

Protego: Efficient, Revocable and Auditable Anonymous Credentials with Applications to Hyperledger Fabric

Aisling Connolly¹, Jerome Deschamps²,
Pascal Lafourcade², and Octavio Perez Kempner^{3,4}

¹ DFINITY*

`aislingmconnolly@gmail.com`

² LIMOS, University Clermont Auvergne, France

`jerome.deschamps@etu.uca.fr`

`pascal.lafourcade@uca.fr`

³ DIENS, École normale supérieure, CNRS, PSL University, Paris, France

⁴ be-ys Research, France

`octavio.perez.kempner@ens.fr`

Abstract. Recent works to improve privacy in permissioned blockchains like Hyperledger Fabric rely on Idemix, the only anonymous credential system that has been integrated to date. The current Idemix implementation in Hyperledger Fabric (v2.4) only supports a fixed set of attributes; it does not support revocation features, nor does it support anonymous endorsement of transactions (in Fabric, transactions need to be approved by a subset of peers before consensus). A prototype Idemix extension by Bogatov *et al.* (CANS, 2021) was proposed to include revocation, auditability, and to gain privacy for users. In this work, we explore how to gain efficiency, functionality, and further privacy, departing from recent works on anonymous credentials based on Structure-Preserving Signatures on Equivalence Classes. As a result, we extend previous works to build a new anonymous credential scheme called Protego. We also present a variant of it (Protego Duo) based on a different approach to hiding the identity of an issuer during showings. We also discuss how both can be integrated into Hyperledger Fabric and provide a prototype implementation. Finally, our results show that Protego and Protego Duo are at least twice as fast as state-of-the-art approaches based on Idemix.

Keywords: anonymous credentials, auditability, Hyperledger Fabric, mercurial signatures, permissioned blockchains.

1 Introduction

When first introduced, the core use of blockchains was in the *permissionless* setting; anyone could join and participate. Over the years, blockchains have also found use within consortiums, where several authorized organizations wish

* Work done while the author was at Wordline Global.

to share information among the group, but not necessarily to the public as a whole. This need gave rise to *permissioned* blockchains whereby authorities are established to define a set of participants. When a federation of authorities (*consortium*), each in control of a subset of participants, shares the blockchain’s governance, the term *federated* is also used to describe such blockchains.

The use of federated blockchains increased to address the need to run a common business logic within a closed environment. As an example, one can consider pharmaceutical companies that would like to trade sensitive information about product developments and agree on supplies or prices in a consortium with partial trust. A recurrent problem in such scenarios is that of privacy while being compliant with regulations and *Know Your Customer* practices. Agreeing with other entities to run a shared business logic should not imply that everything needs be public within the consortium. Privacy still needs to be provided without affecting existing regulations, *e.g.*, when considering bilateral agreements.

The most developed permissioned platform is Hyperledger Fabric (or simply Fabric). In Fabric, users submit transaction proposals to a subset of peers (called *endorsers*) that vouch for their execution. By default, it provides no privacy features as everything (users and transactions) is public. Reading the blockchain anyone can know (1) who triggered a smart contract using which arguments (transaction proposals are signed by the clients), (2) who vouched for its execution (endorsers also sign their responses) concerning reading and writing sets; and (3) why a given transaction was marked as invalid (either because of invalid read/write sets or because the endorsement policy check failed). Furthermore, checking access control and endorsement policies links different organizations, users and their attributes to concrete actions on the system.

Such limitations severely restrict the use of Fabric. From the user perspective this impacts the enforcement of different regulations. For organizations, the case is similar. Consider a consortium of pharmaceuticals that run a common business logic to exchange information on medical research. If the entity behind a request is known, other organizations can infer (based on the request) which drug the entity in question is trying to develop. If the endorsers are known, information about who executes what can disclose business relations.

Motivated by the need to protect business interests and to meet regulatory requirements, some privacy features were integrated using the Identity Mixer [CV02, Zur13] (or Idemix for short). This anonymous attribute-based credential (ABC) scheme gave the first glimpse of privacy for users within a consortium. Idemix allows a Membership Service Provider (MSP) to issue credentials enabling users to sign transactions anonymously. In brief, users generate a zero-knowledge proof attesting that the MSP issued them a credential on its attributes to sign a transaction. Fabric’s support for Idemix was added in v1.3, providing the first solution to tackle the problem of *participant* privacy. Unfortunately, as for v2.4 the Idemix implementation still suffers severe limitations:

1. It supports a fixed set of only four attributes.
2. It does not support revocation features.

3. Credentials leak the MSP ID, meaning that anonymity is local to users within an organization. For this reason, current deployments can only use a single MSP for the whole network, introducing a single point of failure.
4. It does not support the issuance of Idemix credentials for the endorsing peers, meaning that the identity of endorsers is always leaked.

The most promising effort to extend the functionality of Idemix appeared in [BDCET21]. Their aim was to extend the original credential system to support *delegatable credentials* [CDD17], while integrating revocation and auditability features (solving three of the four limitations). Below we outline the main ideas introduced in [BDCET21].

Delegatable Credentials. In a bid to overcome the issue of Idemix credentials leaking the MSP ID and thus the affiliation of the user, a trusted root authority provides credentials to intermediate authorities. This way users can obtain credentials from intermediate authorities. To sign a transaction, the user must generate a zero-knowledge proof attesting that (1) the signer owns the credential; (2) the signature is valid; (3) all adjacent delegation levels are legitimate; and (4) that the top-level public key belongs to the root authority.

Revocation and Auditability. To generate efficient proofs of non-revocation, the system timeline is divided into *epochs*. Issued credentials are only valid for a given epoch, and must be reissued as the timeline advances. For each epoch, a user requests a revocation handle that binds their public key to the epoch. When presenting a credential, the user also provides a proof of non-revocation. To enable auditing of a transaction, users verifiably encrypt their public key under an authorized auditor’s public key.

To date, some functionalities remain limited. (1) There is still no notion of privacy for endorsers. (2) Delegatable credentials require proving knowledge of a list of keys. (3) The root authority is still a single point of failure. (4) Selective disclosure of attributes requires computation linear in the size of all the attributes encoded in the credential. (5) Many zero-knowledge proofs need to be generated for each transaction. (6) Many pairings need to be computed for verification.

Recent results [FHS19, DHS15, CLPK22] introduced newer models based on Structure-Preserving Signatures on Equivalence Classes (SPS-EQ) to build ABC’s, providing a host of extra functionalities and more efficient constructions. The main goal of this work is to leverage such results, position them in the blockchain scenario and provide an alternative to Idemix (and its extension) to overcome existing privacy and functional limitations while also improving efficiency.

Contributions. To build an ABC scheme that overcomes the inherent limitations from Idemix and its extension, we argue that changing some of the underlying building blocks is necessary. Therefore, we take the framework from [FHS19] as our starting point, incorporate the recent improvements from [CLPK22], and include the revocation extension originally proposed in [DHS15]. From there, we

extend the ABC model to support auditability features and adapt it to non-interactive showings. To do so, we rely on the random oracle model (already present in the blockchain setting). We also present and discuss two alternatives to the use of delegatable credentials (as used in [BDCET21]) to hide the identity of credential issuers, following the formalizations from [CLPK22] and [BEK⁺21] but using new approaches.

When compared to [CLPK22], the modifications are: (1) we adapt the ABC model to non-interactive showings, (2) we extend the model defining a revocation authority as in [DHS15], and an auditing authority (not considered in the previous works), (3) we keep the SCDS scheme from [CLPK22] as it is but replace the signature scheme with the one given in [CL19], and (4) we build a *malleable* NIZK argument that can be pre-computed to obtain a more efficient issuer-hiding feature.

As a result, we build Protego, an ABC suitable for permissioned blockchains. We also present Protego Duo, a variant based on a different approach to hide the identity of credential issuers. Both support revocation and auditing features, which are important to enable a broader variety of use cases for permissioned blockchains. We discuss how to integrate our work with Fabric, compare it with Idemix and its recent extensions, and provide a prototype implementation showing that Protego and Protego Duo are faster than the most recent Idemix extension (see Section 5 for a detailed evaluation and benchmarks). Furthermore, a showing proof in Protego Duo is constant-size (8.3 kB), surpassing [BDCET21] in which the proof size grows linearly with the number of *attributes and delegation levels*.

Related work. We describe the related work following two main streams; the results addressing privacy concerns in Fabric, and parallel research developments.

Privacy concerns in Fabric. The most closely related work appears with the introduction of Idemix [CV02, Zur13] and its extension to include revocation and auditability [BDCET21]. Adding auditability is crucial for permissioned blockchains as they are often used in heavily regulated industries. Privacy-preserving auditing for distributed ledgers was introduced in [NVV18] under the guise of zkLedger. This general solution offered great functionality in that it provided confidentiality of transactions, and privacy of the users within the transaction. However, it assumed low transaction volume between few participants and as such is quite limited in scalability. Fabric-friendly auditable token payments were introduced in [ACC⁺20] and were based on threshold blind signatures. The core idea to achieve auditability was to encrypt the user’s public key under the public key of an auditor. This is the same approach in [BDCET21], which we also use in this work. Although the auditing ideas are similar, the construction pertains solely to transaction privacy and offers no identity privacy for a user. Following the approach of gaining auditability of transactions, auditable smart contracts were captured by FabZK [KDJL⁺19] which is based on Pedersen commitments and zero-knowledge proofs. To achieve auditability, the structure of the ledger is modified, and as such, would need to make considerable changes to existing used permissioned blockchain platforms.

One of the limitations in Idemix and its extension is the lack of privacy or anonymity for endorsing peers. A potential solution to this was proposed in [MR19], where the endorsement policy is based on a ring signature scheme such that the endorsement set itself is not revealed, but only that sufficiently many endorsement signatures were obtained. Another approach to obtain privacy-preserving endorsements was described in [ADCNS19], leveraging Idemix credentials to gain endorser-privacy, and as such, inherits the limitations (notably leaking the endorser’s organization) that come with Idemix.

Attribute-based credentials. Early anonymous credential schemes were built from blind signatures, whereby a user obtained a blind signature from an issuer on the user’s commitment to its attributes. When the user later authenticates, they provide the signature, the shown attributes, and a proof of knowledge of all unshown attributes. These schemes are limited as they can only be shown once. Subsequent work like the one underlying Idemix [CL04] allowed for an arbitrary number of unlinkable showings. A user obtains a signature on a commitment on attributes, randomizes the signature, and proves in zero-knowledge that the randomized signature corresponds to the shown and unshown attributes.

Recent work from [FHS19] circumvented inefficiencies in the above ideas by coining two new primitives: set-commitment schemes, and SPS-EQ. As a result, authors obtained a scheme allowing to randomize both the signature and the commitment on a set of attributes. Furthermore, a subset-opening of the set-commitment yielded constant-size selective showing of attributes.

