . Use decryption key dk_{ACC_2} to decrypt all S_i and let s be the sum of the decrypted amounts. Compute $e = \text{Enc}_{\text{ek}_{\text{ACC}_2}}(s)$ and let $S' = \{e\}$. Compute a NIZK proof π showing that e contains the sum of the amounts in the S_i 's.
2. publish a compression transaction that contains the identity of the account and S', π .

When a compression transaction appears, the ledger checks the proof π and if it is valid, then update the account of ACC_2 to contain S' instead of \mathcal{S} .

F Universally Composable Non-interactive Zero-Knowledge

F.1 Simulation Extractable NIZK

In this section we briefly recall how to construct a simulation extractable NIZK $\Pi' = (\text{CRSGen}', \text{Prove}', \text{Verify}')$ from a simulation-sound NIZK $\Pi = (\text{CRSGen}, \text{Prove}, \text{Verify})$ and a CPA-secure encryption scheme $\text{ES} = (\text{KeyGen}, \text{Enc}, \text{Dec})$. While this is a folklore compiler and similar result appears in [Bag19], we give the complete proof here for the sake of completeness.

In Figure 7 is described the construction of the simulation-sound NIZK $\Pi' = (\text{CRSGen}', \text{Prove}', \text{Verify}')$ for the NP-language \mathcal{L} . $\Pi = (\text{CRSGen}, \text{Prove}, \text{Verify})$ is associated to the NP-language \mathcal{L}' and the relation $\mathcal{R}' = \{(x, \text{enc}, \text{pk}) : (\text{sk}, w) \text{ s.t. } w \leftarrow \text{Dec}(\text{enc}, \text{pk}, \text{sk}) \wedge (x, w) \in \mathcal{R}\}$.

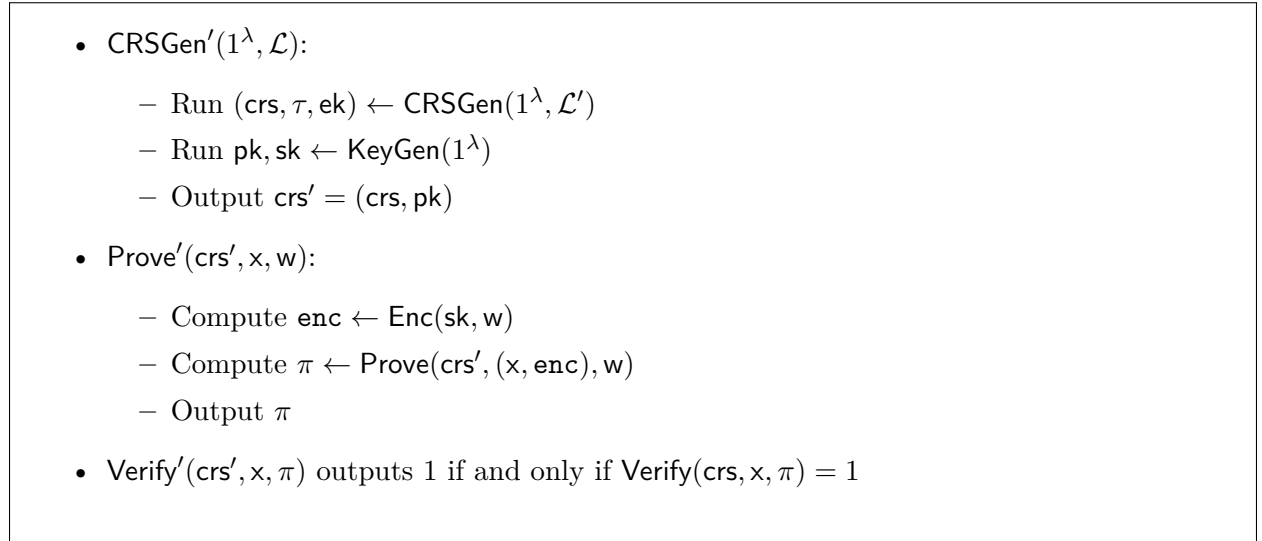


Figure 7: Simulation-extractable NIZK.