New work from [CLPK22] extended [FHS19], improving the expressivity, efficiency trade-offs and introducing the notion of *signer-hiding* (also known as *issuer-hiding* [BEK⁺21]) to allow users to easily randomize the public key used to generate a signature to hide the identity of credential issuers. Authors achieve the previous points using a *Set-Commitment scheme supporting Disjoint Sets* (SCDS) and mercurial signatures [CL19]. The latter primitive extends SPS-EQ to consider equivalence classes not only on the message space but also on the key space.

We build on top of the above-mentioned works but unlike [CLPK22], we work with the generic group model [Sho97] as our main motivation is the proposal of efficient alternatives. For this reason, we use the mercurial signature scheme from [CL19].

2 Cryptographic Background

Below, we walk through the different building blocks mentioning how and why these components yield greater functionality and efficiency for a credential system in the permissioned blockchain setting like Fabric. Subsequently, we introduce the necessary notation and syntax for the following sections.

SCDS. Using commitment schemes that allow to commit to *sets* of attributes enables constant-size openings of subsets (selective disclosure) of the committed sets. These schemes support commitment randomization without the need to

rely on zero-knowledge proofs of correct randomization, as the corresponding witness for openings can be adapted accordingly with respect to the randomization of the committed set. The set-commitment scheme presented in [CLPK22] extends [FHS19] to support openings of attribute sets *disjoint* from the committed set. This is particularly useful in the permissioned blockchain setting, *e.g.*, to model access control policies. Furthermore, the scheme from [CLPK22] also supports the use of proof of exponentiations to outsource some of the computational cost from the verifier to the prover. In the case of Fabric, this is a particularly interesting feature to make the endorser’s verification faster when validating a transaction proposal.

Mercurial Signatures. The introduction of SPS-EQ in [FHS19] allowed to adapt a signature on a representative message to a signature on a different representative (in a given equivalence class) without knowledge of the secret key. If the adapted signature is indistinguishable from a fresh signature on a random message, the scheme satisfies the notion of *perfect adaption*. This, together with the randomizability of the set-commitment scheme, allows to consistently and efficiently update the signature of a credential, bypassing the need to generate and keep account of pseudonyms and NIZK proofs that are required in all previous works based on Idemix. Using mercurial signatures as in [CLPK22] allows to easily randomize the corresponding public keys while consistently adapting the signatures.

Issuer-hiding In [CLPK22], since users can consistently randomize the signature on their credential and the issuer’s public key (as previously mentioned), a fully adaptive NIZK argument is used to prove that a randomized issuer key belongs to the equivalence class of one of the keys contained in a list of issuers keys. This way, the randomized issuer key can be used to verify the credential while hiding the issuer’s identity (like in a ring signature). In permissioned blockchains where there are multiple organizations that issue credentials, such a NIZK allows users holding valid signatures to pick any subset of issuer’s public keys to generate a proof. Another approach following the work from [BEK⁺21] (briefly discussed in [CLPK22]) is to consider issuer-policies. An issuer-policy is a set $\{(\sigma_i, \text{opk}_i)_{i \in [n]}\}$ of signatures on issuer’s public keys generated by some verification secret key vsk . To hide the identity of an issuer j , a user consistently randomizes the pair (σ_j, opk_j) to obtain a randomized public key opk'_j . It then adapts the signature σ on its credential the same way, and presents opk'_j to the verifier. If the verifier accepts the signature σ_j on opk'_j (using vpk), it proceeds to verify σ using opk'_j . Issuer-policies can be specified by the entity that created the smart contract and defined within using the entity’s verification key pair. Unlike the first approach where users choose the issuer’s anonymity set, here it is determined by the policy maker.

2.1 Notation

Let BGGen be a p.p.t algorithm that on input 1^λ with λ the security parameter, returns a description $\text{BG}=(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, P_1, P_2, e)$ of an asymmetric (Type-3)

bilinear group where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are cyclic groups of prime order p with $\lceil \log_2 p \rceil = \lambda$, P_1 and P_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 , and $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is an efficiently computable (non-degenerate) bilinear map. For all $a \in \mathbb{Z}_p$, $[a]_s = aP_s \in \mathbb{G}_s$ denotes the implicit representation of a in \mathbb{G}_s for $s \in \{1, 2\}$. For vectors \mathbf{a}, \mathbf{b} we extend the pairing notation to $e([\mathbf{a}]_1, [\mathbf{b}]_2) := [\mathbf{ab}]_T \in \mathbb{G}_T$. $r \stackrel{\$}{\leftarrow} \mathcal{S}$ denotes sampling r from set \mathcal{S} uniformly at random. $\mathcal{A}(x; y)$ indicates that y is passed directly to \mathcal{A} on input x . Hash functions are denoted by \mathcal{H} .

2.2 Set-Commitment scheme supporting Disjoint Sets [CLPK22]

Syntax. $\text{Setup}(1^\lambda, 1^q)$ takes as input a security parameter λ and an upper bound q for the cardinality of committed sets. It outputs public parameters pp , discarding the trapdoor key s . $\text{TSetup}(1^\lambda, 1^q)$ is like Setup but returns the trapdoor key. $\text{Commit}(\text{pp}, \mathcal{X})$ takes as input pp and a set \mathcal{X} with $1 \leq |\mathcal{X}| \leq q$. It outputs a commitment C on \mathcal{X} and opening information O . $\text{Open}(\text{pp}, C, \mathcal{X}, O)$ takes as input pp , a commitment C , a set \mathcal{X} , and opening information O . It outputs 1 if and only if O is a valid opening of C on \mathcal{X} . $\text{OpenSS}(\text{pp}, C, \mathcal{X}, O, \mathcal{S})$ takes as input pp , a commitment C , a set \mathcal{X} , opening information O , and a non-empty set \mathcal{S} . If \mathcal{S} is a subset of \mathcal{X} committed to in C , it outputs a witness wss that attests to it. Otherwise, outputs \perp . $\text{OpenDS}(\text{pp}, C, \mathcal{X}, O, \mathcal{D})$ takes as input pp , a commitment C , a set \mathcal{X} , opening information O , and a non-empty set \mathcal{D} . If \mathcal{D} is disjoint from \mathcal{X} committed to in C , it outputs a witness wds that attests to it. Otherwise, outputs \perp . $\text{VerifySS}(\text{pp}, C, \mathcal{S}, \text{wss})$ takes as input pp , a commitment C , a non-empty set \mathcal{S} , and a witness wss . If wss is a valid witness for \mathcal{S} a subset of the set committed to in C , it outputs 1 and otherwise \perp . $\text{VerifyDS}(\text{pp}, C, \mathcal{D}, \text{wds})$ takes as input pp , a commitment C , a non-empty set \mathcal{D} , and a witness wds . If wds is a valid witness for \mathcal{D} being disjoint from the set committed to in C , it outputs 1 and otherwise \perp . $\text{PoE}(\text{pp}, \mathcal{X}, \alpha)$ takes as input pp , a non-empty set \mathcal{X} , and a randomly-chosen value α . It computes a proof of exponentiation with respect to \mathcal{X} and outputs a proof π_Q and a witness Q .

Security Properties. Correctness requires that (1) for a set \mathcal{X} , $\text{Open}(\text{Commit}(\text{pp}, \mathcal{X}), \mathcal{X}) = 1$, (2) for \mathcal{S} a subset of \mathcal{X} , $\text{VerifySS}(\mathcal{S}, \text{OpenSS}(\text{Commit}(\text{pp}, \mathcal{X}))) = 1$, and (3) for all possible sets \mathcal{D} disjoint from \mathcal{X} , $\text{VerifyDS}(\mathcal{D}, \text{OpenDS}(\text{Commit}(\text{pp}, \mathcal{X}))) = 1$. The scheme should also be (1) *binding* whereby each commitment overwhelmingly pertains to one particular set of attributes, (2) *hiding* whereby an adversary, given access to opening oracles, should not be able to distinguish which of two sets a commitment was generated on, and (3) *sound* in that sets which are not subsets of the committed set do not verify under VerifySS , and sets that are not disjoint from the committed set do not verify under VerifyDS .

2.3 Structure-Preserving Signatures on Equivalence Classes [CLPK22]

Syntax. $\text{ParGen}(1^\lambda)$ takes as input a security parameter λ and returns public parameters pp with an asymmetric bilinear group BG . $\text{KGen}(\text{pp}, \ell)$ takes as input

pp , a message-length ℓ , and outputs a key pair (sk, pk) . $\text{Sign}(\text{sk}, m)$ takes as input sk and message m . It outputs a signature σ on m where $m \in (\mathbb{G}_i^*)^\ell$ is a representative for a class $[m]_{\mathcal{R}}$. $\text{ChgRep}(m, \sigma, \mu, \text{pk})$ takes as input $m, \sigma, \mu, \text{pk}$. It computes an updated signature σ' on new representative $m^* = \mu m$ and returns (m^*, σ') . $\text{Verify}(m, \sigma, \text{pk})$ takes m, σ and pk , and outputs 1 iff σ is valid for m . For mercurial signatures, the algorithms $\text{ConvertPK}(\text{pk}, \rho)$ and $\text{ConvertSK}(\text{sk}, \rho)$ are included to compute new representatives for public and secret keys, while ChgRep is extended to adapt signatures with respect to new key representatives.

Security Properties. SPS-EQ should be correct, existentially unforgeable against chosen-message attacks and have perfect adaption (in the vein of [CLPK22]).

3 Our ABC Model

We can rely on the random oracle model and apply the Fiat-Shamir transform to the ABC scheme from [CLPK22] (the showing protocol is a three move public coin one). However, in the previous ABC, interaction is required in the showing protocol to provide freshness (*i.e.*, to avoid replay attacks). To overcome this issue, we require the user to send the transaction proposal during the first move. Thus, applying the Fiat-Shamir transform to the first move bounds the credential showing to that particular transaction so that it cannot be replayed. Security is defined following the usual properties from [FHS19, CLPK22, DHS15]. In addition, we also consider the issuer-hiding notion from [CLPK22] and introduce a new one for auditability. However, we do not consider replay-attacks as in the previous models since they can be trivially detected for the same transaction.

ABC Syntax. An ABC consists of the following p.p.t algorithms:

$\text{Setup}(1^\lambda, \text{aux})$ takes a security parameter λ and some optional auxiliary information aux (which may fix an universe of attributes, attribute values and other parameters) and outputs public parameters pp discarding any trapdoor.

$\text{TSetup}(1^\lambda, \text{aux})$ like Setup but returns a trapdoor.

$\text{OKGen}(\text{pp})$ takes pp and outputs an organization key pair (osk, opk) .

$\text{UKGen}(\text{pp})$ takes pp and outputs a user key pair (usk, upk) .

$\text{AAKGen}(\text{pp})$ takes pp and outputs an auditor key pair (ask, apk) .

$\text{RAKGen}(\text{pp})$ takes pp and outputs a revocation key pair (rsk, rpk) .

$\text{Obtain}(\text{pp}, \text{usk}, \text{opk}, \text{apk}, \mathcal{X}, \text{nym})$ and $\text{Issue}(\text{pp}, \text{upk}, \text{osk}, \text{apk}, \mathcal{X}, \text{nym})$ are run by a user and the organization respectively, who interact during execution. Obtain takes pp , the user's secret key usk , an organization's public key opk , an auditor's public key apk , an attribute set \mathcal{X} of size $|\mathcal{X}| < q$, and a pseudonym nym used for revocation. Issue takes pp , a public key upk , a secret key osk , an auditor's public key apk , an attribute set \mathcal{X} of size $|\mathcal{X}| < q$, and a pseudonym nym . At the end of this protocol, Obtain outputs a credential cred on \mathcal{X} for the user or \perp if the execution failed.

$\text{Show}(\text{pp}, \text{opk}, \text{upk}, \text{usk}, \text{cred}, \mathcal{X}, \mathcal{S}, \mathcal{D}, \text{aux})$ takes pp , a public key opk , a key pair (usk, upk) , a credential cred for the attribute set \mathcal{X} , potentially non-empty

sets $\mathcal{S} \subseteq \mathcal{X}$, $\mathcal{D} \not\subseteq \mathcal{X}$ representing attributes sets being a subset (\mathcal{S}) or disjoint (\mathcal{D}) to the attribute set (\mathcal{X}) committed in the credential, and auxiliary information aux . It outputs a proof π .

$\text{Verify}(\text{pp}, \text{opk}, \mathcal{X}, \mathcal{S}, \mathcal{D}, \pi, \text{aux})$ takes pp , the (potentially empty) sets \mathcal{S} and \mathcal{D} , a proof π and auxiliary information aux . It outputs 1 or 0 indicating whether the credential showing proof π was accepted or not.

$\text{RSetup}(\text{pp}, (\text{rsk}, \text{rpk}), \text{NYM}, \text{RNYM})$ takes pp , a revocation key pair (rsk, rpk) and two disjoint lists NYM and RNYM (holding valid and revoked pseudonyms).

It outputs auxiliary information aux_{rev} for the revocation authority and revocation information $\mathbb{R} = (\mathbb{R}_V, \mathbb{R}_S)$. \mathbb{R}_V is needed for verifying the revocation status and \mathbb{R}_S is a list holding the revocation information per nym .

$\text{Revoke}(\text{pp}, (\text{rsk}, \text{rpk}), \text{aux}_{\text{rev}}, \mathbb{R}, b)$ takes pp , (rsk, rpk) , aux_{rev} , \mathbb{R} and a bit b indicating revoked/unrevoked. It outputs information \mathbb{R}' and aux'_{rev} .

$\text{AuditEnc}(\text{upk}, \text{apk})$ takes upk and apk . It outputs an encryption enc of upk under apk and auxiliary information α .

$\text{AuditDec}(\text{enc}, \text{ask})$ takes enc and ask . It outputs a decryption of enc using ask .

$\text{AuditPrv}(\text{enc}, \alpha, \text{usk}, \text{apk})$ takes enc , α , usk , and apk . It generates a proof for enc being the encryption of upk under apk and outputs a proof π .

$\text{AuditVerify}(\text{apk}, \pi)$ takes apk and a proof π for the correct encryption of a user's public key under apk and outputs 1 if and only if the proof verifies.

To introduce the formal security model, we consider a single revocation, issuing and auditability authority. Extension to the multi-issuing and multi-auditing setting is straightforward as each key can be generated independently. For multiple revocation authorities, one needs to consider multiple revocation accumulators and thus adapt the scheme accordingly. Issuer-hiding and auditability properties are considered independently as extensions. Let us denote by Tx the universe of transactions tx represented as bitstrings. We also use the following auxiliary lists, sets and global variables in oracles and formal definitions. N represents the set of all pseudonyms nym while the sets NYM and RNYM represent the subsets of unrevoked and revoked pseudonyms respectively. Therefore, we have that $\text{NYM} \cap \text{RNYM} = \emptyset \wedge \text{NYM} \cup \text{RNYM} = \text{N}$. NYM , HU and CU are lists that keep track of which nym is assigned to which user, honest users and corrupt users, respectively. The global variables RI and NYM_{LoR} (initially set to \perp) store the revocation information $(\mathbb{R}_S, \mathbb{R}_V)$ and the pseudonyms used in \mathcal{O}_{LoR} respectively. The oracles are defined as follows:

$\mathcal{O}_{\text{HU}}(i)$ takes as input a user identity i . If $i \in \text{HU} \cup \text{CU}$, it returns \perp . Otherwise, it creates a new honest user i by running $(\text{USK}[i], \text{UPK}[i]) \stackrel{\$}{\leftarrow} \text{UKGen}(\text{opk})$, adding i to the honest user list HU and returning $\text{UPK}[i]$.

$\mathcal{O}_{\text{CU}}(i, \text{upk})$ takes as input a user identity i and (optionally) upk ; if user i does not exist, a new corrupt user with public key upk is registered, while if i is honest, its secret key and all credentials are leaked. If $i \in \text{CU}$, $i \in I_{\text{LoR}}$ (that is, i is a challenge user in the anonymity game) or if $\text{NYM}_{\text{LoR}} \cap \text{N}[i] \neq \emptyset$ then the oracle returns \perp . If $i \in \text{HU}$ then the oracle removes i from HU and adds it to CU ; it returns $\text{USK}[i]$ and $\text{CRED}[j]$ for all j with $\text{OWNR}[j] = i$. Otherwise (*i.e.*, $i \notin \text{HU} \cup \text{CU}$), it adds i to CU and sets $\text{UPK}[i] \leftarrow \text{upk}$.

- $\mathcal{O}_{\text{RN}}(\text{rsk}, \text{rpk}, \text{REV})$ takes as input the revocation secret key rsk , the revocation public key rpk and a list REV of pseudonyms to be revoked. If $\text{REV} \cap \text{RNYM} \neq \emptyset$ or $\text{REV} \not\subseteq \mathbf{N}$ return \perp . Otherwise, set $\text{RNYM} \leftarrow \text{RNYM} \cup \text{REV}$ and $\text{RI} \leftarrow \text{Revoke}(\text{pp}, (\text{rsk}, \text{rpk}), \text{RNYM}, \text{RI}, 1)$.
- $\mathcal{O}_{\text{Obtss}}(i, \mathcal{X})$ takes as input a user identity i , a pseudonym nym and a set of attributes \mathcal{X} . If $i \notin \text{HU}$ or $\exists j : \text{NYM}[j] = \text{nym}$, it returns \perp . Otherwise, it issues a credential to i by running $(\text{cred}, \top) \stackrel{\$}{\leftarrow} \text{Obtain}(\text{pp}, \text{USK}[i], \text{opk}, \text{apk}, \mathcal{X}, \text{nym})$, $\text{Issue}(\text{pp}, \text{UPK}[i], \text{osk}, \text{apk}, \mathcal{X}, \text{nym})$. If $\text{cred} = \perp$, it returns \perp . Else, it appends $(i, \text{cred}, \mathcal{X}, \text{nym})$ to $(\text{OWNR}, \text{CRED}, \text{ATTR}, \text{NYM})$ and returns \top .
- $\mathcal{O}_{\text{Obtain}}(i, \mathcal{X})$ lets the adversary \mathcal{A} , who impersonates a malicious organization, issue a credential to an honest user. It takes as input a user identity i , a pseudonym nym and a set of attributes \mathcal{X} . If $i \notin \text{HU}$, it returns \perp . Otherwise, it runs $(\text{cred}, \cdot) \stackrel{\$}{\leftarrow} \text{Obtain}(\text{pp}, \text{USK}[i], \text{opk}, \text{apk}, \mathcal{X}, \text{nym}), \cdot)$, where the Issue part is executed by \mathcal{A} . If $\text{cred} = \perp$, it returns \perp . Else, it appends $(i, \text{cred}, \mathcal{X}, \text{nym})$ to $(\text{OWNR}, \text{CRED}, \text{ATTR}, \text{NYM})$ and returns \top .
- $\mathcal{O}_{\text{Issue}}(i, \mathcal{X})$ lets the adversary \mathcal{A} , who impersonates a malicious user, obtain a credential from an honest organization. It takes as input a user identity i , a pseudonym nym and a set of attributes \mathcal{X} . If $i \notin \text{CU}$, it returns \perp . Otherwise, it runs $(\cdot, I) \stackrel{\$}{\leftarrow} (\cdot, \text{Issue}(\text{pp}, \text{UPK}[i], \text{osk}, \text{apk}, \mathcal{X}, \text{nym}))$, where the Obtain part is executed by \mathcal{A} . If $I = \perp$, it returns \perp . Else, it appends $(i, \perp, \mathcal{X}, \text{nym})$ to $(\text{OWNR}, \text{CRED}, \text{ATTR}, \text{NYM})$ and returns \top .
- $\mathcal{O}_{\text{Show}}(j, \mathcal{S}, \mathcal{D})$ lets the adversary \mathcal{A} play a dishonest verifier during a showing by an honest user. It takes as input an index of an issuance j and attributes sets \mathcal{S} and \mathcal{D} . Let $i \stackrel{\$}{\leftarrow} \text{OWNR}[j]$. If $i \notin \text{HU}$, it returns \perp . Otherwise, it runs $(\mathcal{S}, \cdot) \stackrel{\$}{\leftarrow} \text{Show}(\text{pp}, \text{USK}[i], \text{UPK}[i], \text{opk}, \text{ATTR}[j], \mathcal{S}, \mathcal{D}, \text{CRED}[j], \text{RI}, \text{apk}, \text{tx}), \cdot)$
- $\mathcal{O}_{\text{LoR}}(j_0, j_1, \mathcal{S}, \mathcal{D})$ is the challenge oracle in the anonymity game where \mathcal{A} runs Verify and must distinguish (multiple) showings of two credentials $\text{CRED}[j_0]$ and $\text{CRED}[j_1]$. The oracle takes two issuance indices j_0 and j_1 and attribute sets \mathcal{S} and \mathcal{D} . If $J_{\text{LoR}} \neq \emptyset$ and $J_{\text{LoR}} \neq \{j_0, j_1\}$, it returns \perp . Let $i_0 \stackrel{\$}{\leftarrow} \text{OWNR}[j_0]$ and $i_1 \stackrel{\$}{\leftarrow} \text{OWNR}[j_1]$. If $J_{\text{LoR}} \neq \emptyset$ then it sets $J_{\text{LoR}} \stackrel{\$}{\leftarrow} \{j_0, j_1\}$ and $I_{\text{LoR}} \stackrel{\$}{\leftarrow} \{i_0, i_1\}$. If $i_0, i_1 \neq \text{HU} \vee \text{N}[i_0] = \perp \vee \text{N}[i_1] = \perp \vee \text{N}[i_0] \in \text{RNYM} \vee \text{N}[i_1] \in \text{RNYM} \vee \mathcal{S} \not\subseteq \text{ATTR}[j_0] \cap \text{ATTR}[j_1] \vee \mathcal{D} \cap \{\text{ATTR}[j_0] \cup \text{ATTR}[j_1]\} \neq \emptyset$, it returns \perp . Else, it adds $\text{N}[i_b]$ to NYM_{LoR} and runs $(\mathcal{S}, \cdot) \stackrel{\$}{\leftarrow} (\text{Show}(\text{pp}, \text{USK}[j_b], \text{UPK}[j_b], \text{opk}, \text{ATTR}[j_b], \mathcal{S}, \mathcal{D}, \text{CRED}[j_b], \text{RI}, \text{apk}, \text{tx}), \cdot)$ (with b set by the experiment)

Intuitively, *correctness* requires that a credential showing with respect to a non-empty sets \mathcal{S} and \mathcal{D} of attributes always verifies if it was issued honestly on some attribute set \mathcal{X} with $\mathcal{S} \subset \mathcal{X}$ and $\mathcal{D} \subseteq \mathcal{X}$.