Theorem F.1. *Assume that $\text{ES} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is a CPA-secure encryption scheme and $\Pi = (\text{CRSGen}, \text{Prove}, \text{Verify})$ is a simulation-sound NIZK for the NP-language \mathcal{L}' . Then $\Pi' = (\text{CRSGen}', \text{Prove}', \text{Verify}')$ for the NP-language \mathcal{L} , described in Figure 7, is simulation extractable.*

Proof. (Sketch.) Let Sim the corresponding simulator associated to Π . Let us now describe the simulator-extractor Sim' of Π' .

CRS generation: Sim' runs Sim to obtain the simulated CRS $\text{crs}, \tau, \text{ek}$ and computes $\text{pk}, \text{sk} \leftarrow \text{KeyGen}(1^\lambda)$. They set the $\text{crs}' = (\text{crs}, \text{pk})$.

Simulation: When \mathcal{A} asks to see a proof for the statement x , Sim' computes an encryption of a dummy value enc and runs Sim on input $(x, \text{enc}, \tau, \text{ek})$ to obtain π .

Extraction: On input $\bar{\pi}' = (\bar{\pi}, \bar{\text{enc}})$ received by \mathcal{A} the simulator Sim runs the decryption algorithm of ES and outputs \bar{w} for theorem \bar{x} .

We first observe that the zero-knowledge property of Π' follows from the zero-knowledge property of Π and the CPA-security of ES .

The proof proceeds by contradiction. Let us assume that \mathcal{A} is able to produce valid proofs but the extraction procedure described above fails with non-negligible probability. If this is the

case, then \mathcal{A} is able to produce valid proofs (different from the one received from the oracle) for the statements of the form $(\bar{x}, \overline{\text{enc}}\bar{c})$, where $\overline{\text{enc}}\bar{c}$ is not an encryption of a valid \bar{w} for $x \in \mathcal{L}$ (by assumption) and therefore we can show a reduction to the simulation-sound property of Π . \square

F.2 SNARK-Lifting transformations via $\text{C}\emptyset\text{C}\emptyset$ framework [KZM⁺15]

To transform a SNARK into a strong simulation extractable SNARK, Kosba et al. [KZM⁺15] makes use of four primitives: a pseudo-random function, a perfectly-binding commitment scheme, a one-time signature (OTS) and a public key encryption (PKE) scheme. Let PRF_K be a pseudorandom function with key K such that without the knowledge of K , $\text{PRF}_K(\cdot)$ behaves like a random function, whereas given K , computing $\text{PRF}_K(\cdot)$ is easy. Let $\text{COM} = (\text{Commit}, \text{OpenVrf})$ be a commitment scheme. Let $\Omega = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a PKE scheme and $\Sigma_{\text{OT}} = (\text{KeyGen}, \text{Sign}, \text{VerifySig})$ be a OTS scheme. Given a language \mathcal{L} with NP relation $\mathcal{R}_{\mathcal{L}}$, let \mathcal{L}' be the language with relation $\mathcal{R}_{\mathcal{L}'}$ defined as follows: $\{(x, c, \mu, \text{pk}_{\text{OT}}, \text{pk}_e, \rho), (w, r_1, r_0, K)\} \in \mathcal{R}_{\mathcal{L}'}$ iff:

$$c = \text{Enc}_{\text{pk}_e}(w; r_1) \wedge \left((x, w) \in \mathcal{R}_{\mathcal{L}} \vee \right. \\ \left. (\mu = \text{PRF}_K(\text{vk}_{\text{OT}}) \wedge \rho = \text{Commit}(K; r_0)) \right)$$

The transformation from a NIZK Π to one that satisfies SSE is then defined in Fig. 8. The following theorem is taken from [KZM⁺15].

Theorem F.2. *Assume that the NIZK scheme $\Pi = (\text{CRSGen}, \text{Prove}, \text{Verify})$ satisfies perfect completeness, computational soundness, computational zero-knowledge, that the encryption scheme $\Omega = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is semantically secure and perfectly correct, that the pseudo-random function family PRF_K is secure, that the commitment scheme $\text{COM} = (\text{Commit}, \text{OpenVrf})$ is perfectly binding and computational hiding, and that the one-time signature scheme is strongly unforgeable. Then the resulting transformation described in Fig. 8 is a zero-knowledge proof system satisfying perfect completeness, computational zero-knowledge, and strongly simulation extractability as defined in Section A.6.*

G Standard Ideal Functionalities

We describe here some of the (standard) ideal functionalities that are used in our protocol design.