Correctness. An ABC system is correct if $\forall \lambda > 0, \forall q, q' > 0, \forall \mathcal{X} : 0 < |\mathcal{X}| \leq q, \forall \emptyset \neq \mathcal{S} \subset \mathcal{X}, \forall \emptyset \neq \mathcal{D} \subseteq \mathcal{X} : 0 < |\mathcal{D}| \leq q, \forall \text{NYM}, \text{RNYM} \subseteq \mathbf{N} : 0 < |\mathbf{N}| \leq q' \wedge \text{NYM} \cap \text{RNYM} = \emptyset, \forall \text{nym} \in \text{NYM}, \forall \text{nym}' \in \text{RNYM}$ it holds that:

$\text{pp} \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda, (1^q, 1^{q'})); (\text{rsk}, \text{rpk}) \stackrel{\$}{\leftarrow} \text{RAKGen}(\text{pp}); (\text{ask}, \text{apk}) \stackrel{\$}{\leftarrow} \text{AAKGen}(\text{pp}); (\mathbb{R}, \text{aux}_{\text{rev}}) \leftarrow \text{RSetup}(\text{pp}, \text{rpk}, \text{NYM}, \text{RNYM}); (\text{osk}, \text{opk}) \stackrel{\$}{\leftarrow} \text{OKGen}(\text{pp}); (\text{usk}, \text{upk}) \stackrel{\$}{\leftarrow} \text{UKGen}(\text{pp}); (\text{cred}, \top) \stackrel{\$}{\leftarrow} (\text{Obtain}(\text{pp}, \text{usk}, \text{opk}, \text{apk}, \mathcal{X}, \text{nym}), \text{Issue}(\text{pp}, \text{upk}, \text{osk},$

$\text{apk}, \mathcal{X}, \text{nym})$; $(\mathbb{R}_S, \mathbb{R}_V) \leftarrow \text{Revoke}(\text{pp}, \mathbb{R}, \text{aux}_{\text{rev}}, \text{nym}', 1)$; $\Omega \leftarrow \text{Show}(\text{pp}, \text{usk}, \text{upk}, \text{opk}, \text{cred}, \mathcal{S}, \mathcal{D}, \mathbb{R}, \text{apk}, \text{tx})$; $1 \leftarrow \text{Verify}(\text{pp}, \mathcal{S}, \mathcal{D}, \text{opk}, \mathbb{R}_V, \text{rpk}, \text{apk}, \text{tx}, \Omega)$

We now provide a formal definition for *unforgeability*. Given at least one non-empty set $\mathcal{S} \subset \mathcal{X}$ or $\mathcal{D} \not\subseteq \mathcal{X}$, a user in possession of a credential for the attribute set \mathcal{X} cannot perform a valid showing for $\mathcal{D} \subset \mathcal{X}$ nor for $\mathcal{S} \not\subseteq \mathcal{X}$. Moreover, revoked users cannot perform valid showings and no coalition of malicious users can combine their credentials and prove possession of a set of attributes which no single member has. This holds even after seeing showings of arbitrary credentials by honest users.

Unforgeability. An ABC system is unforgeable, if $\forall \lambda, q, q' > 0$ and p.p.t adversaries \mathcal{A} having oracle access to $\mathcal{O} := \{\mathcal{O}_{\text{HU}}, \mathcal{O}_{\text{CU}}, \mathcal{O}_{\text{RN}}, \mathcal{O}_{\text{ObtIss}}, \mathcal{O}_{\text{Issue}}, \mathcal{O}_{\text{Show}}\}$ the following probability is negligible.

$$\Pr \left[\begin{array}{l} \text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda, (1^q, 1^{q'})); (\text{rsk}, \text{rpk}) \xleftarrow{\$} \text{RAKGen}(\text{pp}); (\text{ask}, \text{apk}) \xleftarrow{\$} \text{AAKGen}(\text{pp}); \\ (\text{osk}, \text{opk}) \xleftarrow{\$} \text{OKGen}(\text{pp}); (\mathcal{S}, \mathcal{D}, \text{st}) \xleftarrow{\$} \mathcal{A}^{\mathcal{O}}(\text{pp}, \text{opk}, \text{rpk}, \text{apk}); \\ (\cdot, b^*) \xleftarrow{\$} (\mathcal{A}(\text{st}), \text{Verify}(\text{pp}, \mathcal{S}, \mathcal{D}, \text{opk}, \text{rpk}, \text{apk}, \text{RI}, \text{tx}, \Omega)) : \\ b^* = 1 \wedge \forall j : \text{OWNR}[j] \in \text{CU} \implies (\text{N}[j] = \perp \vee (\text{N}[j] \neq \perp \wedge (\mathcal{S} \not\subseteq \text{ATTR}[j] \vee \\ \mathcal{D} \subseteq \text{ATTR}[j] \vee \text{N}[j] \in \text{RNYM})) \end{array} \right]$$

For *anonymity*, during a showing, no verifier and no (malicious) organization (even if they collude) should be able to identify the user or learn anything about the user, except that she owns a valid credential for the shown attributes. Furthermore, different showings of the same credential are unlinkable.

Anonymity. An ABC system is anonymous, if $\forall \lambda, q, q' > 0$ and all p.p.t adversaries \mathcal{A} having oracle access to $\mathcal{O} := \{\mathcal{O}_{\text{HU}}, \mathcal{O}_{\text{CU}}, \mathcal{O}_{\text{RN}}, \mathcal{O}_{\text{Obtain}}, \mathcal{O}_{\text{Show}}, \mathcal{O}_{\text{LoR}}\}$ the following probability is negligible.

$$\Pr \left[\begin{array}{l} \text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda, (1^q, 1^{q'})); (\text{ask}, \text{apk}) \xleftarrow{\$} \text{AAKGen}(\text{pp}); : b^* = b \\ b \xleftarrow{\$} \{0, 1\}; (\text{opk}, \text{rpk}, \text{st}) \xleftarrow{\$} \mathcal{A}(\text{pp}); b^* \xleftarrow{\$} \mathcal{A}^{\mathcal{O}}(\text{st}) \end{array} \right] - \frac{1}{2}$$

Issuer-hiding states that no adversary (*i.e.*, a malicious verifier) cannot tell with high probability who is the issuer of a credential issued to an honest user.

Issuer-hiding. An ABC system supports issuer-hiding if for all $\lambda > 0$, all $q > 0$, all $n > 0$, all $t > 0$, all \mathcal{X} with $0 < |\mathcal{X}| \leq t$, all $\emptyset \neq \mathcal{S} \subset \mathcal{X}$ and $\emptyset \neq \mathcal{D} \not\subseteq \mathcal{X}$ with $0 < |\mathcal{D}| \leq t$, and p.p.t adversaries \mathcal{A} , the following holds.

$$\Pr \left[\begin{array}{l} \text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda, 1^q); \forall i \in [n] : (\text{osk}_i, \text{opk}_i) \xleftarrow{\$} \text{OKGen}(\text{pp}); \\ (\text{usk}, \text{upk}) \xleftarrow{\$} \text{UKGen}(\text{pp}); j \xleftarrow{\$} [n]; \\ (\text{cred}, \top) \xleftarrow{\$} (\text{Obtain}(\text{usk}, \text{opk}_j, \mathcal{X}), \text{Issue}(\text{upk}, \text{osk}_j, \mathcal{X})); \\ j^* \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_{\text{Show}}}(\text{pp}, \mathcal{S}, \mathcal{D}, (\text{opk}_i)_{i \in [n]}) \end{array} : j^* = j \right] \leq \frac{1}{n}$$

Finally, *auditability* requires that showings correctly encrypt users' keys.

Auditability. An ABC system is auditable if $\forall \lambda > 0, \forall q, q' > 0, \forall \mathcal{X} : 0 < |\mathcal{X}| \leq q, \forall \emptyset \neq \mathcal{S} \subset \mathcal{X}, \forall \emptyset \neq \mathcal{D} \not\subseteq \mathcal{X} : 0 < |\mathcal{D}| \leq q, \forall \text{NYM}, \text{RNYM} \subseteq \mathbb{N} : 0 < |\mathbb{N}| \leq q' \wedge \text{NYM} \cap \text{RNYM} = \emptyset, \forall \text{nym} \in \text{NYM}, \forall \text{nym}' \in \text{RNYM}$ it holds that:

$$\begin{aligned} & \text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda, (1^q, 1^{q'})); (\text{rsk}, \text{rpk}) \xleftarrow{\$} \text{RAKGen}(\text{pp}); (\text{ask}, \text{apk}) \xleftarrow{\$} \text{AAKGen}(\text{pp}); \\ & (\mathbb{R}, \text{aux}_{\text{rev}}) \leftarrow \text{RSetup}(\text{pp}, \text{rpk}, \text{NYM}, \text{RNYM}); (\text{osk}, \text{opk}) \xleftarrow{\$} \text{OKGen}(\text{pp}); (\text{usk}, \text{upk}) \\ & \xleftarrow{\$} \text{UKGen}(\text{pp}); (\text{cred}, \top) \xleftarrow{\$} (\text{Obtain}(\text{pp}, \text{usk}, \text{opk}, \text{apk}, \mathcal{X}, \text{nym}), \text{Issue}(\text{pp}, \text{upk}, \text{osk}, \\ & \text{apk}, \mathcal{X}, \text{nym})); (\mathbb{R}_S, \mathbb{R}_V) \leftarrow \text{Revoke}(\text{pp}, \mathbb{R}, \text{aux}_{\text{rev}}, \text{nym}', 1); (\text{enc}, \Omega) \leftarrow \text{Show}(\text{pp}, \text{usk}, \\ & \text{upk}, \text{opk}, \text{cred}, \mathcal{S}, \mathcal{D}, \mathbb{R}, \text{apk}, \text{tx}); 1 \leftarrow \text{Verify}(\text{pp}, \mathcal{S}, \mathcal{D}, \text{opk}, \mathbb{R}_V, \text{rpk}, \text{apk}, \text{tx}, \text{enc}, \Omega); \\ & \text{upk} \leftarrow \text{AuditDec}(\text{enc}, \text{ask}) \end{aligned}$$

4 Protego

We elaborate on the decisions that led to the design of our ABC scheme Protego. Subsequently, we discuss our construction and the integration with Fabric.