Functionality \mathcal{F}_{crs}

The functionality is parametrized by a distribution \mathcal{D} and proceeds as follows.

- Choose a value $\text{crs} \leftarrow \mathcal{D}$.
- On input (CRS) from a party P , return (CRS, crs) to P .

<p>CRSGen($1^\lambda, \mathcal{L}$)</p> <hr/> <ul style="list-style-type: none"> - $\Pi.crs \leftarrow \Pi.CRSGen(1^\lambda, \mathcal{L}'); (pk_e, sk_e) \leftarrow \Omega.KeyGen(1^\lambda);$ - $\tau := (s_0, r_0) \leftarrow_{\\$} \{0, 1\}^\lambda; \rho \leftarrow \text{Commit}(s_0; r_0);$ - return $(crs := (\Pi.crs, pk_e, \rho), \tau, ek := sk_e)$
<p>Prove(crs, x, w)</p> <hr/> <ul style="list-style-type: none"> - Parse $crs := (\Pi.crs, pk_e, \rho);$ - $(vk_{OT}, sk_{OT}) \leftarrow \Sigma_{OT}.KeyGen(1^\lambda); r_1, z_0, z_1, z_2 \leftarrow_{\\$} \{0, 1\}^\lambda;$ - $c = \Omega.Enc(pk_e, w; r_1);$ - $\pi_\Pi \leftarrow \Pi.Prove(\Pi.crs, (x, c, z_0, pk_e, vk_{OT}, \rho), (w, r_1, z_1, z_2));$ - $\sigma_{OT} \leftarrow \Sigma_{OT}.Sign(sk_{OT}, (x, c, z_0, \pi_\Pi));$ - return $\pi := (c, z_0, \pi_\Pi, vk_{OT}, \sigma_{OT});$
<p>Verify(crs, x, π)</p> <hr/> <ul style="list-style-type: none"> - Parse $crs := (\Pi.crs, pk_e, \rho)$ and $\pi := (c, \mu, \pi_\Pi, vk_{OT}, \sigma_{OT});$ - if $\Sigma_{OT}.Verify(vk_{OT}, (x, c, \mu, \pi_\Pi), \sigma_{OT}) = 0$ - $\vee \Pi.Verify(\Pi.crs, (x, c, \mu, pk_e, vk_{OT}, \rho), \pi_\Pi) = 0$ - then return 0 else return 1;
<p>SimProve(crs, τ, x)</p> <hr/> <ul style="list-style-type: none"> - Parse $crs := (\Pi.crs, pk_e, \rho)$ and $\tau := (s_0, r_0);$ - $(vk_{OT}, sk_{OT}) \leftarrow \Sigma_{OT}.KeyGen(1^\lambda); \mu = \text{PRF}_{s_0}(vk_{OT});$ - $r_1, z_3 \leftarrow_{\\$} \{0, 1\}^\lambda; c \leftarrow \Omega.Enc(pk_e, z_3; r_1);$ - $\pi_\Pi \leftarrow \Pi.Prove(\Pi.crs, (x, c, \mu, pk_e, vk_{OT}, \rho), (r_1, r_0, z_3, s_0));$ - $\sigma_{OT} \leftarrow \Sigma_{OT}.Sign(sk_{OT}, (x, c, \mu, \pi_\Pi));$ - return $\pi := (c, \mu, \pi_\Pi, vk_{OT}, \sigma_{OT});$
<p>Ext(crs, x, π, ek)</p> <hr/> <ul style="list-style-type: none"> - Parse $\pi := (c, \mu, \pi_\Pi, vk_{OT}, \sigma_{OT});$ - return $w \leftarrow \Omega.Dec(ek, c);$

Figure 8: The strong version of $C\emptyset C\emptyset$ transformation [KZM⁺15]

Functionality \mathcal{F}_{smt}

The functionality models a secure channel between a sender S and a receiver R .

- Upon input (SEND, R, m) from a party S , if both S and R are honest, output (SENT, S, m) to R and $(\text{SENT}, S, R, |m|)$ to \mathcal{A} . Otherwise, if at least one of S and R is corrupt, output (SENT, S, R, m) to \mathcal{A} .