Revocation. We opt to integrate the work from [DHS15] as pointed out in [CLPK22]. The revocation system from [DHS15] defines a revocation authority responsible for managing an allow and deny list of revocation handlers. The authority publishes an accumulator RevAcc representing the deny list, and maintains a public list of non-membership witnesses for unrevoked users. During the issuing protocol, users are given a revocation handle that is encoded in the credential. To prove that they are not revoked during a showing, the user consistently randomizes its credential with the accumulator and the corresponding non-membership witness. Then the verifier checks that the (randomized) witness is valid for the revocation handle (encoded in the user credential), and with respect to the (randomized) accumulator. To work, the user must compute a Zero-Knowledge Proof of Knowledge (ZKPoK) on the correct randomization of the non-membership witness and the accumulator. As explained in [DHS15], the revocation handle encoded in the user's credential is of the form $\text{usk}_2(b + \text{nym})P_1$, where usk_2 is an additional user secret key required for anonymity and nym is the pseudonym used for revocation. For this reason, users are required to manage augmented keys of the form $\text{upk} = (\text{upk}_1, \text{upk}_2)$, $\text{usk} = (\text{usk}_1, \text{usk}_2)$. Furthermore, for technical reasons, another component usk_2Q , where Q is a random element in \mathbb{G}_1 with unknown discrete logarithm, must be included in the credential.

Auditability. A credential in [FHS19, CLPK22] and [DHS15] contains a tuple (C, rC, P_1) where C is the set commitment on the user attributes, r is a random value used for technical purposes and P_1 is used to compute a ZKPoK of the randomizer μ in $(\mu C, \mu rC, \mu P_1)$ during a showing. We borrow the idea of using a verifiable variant of ElGamal from [BDCET21] to prove the well-formedness of a ciphertext (encrypting the user's public key) with respect to the auditor's key. Therefore, we add the user's public key upk_1 and the auditor's public key apk as the sixth and seventh components to the credential. Thus, we now have revocable credentials of the form $(C, rC, P_1, \text{usk}_2(b + \text{nym})P_1, \text{usk}_2Q, \text{upk}_1, \text{apk})$, which can be randomized to obtain a tuple $(\mu C, \mu rC, \mu P_1, \mu \text{usk}_2(b + \text{nym})P_1,$

$\mu\text{usk}_2Q, \mu\text{usk}_1P_1, \mu\text{apk}$). We exploit this fact to allow the user to generate an *audit* proof that can be publicly verified without leaking information about the user’s public key. This way, verifiers can check a proof using the sixth and seventh component in the credential to be sure that (1) the user encrypted a public key for which it has the corresponding secret key, and (2) using the correct one. Since the issuing authority signs the credential, the randomization needs to be consistent. Modifications required to implement our auditability approach are as follows:

1. The user randomizes its credential as usual to obtain a new one of the form $(C'_1, C'_2, C'_3, C'_4, C'_5, C'_6, C'_7) = (\mu C_1, \mu C_2, \mu P_1, \mu C_4, \mu C_5, \mu \text{upk}_1, \mu \text{apk})$. Since only the user knows the randomizer μ , its public key remains hidden.
2. The user picks $\alpha \in \mathbb{Z}_p$ and encrypts its own public key using ElGamal encryption with auditor’s public key apk and randomness α to obtain a ciphertext $\text{enc} = (\text{enc}_1, \text{enc}_2) = (\text{upk}_1 + \alpha \text{apk}, \alpha P_1)$.
3. The user runs the algorithm `AuditPrv` (Figure 1) with input $(\text{enc}, \alpha, \text{usk}_1, \text{apk})$ to obtain c, z_1 and z_2 .
4. Then, the user picks $\beta \leftarrow \mathbb{Z}_p$, computes $t_1 = \beta P_2, t_2 = \beta \mu P_2, t_3 = \alpha \beta P_2$ and sends $(\text{enc}, c, z_1, z_2, t_1, t_2, t_3)$ to the verifier alongside the randomized credential from step 1.
5. The verifier checks the well-formedness of the ElGamal encryption pair running the algorithm `AuditVerify` (Figure 1) with input $(c, \text{enc}, z_1, z_2)$. If the check succeeds, it checks the following pairing equations to verify that the encrypted public key is the one in the credential:

$$e(\text{enc}_2, t_2) = e(C'_3, t_3) \wedge e(\text{enc}_2, t_1) = e(P_1, t_3) \wedge e(\text{enc}_1, t_2) = e(C'_6, t_1) + e(C'_7, t_3)$$

Observe that the verifier knows $\mu P_1 = C'_3, \mu \text{usk}_1 P_1 = C'_6, \mu \text{ask} P_1 = C'_7, (\text{usk}_1 + \alpha \text{ask}) P_1 = \text{enc}_1, \alpha P_1 = \text{enc}_2, \beta P_2 = t_1, \beta \mu P_2 = t_2$ and $\alpha \beta P_2 = t_3$. With β the user is able to randomize the other values so that the pairing equation can be checked to verify the relation between the ElGamal ciphertext and the randomized public key in C'_6 , without leaking information about the user’s public key. Furthermore, the first two pairing equations verify the well-formedness of t_1, t_2 and t_3 with respect to the user’s credential and the ciphertext. Hence, the verifier will not be able to recover the user’s public key nor the user cheat.

The proposed solution only adds two elements to the credential, while requiring the user to send two more elements in \mathbb{G}_1 , three in \mathbb{Z}_p and three in \mathbb{G}_2 , for a total of eight. Computational cost remains low as it just involves the computation of seven pairings, the ElGamal encryption and two Schnorr proofs [Sch90].

Issuer-hiding. We incorporate the issuer-hiding approaches discussed in Section 2. The work from [CLPK22] relies on a NIZK argument to prove that a randomized public key belongs to the equivalence class of one of the public keys contained in a list of issuer keys. We adapt the proof system to the signature used in this work, and extend it to make it *malleable* (see Appendix A) so that users can compute the proof once and then adapt it during showings with little computational cost (instead of having to compute it from scratch). This efficiency improvement is very useful in permissioned blockchains as the set of authorities tends to stay the same over time. For both approaches we observe that

the mercurial signature used in this work only provides a weak form of issuer-hiding. Given a signature that has been adapted to verify under a randomized public key \mathbf{pk}' in the equivalence class of \mathbf{pk} , the owner of \mathbf{pk} can recognize it. Thus, issuers can know which transactions belong to users from their organizations (but not to which particular user) and which ones don't by reading the non-interactive showing proof (it contains the issuer's randomized public key). However, we argue that in the permissioned blockchain setting this provides a fair trade-off as a minimum traceability level is important for compliance and auditability purposes.

Our construction. Compared to [CLPK22], we make use of a hash function to apply the (strong) Fiat-Shamir transform while adding the previously discussed auditability and revocation features. Therefore we implement the ZKPoK's as Schnorr proofs (unlike [CLPK22] which followed Remark 1 from [FHS19]).

In Figure 1 we present the setup, key generation, revocation and auditing algorithms. The setup algorithm also takes a bound q' on the maximum number of revoked pseudonyms for the revocation accumulator. The revocation authority is responsible for running the Revoke algorithm and updating the accumulator.

Obtain and Issue have constant-size communication and are given in Figure 2. For Show and Verify we present *Protego* and *Protego Duo*, depending on the issuer-hiding approach. *Protego* is given in Figure 3 and produces a variable-length proof as it relies on the (malleable) NIZK proof. *Protego Duo* produces a constant-size proof and is depicted in Figure 4. The differences are highlighted with grey. For both, after the credential is updated, the user randomizes the revocation accumulator, witnesses, and generates the Schnorr proofs. Following the auditing proof, the Fiat-Shamir transform is applied, the ZKPoK's and PoE's are computed, returning the showing proof. *Verify* takes a proof (depending on the case), computes the challenge and verifies each of the statements.

Integration with Fabric. A multi-party computation ceremony can be run for the CRS generation of the Setup algorithm to ensure that no organization knows the trapdoors of the different components. As we are in the permissioned setting it is plausible to assume that at least one of the organizations is honest. By allowing *users and endorsers* to obtain credentials, both can produce showing proofs. Users can generate showing proofs to prove that they satisfy the access policy for the execution of a particular transaction proposal. Furthermore, by computing the PoE's, the verification time for endorsers improves substantially. Similarly, endorsers can prove that they satisfy a given endorsement policy attaching a showing proof to their endorsements. Even if the endorsement policies are defined in a privacy-preserving way as suggested in [ADCNS19], endorsers can still compute selective AND and NAND clauses for the respective pseudonyms defined by the policy using their credentials. Endorsers should also use the read and write sets to from the transaction proposals to generate showing proofs.

Security Proofs. We present the main theorems and proofs for *Protego* (which are analogous for *Protego Duo*). Correctness and issuer-hiding follow from [CLPK22].

Setup($1^\lambda, \text{aux}$):
 $(q, q') \leftarrow \text{aux}$; **pick** $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$; $Q \xleftarrow{\$} \mathbb{G}_1$; $(\text{rev}_{\text{pp}}, \text{rev}_{\text{td}}) \xleftarrow{\$} \text{RevAcc.Setup}(1^\lambda, q')$
 $(\text{BG}, \text{scds}_{\text{pp}}, \text{scds}_{\text{td}}) \xleftarrow{\$} \text{SCDS.Setup}(1^\lambda, q)$; $(\text{sps}_{\text{pp}}, \text{sps}_{\text{td}}) \xleftarrow{\$} \text{SPS-EQ.ParGen}(1^\lambda; \text{BG})$
return $(\mathcal{H}, \text{BG}, \text{rev}_{\text{pp}}, Q, \text{scds}_{\text{pp}}, \text{sps}_{\text{pp}})$
TSetup($1^\lambda, \text{aux}$):
 $(q, q') \leftarrow \text{aux}$; **pick** $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$; $Q \xleftarrow{\$} \mathbb{G}_1$; $(\text{rev}_{\text{pp}}, \text{rev}_{\text{td}}) \xleftarrow{\$} \text{RevAcc.Setup}(1^\lambda, q')$
 $(\text{BG}, \text{scds}_{\text{pp}}, \text{scds}_{\text{td}}) \xleftarrow{\$} \text{SCDS.Setup}(1^\lambda, q)$; $(\text{sps}_{\text{pp}}, \text{sps}_{\text{td}}) \xleftarrow{\$} \text{SPS-EQ.ParGen}(1^\lambda; \text{BG})$
 $\text{td} = (\text{rev}_{\text{td}}, \text{scds}_{\text{td}}, \text{sps}_{\text{td}})$; **return** $(\mathcal{H}, \text{BG}, \text{rev}_{\text{pp}}, Q, \text{scds}_{\text{pp}}, \text{sps}_{\text{pp}}, \text{td})$
RevAcc.Setup($1^\lambda, 1^q$): $\text{BG} \xleftarrow{\$} \text{BGGen}(1^\lambda)$; $b \xleftarrow{\$} \mathbb{Z}_p^*$; **return** $(\text{BG}, (b^i P_1, b^i P_2)_{i \in [q]})$
AAKGen(pp): $\text{ask} \xleftarrow{\$} \mathbb{Z}_p^*$; $\text{apk} \leftarrow \text{ask} P_1$; **return** (ask, apk)
RAKGen(pp): $\text{rsk} \xleftarrow{\$} \mathbb{Z}_p^*$; $\text{rpk} \leftarrow \text{rsk} P_2$; **return** (rpk, rsk)
OKGen(pp): **return** $\text{SPS-EQ.KGen}(\text{BG}, \text{sps}_{\text{pp}}, 7)$
UKGen(pp): $\text{usk}_1, \text{usk}_2 \xleftarrow{\$} \mathbb{Z}_p^*$; $(\text{upk}_1, \text{upk}_2) \leftarrow (\text{usk}_1 P_1, \text{usk}_2 P_1)$
return $((\text{usk}_1, \text{usk}_2), (\text{upk}_1, \text{upk}_2))$
RSetup(pp, (rsk, rpk), NYM, RNYM):
 $(\Pi_{\text{rev}}, \text{aux}_{\text{rev}}) \leftarrow \text{RevAcc.Commit}(\text{rev}_{\text{pp}}, \text{RNYM})$
foreach $\text{nym} \in \text{NYM}$ **do** $\text{WIT}[\text{nym}] \leftarrow \text{RevAcc.NonMemWit}(\text{pp}, \Pi_{\text{rev}}, \text{aux}_{\text{rev}}, \text{nym})$
return $((\Pi_{\text{rev}}, \text{WIT}), \text{aux}_{\text{rev}})$
RevAcc.Commit(pp, \mathcal{X} ; rsk):
check $|\mathcal{X}| \leq q \wedge \forall b' \in \mathcal{X} : b' P_1 = b P_1$; $\Pi_{\text{rev}} \leftarrow \text{rsk}^{-1} \cdot \text{Ch}_{\mathcal{X}}(s) P_1$; $\text{aux}_{\text{rev}} \leftarrow \mathcal{X}$
return $(\Pi_{\text{rev}}, \text{aux}_{\text{rev}})$
Revoke(pp, \mathbb{R} , aux_{rev} , nym, b):
parse $\mathbb{R} = (\Pi_{\text{rev}}, \text{WIT})$; **parse** $\text{aux}_{\text{rev}} = \text{RNYM}$
if $b = 1$
 $\text{NYM} \leftarrow \text{NYM} \setminus \{\text{nym}\}$; $\text{RNYM} \leftarrow \text{RNYM} \cup \{\text{nym}\}$
 $(\Pi'_{\text{rev}}, \text{aux}'_{\text{rev}}) \leftarrow \text{RevAcc.Add}(\text{pp}, \Pi_{\text{rev}}, \text{RNYM}, \text{nym})$
else
 $\text{NYM} \leftarrow \text{NYM} \cup \{\text{nym}\}$; $\text{RNYM} \leftarrow \text{RNYM} \setminus \{\text{nym}\}$
 $(\Pi'_{\text{rev}}, \text{aux}'_{\text{rev}}) \leftarrow \text{RevAcc.Del}(\text{pp}, \Pi_{\text{rev}}, \text{RNYM}, \text{nym})$
foreach $\text{nym}' \in \text{NYM}$ **do** $\text{WIT}[\text{nym}'] \leftarrow \text{RevAcc.NonMemWit}(\text{pp}, \Pi'_{\text{rev}}, \text{aux}'_{\text{rev}}, \text{nym}')$
return $((\Pi'_{\text{rev}}, \text{WIT}), \text{aux}'_{\text{rev}})$
RevAcc.Add(pp, rsk, Π_{rev} , aux_{rev} , nym):
parse $\text{aux}_{\text{rev}} = \mathcal{X}$; $\mathcal{X} \leftarrow \mathcal{X} \cup \{\text{nym}\}$; **return** $\text{RevAcc.Commit}(\text{pp}, \mathcal{X}; \text{rsk})$
RevAcc.Del(pp, rsk, Π_{rev} , aux_{rev} , nym):
parse $\text{aux}_{\text{rev}} = \mathcal{X}$; $\mathcal{X} \leftarrow \mathcal{X} \setminus \{\text{nym}\}$; **return** $\text{RevAcc.Commit}(\text{pp}, \mathcal{X}; \text{rsk})$
RevAcc.NonMemWit(pp, Π_{rev} , aux_{rev} , nym):
 $\mathcal{X} \leftarrow \text{aux}_{\text{rev}}$; **check** $\text{nym} \notin \mathcal{X}$; **Let** $q(X)$ and $d \in \mathbb{Z}_p^*$ s.t. $\text{Ch}_{\mathcal{X}}(X) = q(X)(X + \text{nym}) + d$
return $(q(b) P_2, d)$
RevAcc.VerifyWit(pp, Π_{rev} , nym, wss_{rev}):
 $(\text{wss}_{\text{rev}}^1, \text{wss}_{\text{rev}}^2) \leftarrow \text{wss}_{\text{rev}}$; **return** $e(\Pi_{\text{rev}}, \text{rpk}) = e((b + \text{nym}) P_1, \text{wss}_{\text{rev}}^1) e(\text{wss}_{\text{rev}}^2 P_1, P_2)$
AuditEnc(upk, apk): $\alpha \leftarrow \mathbb{Z}_p$; $\text{enc} \leftarrow (\text{upk} + \alpha \text{apk}, \alpha P_1)$; **return** (enc, α)
AuditDec(enc, ask): $(\text{enc}_1, \text{enc}_2) \leftarrow \text{enc}$; **return** $(\text{enc}_1 - \text{ask} \cdot \text{enc}_2)$
AuditPrv(enc, α , usk, apk):
 $r_1, r_2 \leftarrow \mathbb{Z}_p$; $\text{com}_1 \leftarrow r_1 P_1 + r_2 \text{apk}$; $\text{com}_2 \leftarrow r_2 P_1$; $c \leftarrow \mathcal{H}(\text{com}_1, \text{com}_2, \text{enc})$
 $z_1 \leftarrow r_1 + c \cdot \text{usk}$; $z_2 \leftarrow r_2 + c \cdot \alpha$; **return** (c, z_1, z_2)
AuditVerify(apk, c, enc, z_1, z_2):
 $\text{com}_1 \leftarrow z_1 P_1 + z_2 \text{apk} - c \text{enc}_1$; $\text{com}_2 \leftarrow z_2 P_1 - c \text{enc}_2$; $c' \leftarrow \mathcal{H}(\text{com}_1, \text{com}_2, \text{enc})$
return $c' = c$

Fig. 1: Protego: setup, key generation, revocation and auditing algorithms.

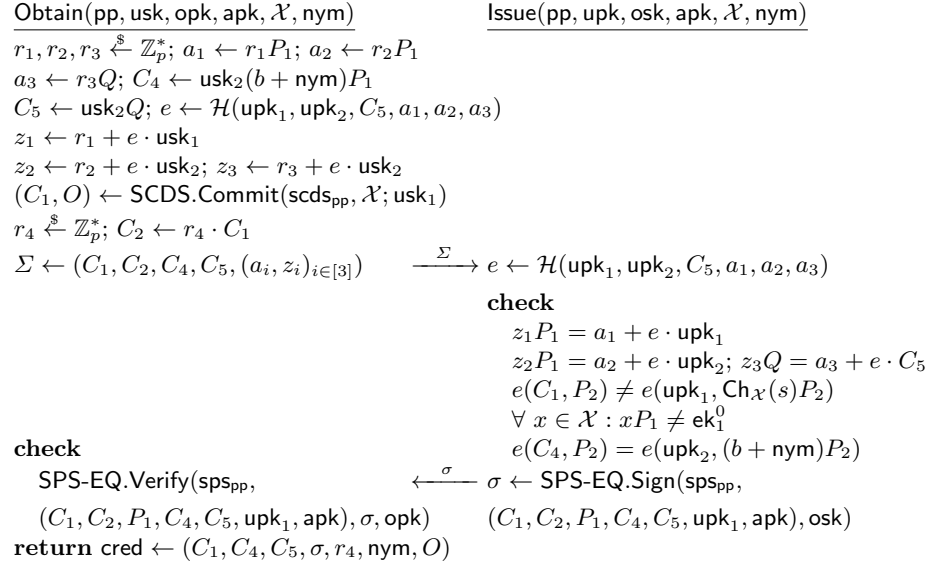


Fig. 2: Protego: obtain and issue algorithms.

Theorem 1. *If the q -co-DL assumption holds, the ZKPoK's have perfect ZK, SCDS is sound, SPS-EQ is EUF-CMA secure, and RevAcc is collision-free then Protego is unforgeable.*

Proof Sketch. The proof follows from [CLPK22] (Th. 6) and [DHS15] (Th. 3) whereby we assume there is an efficient adversary \mathcal{A} winning the unforgeability game with non-negligible probability. We use \mathcal{A} considering the following types of attacks:

- Type 1. Adversary \mathcal{A} conducts a valid showing so that $\text{nym}^* = \perp$. Then we construct an adversary \mathcal{B} that uses \mathcal{A} to break the EUF-CMA security.
- Type 2. Adversary \mathcal{A} manages to conduct a showing accepted by the verifier using the credential of user i^* under nym^* with respect to \mathcal{S}^* such that $\mathcal{S}^* \not\subseteq \text{ATTR}[\text{nym}]$ or with respect to \mathcal{D}^* such that $\mathcal{D}^* \subseteq \text{ATTR}[\text{nym}]$ holds. Then we construct an adversary \mathcal{B} that uses \mathcal{A} to break the soundness of the set-commitment scheme SCDS.
- Type 3. Adversary \mathcal{A} manages to conduct a showing accepted by the verifier reusing a showing based on the credential of a user i^* under nym^* with $i^* \in \text{HU}$, whose secret usk_{i^*} and credentials it does not know.
- Type 4. Adversary \mathcal{A} manages to conduct a showing accepted by the verifier using some credential corresponding to a revoked pseudonym $\text{nym}^* \in \text{RNYM}$. Then, we construct an adversary \mathcal{B} that uses \mathcal{A} to break the binding property of the revocation accumulator RevAcc.