G.0.1 Registration Functionality

The functionality allows identity providers and anonymity revokers to input a key pair $(\text{sk}, \text{pk}) \in \{0, 1\}^* \times \{0, 1\}^*$ once such that they would not be allowed to modify or delete it later. The account holders can retrieve the public keys of registered parties by the `RETRIEVE` command.

Functionality \mathcal{F}_{reg}

Register

Upon receiving a message $(\text{REGISTER}, \text{sk}_P, \text{pk}_P)$ from party P , if this is the first request from P , then keep the record $(P, \text{sk}_P, \text{pk}_P)$ and return $(\text{INITIALIZED}, P)$. Otherwise, ignore the message.

Retrieve

Upon receiving a message $(\text{RETRIEVE}, P_i)$ from some party P_j or the adversary, output $(\text{RETRIEVE}, P_i, \text{pk}_{P_i})$ to P_j where $\text{pk}_{P_i} = \perp$ if no record $(P_i, \text{sk}_{P_i}, \text{pk}_{P_i})$ exists.

G.0.2 NIZK Functionality

We use $\mathcal{F}_{\text{nizk}}$ as defined by Groth et al. in [GOS12].

Functionality $\mathcal{F}_{\text{nizk}}$

The functionality is parametrized by a relation \mathcal{R} for which we can efficiently check membership.

Proof

- On input (PROVE, x, w) from party P , ignore if $(x, w) \notin \mathcal{R}$. Send (PROVE, x) to \mathcal{A} .
- Upon receiving the answer (PROOF, π) from \mathcal{A} , store (x, π) and send (PROOF, π) to P .

Verify

- On input (VERIFY, x, π) from V check whether (x, π) is stored. If not send (VERIFY, x, π) to \mathcal{A} .
- Upon receiving the answer $(\text{WITNESS}, w)$ from \mathcal{A} , check $(x, w) \in \mathcal{R}$ and if so, store (x, π) . If (x, π) has been stored, output $(\text{VERIFICATION}, 1)$ to V , else output $(\text{VERIFICATION}, 0)$.

G.0.3 Ledger Functionality

The functionality $\mathcal{F}_{\text{ledger}}$ (which is an adapted and simplified version of the functionality given by [KZZ16]) abstracts a public ledger which can be accessed globally by protocol parties or en-

vironment \mathcal{Z} . When accounts are created and posted by account holders, the ideal functionality first validates the account and then appends them to the buffer. The environment can later specify when to free the buffer and append it to the ledger.

Functionality $\mathcal{F}_{\text{ledger}}$

The functionality is globally available to all parties and is parameterized by a predicate VALIDATE , an initially empty list L of bit strings and a variable buffer initially set to ε .

Append

Upon input (APPEND, x) from party P , if $\text{VALIDATE}(\text{state}, (\text{buffer}, x)) = 1$, then set $\text{buffer} \leftarrow \text{buffer} || x$.

Retrieve

Upon input (RETRIEVE) from a party P or \mathcal{A} , return $(\text{RETRIEVE}, L)$ to the requestor in case of an honest party or $(\text{RETRIEVE}, L, \text{buffer})$ in case of a corrupt party.

Buffer Release

Upon input $(\text{RELEASE}, \Pi)$ from \mathcal{A} , append a permutation Π of buffer to L by setting $L \leftarrow L || \Pi(\text{buffer})$ and set $\text{buffer} := \varepsilon$.

G.0.4 Functionality for MPC of PRF output

This functionality allows a set of parties who own secret shares of a PRF key to evaluate the PRF on a fixed set of inputs.

Functionality $\mathcal{F}_{\text{mpc-prf}}$

The functionality interacts with parties P_1, \dots, P_n and is parameterized by a pseudo-random function PRF , a constant Max_{ACC} , and a (n, d) -secret sharing scheme SS .

Compute

Upon input $(\text{COMPUTE}, K_i)$ from at least $d + 1$ distinct P_i , proceed as follows:

- compute $K = \text{Reconstruct}^{n,d}(K_{i_1}, \dots, K_{i_{d+1}})$.
- evaluate $\text{PRF}_K(x)$ on all inputs $x \in [\text{Max}_{\text{ACC}}]$ and return the list of outputs to all the requesters P_i for $i \in n$.