Types 1 and 2 follow the proofs of [CLPK22] (Th. 6) as the underlying primitives remain unchanged. For Type 3, we leverage the fact that reusing a showing would only allow the adversary to generate a valid showing for *the same* original transaction tx (that is timestamped), and hence, we do not consider it


```

Show(pp, usk, upk, opkj, cred, S, D, (opki)i∈[n], (opkji, wji)i∈[2], Ω, ℝ, apk, tx)
(C1, C4, C5, σ, r, nym, O) ← cred; (Πrev, WIT) ← ℝ; β, μ, ρ, γ, τ, (ri)i∈[5]  $\xleftarrow{\$}$  ℤp*
if O = (1, (o1, o2)) then O' = (1, (μ · o1, o2)) else O' = μ · O
opk'j ← ConvertPK(opkj, ρwj1 + γwj2); Ω' ← SH.ZKEval(opk'j, opkj2, Ω; ρ, γ)
σ'  $\xleftarrow{\$}$  SPS-EQ.ChgRep(spspp, (C1, rC1, P1, C4, C5, upk1, apk), σ, μ, ρwj1 + γwj2, opkj)
cred' ← ((Ci)i∈[7] = μ · (C1, rC1, P1, C4, C5, upk1, apk), σ')
wss ← SCDS.OpenSS(scdspp, C1, S, O'); wds ← SCDS.OpenDS(scdspp, C1, D, O')
wssrev ← WIT[nym]; wss'rev ← (τwssrev1, usk2μτwssrev2 P1)
a1 ← r1C1; a2 ← r2P1; a3 ← r3Πrev; a4 ← r4Q; a5 ← r5P1; Π'rev ← (usk2μτ)Πrev
(enc, α) ← AuditEnc(apk, upk1); t1 = βP2; t2 = βμP2; t3 = αβP2
Π ← AuditPrv(enc, α, usk, apk)
e ← ℋ(S, D, apk, tx, enc, Π, opk'j, (opki)i∈[n], Ω', (ai)i∈[5], (ti)i∈[3], cred', wss, wds, tx,
C2, C3, Π'rev, C5, wss'rev)
z1 ← r1 + e · r; z2 ← r2 + e · μ; z3 ← r3 + e · (usk2μτ); z4 ← r4 + e · (usk2μ)
z5 ← r5 + e · (usk2μτwssrev2)
π1 ← SCDS.PoE(scdspp, S, e); π2 ← SCDS.PoE(scdspp, D, e)
return (enc, (ti)i∈[3], opk', (opki)i∈[n], Ω', cred', wss, wds, wss'rev, Π'rev, Π, π1, π2, (ai, zi)i∈[5])
Verify(pp, S, D, Πrev, rpk, apk, tx, Ω)
(enc, (ti)i∈[3], opk', (opki)i∈[n], Ω', cred', wss, wds, wss'rev, Π'rev, Π, π1, π2, (ai, zi)i∈[5]) ← Ω
(C1, C2, C3, C4, C5, C6, C7, σ) ← cred'
e ← ℋ(S, D, apk, tx, enc, Π, opk', (opki)i∈[n], Ω', (ai)i∈[5], (ti)i∈[3], cred', wss, wds, tx,
C2, C3, Π'rev, C5, wss'rev)
check
z1C1 = a1 + eC2; z2P1 = a2 + eC3; z3Πrev = a3 + eΠ'rev; z4Q = a4 + eC5
z5P1 = a5 + ewss'rev; RevAcc.VerifyWit(Π'rev, C4, wss'rev); AuditVerify(enc, Π2)
e(enc1, t2) = e(C6, t1) + e(C7, t3); e(enc2, t2) = e(C3, t3); e(enc2, t1) = e(P1, t3)
SCDS.VerifySS(C1, S, wss; π1, e); SCDS.VerifyDS(C1, D, wds; π2, e)
SH.Verify((opki)i∈[n], opk', Ω'); SPS-EQ.Verify(cred', opk')

```

Fig. 3: Protego: show and verify algorithms.

as an attack. Observe that any modification done to the original tx will lead to a different challenge and thus the rest of the proofs (showing, revocation and auditing) will not pass. Finally, Type 4 follows from [DHS15] (Th. 3). \square

Theorem 2. *If the DDH assumption holds, the SPS-EQ perfectly adapts signatures, and \mathcal{H} is assumed to be a random oracle, then Protego is anonymous.*

Proof Sketch. The proof follows from [CLPK22] (Th. 7) and [DHS15] (Th. 4). However, we must also take into account the RO model and the addition of the auditing features. The extra credential components for the auditing are randomized during every credential showing like the rest of the components. Similarly, the user generates a new encryption of the auditor's public key with a fresh α , while a fresh β is used to randomize the values t_i . Since ElGamal encryption is IND-CPA secure and key-private [BBDP01], the ciphertexts produced by the user are indistinguishable and do not leak information about the user's public key nor the auditor's. \square

```

Show(pp, usk, upk, opkj, cred, S, D, opkj, σj, ℔, apk, tx)
(C1, C4, C5, σ, r, nym, O) ← cred; (Πrev, WIT) ← ℔; β, μ, ρ, γ, τ, (ri)i∈[5]  $\xleftarrow{\$}$  ℤp*
if O = (1, (o1, o2)) then O' = (1, (μ · o1, o2)) else O' = μ · O
opk'j ← ConvertPK(opkj, ρ); σ'j  $\xleftarrow{\$}$  SPS-EQ.ChgRep(spspp, opkj, σj, ρ)
σ'  $\xleftarrow{\$}$  SPS-EQ.ChgRep(spspp, (C1, rC1, P1, C4, C5, upk1, apk), σ, μ, ρ, opkj)
cred' ← ((Ci)i∈[7] = μ · (C1, rC1, P1, C4, C5, upk1, apk), σ')
wss ← SCDS.OpenSS(scdspp, C1, S, O'); wds ← SCDS.OpenDS(scdspp, C1, D, O')
wssrev ← WIT[nym]; wss'rev ← (τwssrev1, usk2μτwssrev2P1)
a1 ← r1C1; a2 ← r2P1; a3 ← r3Πrev; a4 ← r4Q; a5 ← r5P1; Π'rev ← (usk2μτ)Πrev
(enc, α) ← AuditEnc(apk, upk1); t1 = βP2; t2 = βμP2; t3 = αβP2
Π ← AuditPrv(enc, α, usk, apk)
e ← ℋ(S, D, apk, tx, enc, Π, opk'j, σ'j, (ai)i∈[5], (ti)i∈[3], cred', wss, wds, tx,
C2, C3, Π'rev, C5, wss'rev)
z1 ← r1 + e · r; z2 ← r2 + e · μ; z3 ← r3 + e · (usk2μτ); z4 ← r4 + e · (usk2μ)
z5 ← r5 + e · (usk2μτwssrev2)
π1 ← SCDS.PoE(scdspp, S, e); π2 ← SCDS.PoE(scdspp, D, e)
return (enc, (ti)i∈[3], opk'j, σ'j, cred', wss, wds, wss'rev, Π'rev, Π, π1, π2, (ai, zi)i∈[5])
Verify(pp, S, D, Πrev, rpk, apk, vpk, tx, Ω)
(enc, (ti)i∈[3], opk', σ', cred', wss, wds, wss'rev, Π'rev, Π, π1, π2, (ai, zi)i∈[5]) ← Ω
(C1, C2, C3, C4, C5, C6, C7, σ) ← cred'
e ← ℋ(S, D, apk, tx, enc, Π, opk', σ', (ai)i∈[5], (ti)i∈[3], cred', wss, wds, tx,
C2, C3, Π'rev, C5, wss'rev)
check
z1C1 = a1 + eC2; z2P1 = a2 + eC3; z3Πrev = a3 + eΠ'rev; z4Q = a4 + eC5
z5P1 = a5 + ewss'rev; RevAcc.VerifyWit(Π'rev, C4, wss'rev); AuditVerify(enc, Π2)
e(enc1, t2) = e(C6, t1) + e(C7, t3); e(enc2, t2) = e(C3, t3); e(enc2, t1) = e(P1, t3)
SCDS.VerifySS(C1, S, wss; π1, e); SCDS.VerifyDS(C1, D, wds; π2, e)
SPS-EQ.Verify(opk', vpk); SPS-EQ.Verify(cred', opk')

```

Fig. 4: Protego Duo: show and verify algorithms.

Theorem 3. *If the algorithms AuditPrv and AuditVerify are a NIZK proof system and the SPS-EQ is EUF-CMA secure then Protego is auditable.*

Proof. If the verification returns true, we have that $\exists (\text{enc}_1^*, \text{enc}_2^*) = ((\delta^* + \alpha^* \text{ask})P_1, \alpha^* P_1)$ for some δ^* and α^* chosen by the adversary. Moreover, because of the unforgeability of the signature scheme, the verification implies that $C_3 = \mu^* P_1$, $C_6 = \mu^* \text{usk}_1 P_1$ and $C_7 = \mu^* \text{ask} P_1$ for some μ^* chosen by the adversary. As a result, we can re-write the pairing equations for the audit proof as:

$$\begin{aligned}
e(\alpha^* P_1, t_2^*) &= e(\mu^* P_1, t_3^*) \\
e(\alpha^* P_1, t_1^*) &= e(P_1, t_3^*) \\
e((\delta^* + \alpha^* \text{ask})P_1, t_2^*) &= e(\mu^* \text{usk}_1 P_1, t_1^*) + e(\mu^* \text{ask} P_1, t_3^*)
\end{aligned}$$

where t_1^* , t_2^* and t_3^* are also chosen by the adversary. We show that $\delta^* = \text{usk}_1$, which implies that $\text{upk}_1 = \text{AuditDec}(\text{enc}, \text{ask})$. Looking at the first two equations in the target group, we have that $\alpha^* t_2^* = \mu^* t_3^*$ and $\alpha^* t_1^* = t_3^*$, concluding that $t_2^* = \mu^* t_1^*$. Replacing t_2^* and t_3^* in third one and simplifying we obtain $(\delta^* +$

Scheme	Revocation				Signature			Issuer-hiding NIZK		
	$n = 10$		$n = 100$		$\ell = 7$ (for Protego)			$n = 5$		
	Prove	Verify	Prove	Verify	Sign	Verify	ChgRep	Prove	Verify	ZKEval
[BDCET21]	28	64	28	64	23	59	N/A	N/A	N/A	N/A
Protego	7.7	4.2	77.4	4.2	3.4	11	8.8	103	118	59

Table 1: Running time for the different algorithms in milliseconds.

$\alpha^* \text{ask}) \mu^* t_1^* = \mu^* \text{usk}_1 t_1^* + \mu^* \text{ask} \alpha^* t_1^*$. Therefore, we have $\mu^* \delta^* t_1^* + \mu^* \alpha^* \text{ask} t_1^* = \mu^* \text{usk}_1 t_1^* + \mu^* \alpha^* \text{ask} t_1^*$, deducing that $\delta^* = \text{usk}_1$. \square

5 Evaluation

We implemented a prototype of Protego and Protego Duo (available in [Pro]), using Rust with the bls12-381 curve and the BLAKE3 hash function. Our signature implementation is based on the one from [Bur20] but using the bls12-381 curve instead of BN curves [BN06]. As a result, we obtain times up to 67% faster when compared to [Bur20]. To run the benchmarks a Macbook Air (Chip M2 & 16GB RAM) was used with no extra optimizations, using the nightly compiler, and the *Criterion* library. For all values, the standard deviation was at most 1 millisecond.

Issue and **Obtain** take roughly 20 ms each when issuing a credential for 10 attributes. Both scale linearly on the number of attributes. To evaluate showing and verification, we considered the PoE in the showing protocol. Therefore, verification running time remains (almost) constant⁵ regardless the number of shown attributes, credential size, and issuer-hiding approach. If the PoE is disabled, showing running time would be smaller while verification would increase linearly with the number of shown attributes. An auditing proof in Protego takes roughly 1 and 1.5 ms for proof generation and verification, surpassing the values from [BDCET21]. In Table 1 we report the revocation and signing algorithms, including our issuer-hiding NIZK with $n = 5$. For Protego, we consider a signature for vectors of length seven (the size of a credential). In our case, the revocation witnesses are computed by the authority (in linear time) and then randomized by the users (in constant time). For this reason we consider the generation of a single witness for a revocation lists of 10 and a hundred elements (although in practice one would expect it to be closer to 10). For [BDCET21], we consider the total time to generate and verify a signature in a user level $L = 2$ (involving two delegations), with revocation times in \mathbb{G}_2 .

Comparison with the Idemix extension from [BDCET21]. The computational cost for the prover and verifier grows linearly with the number of attributes in the credential and delegation levels for [BDCET21]. In Protego Duo, the prover computational cost is $O(n - k)$ for showings involving k -attributes out of n , which in practice is much better. Verification cost in Protego and Protego Duo is almost constant (or $O(k)$ if the PoE is disabled). The two works are

⁵ Asymptotic complexity is $O(1)$ (considering exponentiations and pairings) but some multiplications depending on the shown attributes are required, hence the difference.




Fig. 5: From left to right, showing and verification times (in seconds) for the different schemes considering credential showings for 2, 4, 6 and 8 attributes.

Scheme	$k = 2$		$k = 4$		$k = 6$		$k = 8$		$k = 10$	
	Show	Verify	Show	Verify	Show	Verify	Show	Verify	Show	Verify
[BDCET21]	141	106	220	170	309	266	388	356	-	-
Protego	86	142	89	140	92	141	93	145	96	145
Protego Duo	29	35	32	35	34	36	37	36	39	38

Table 2: Protocols’ comparison showing the running times in milliseconds.

compared in Figure 5 using the same hardware (exact times are also given in Table 2). For [BDCET21], we consider a delegation level $L = 2$, which corresponds to a user level given that the root authority is at $L = 0$ and organizations start at $L = 1$. Regarding the attributes, [BDCET21] we could only retrieve information considering proofs for credential possession below ten attributes (assuming a minimal overhead when all attributes are shown as authors suggest). Therefore, we report credential possessions for [BDCET21] considering up to 8 attributes, and selective disclosures of k -out-of-10 attributes in ours. For Protego, we consider five authorities for the NIZK proof, which would suffice for practical scenarios like a consortium of pharmaceuticals.

6 Conclusions & Future Work

We presented here the first SPS-EQ credential scheme modified to work with permissioned blockchains. The versatility of Protego alongside the efficiency gains (at least twice as fast as the most recent Idemix extension), enables a broader scope of applications in such a setting. Depending on the context, the PoE’s can be computed or not, the credential issuer can be hidden or not, and one can select only subsets or disjoint sets to generate the proofs. Similarly, auditability and revocation features can be considered as optional, showing its flexibility.

As future directions to explore, we consider the following points: (1) adding confidentiality of transactions to a Protego-like credential scheme, (2) adding more power to the users (*i.e.*, how to define precise notions of user-invoked regulatory measures), and (3) extend our results to the multi-authority setting, where users can get attributes from multiple authorities instead of a single one.

A Our NIZK Argument for Issuer-hiding

We refer the reader to [CLPK22] (Section 3.1) for the basic syntax and security properties of malleable NIZK proof systems. In Figure 6 we build a fully adaptive

malleable NIZK argument following the construction from [CLPK22]. The main idea is that given two proofs π_1 and π_2 for statements $\mathbf{x}_1 = w_1 \mathbf{v}_i$ and $\mathbf{x}_2 = w_2 \mathbf{v}_i$, one can compute a valid proof π for the statement $\mathbf{x} = (\alpha w_1 + \beta w_2) \mathbf{v}_i$ with fresh α and β . The derivation privacy property of the proof system ensures that π looks like a freshly computed proof. Security follows from theorems 2 and 8 from [CLPK22].

$\text{SH.PGen}(1^\lambda):$ $\text{BG} \xleftarrow{\$} \text{BGGen}(1^\lambda); z \xleftarrow{\$} \mathbb{Z}_p$ $\text{return } (\text{BG}, [z]_1)$	$\text{SH.TPGen}(1^\lambda):$ $\text{BG} \xleftarrow{\$} \text{BGGen}(1^\lambda); z \xleftarrow{\$} \mathbb{Z}_p; \text{td} \leftarrow z$ $\text{return } (\text{BG}, [z]_1, \text{td})$
$\text{SH.PSim}(\text{crs}, \text{td}, (\mathbf{v}_i)_{i \in [n]}, [\mathbf{x}_1]_2, [\mathbf{x}_2]_2):$ $\delta, z_1, \dots, z_{n-1} \xleftarrow{\$} \mathbb{Z}_p^*$ $z_n \leftarrow \delta \text{td} - \sum_{i=1}^{n-1} z_i$ $\text{for all } i \in [n] \text{ do}$ $d_i \xleftarrow{\$} \mathbb{Z}_p; [\mathbf{a}_i]_2 \leftarrow d_i \cdot \mathbf{v}_i - z_i \cdot \mathbf{x}$ $\text{return } (([\mathbf{a}_n]_2, [d_n]_1, [z_n]_1)_{n \in [n]}, \delta P_2)$	$\text{SH.Prove}(\text{crs}, ([\mathbf{v}_i]_2)_{i \in [n]}, ([\mathbf{x}_j]_2, w_j)_{j \in [2]}):$ $// [\mathbf{x}_1]_2 = w_1 [\mathbf{v}_1]_2, [\mathbf{x}_2]_2 = w_2 [\mathbf{v}_1]_2$ $\delta, r_1, r_2, z_1, \dots, z_{n-1} \xleftarrow{\$} \mathbb{Z}_p^*$ $[z_n]_1 \leftarrow \delta [z]_1 - \sum_{i=1}^{n-1} [z_i]_1$ $([\mathbf{a}_i^j]_2, [d_i^j]_1) \leftarrow (r_j [\mathbf{v}_i]_2, w_j [z_i]_1 + [r_j]_1)$ $\text{for all } k \neq i \in [n], j \in [2] \text{ do}$ $d_k^j \xleftarrow{\$} \mathbb{Z}_p; [\mathbf{a}_k^j]_2 \leftarrow d_k^j [\mathbf{v}_k]_2 - z_k [\mathbf{x}_j]_2$ $\text{return } (([\mathbf{a}_n^j]_2, [d_n^j]_1, [z_n]_1)_{n \in [n]}^{j \in [2]}, \delta P_2)$
$\text{SH.ZKEval}(\text{crs}, [\mathbf{x}_1]_2, [\mathbf{x}_2]_2, \pi; \alpha, \beta):$ $// [\mathbf{x}']_2 = (\alpha w_1 + \beta w_2) [\mathbf{v}_1]_2$ $(([\mathbf{a}_n^j]_2, [d_n^j]_1, [z_n]_1)_{n \in [n]}^{j \in [2]}, Z_2) \leftarrow \pi$ $\delta \xleftarrow{\$} \mathbb{Z}_p^*; Z'_2 \leftarrow \delta Z_2$ $\text{for all } i \in [n] \text{ do}$ $[z'_i]_1 \leftarrow \delta [z_i]_1;$ $[d'_i]_2 \leftarrow \delta \alpha [d_i^1]_2 + \delta \beta [d_i^2]_2;$ $[\mathbf{a}'_i]_2 \leftarrow \delta \alpha [\mathbf{a}_i^1]_2 + \delta \beta [\mathbf{a}_i^2]_2;$ $\text{return } (([\mathbf{a}'_n]_2, [d'_n]_1, [z'_n]_1)_{n \in [n]}, Z'_2)$	$\text{SH.Verify}(\text{crs}, ([\mathbf{v}_i]_2)_{i \in [n]}, [\mathbf{x}]_2, \pi):$ $(([\mathbf{a}_n]_2, [d_n]_1, [z_n]_1)_{n \in [n]}, Z_2) \leftarrow \pi$ $\text{check } e([z]_1, Z_2) = e(\sum_{i=1}^n [z_i]_1, [1]_2)$ $\text{for all } i \in [n] \text{ do}$ $\text{check } e([d_i]_1, [\mathbf{v}_i]_2) = e([z_i]_1, [\mathbf{x}]_2) + e([1]_1, [\mathbf{a}_i]_2)$

Fig. 6: Our fully adaptive malleable NIZK argument

References

- ACC⁺20. Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhiyaoui, 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, AFT '20*, page 255–267, New York, NY, USA, 2020. Association for Computing Machinery.
- ADCNS19. Elli Androulaki, Angelo De Caro, Matthias Neugschwandtner, and Alessandro Sorniotti. Endorsement in hyperledger fabric. In *2019 IEEE International Conference on Blockchain (Blockchain)*, pages 510–519, 2019.
- BBDP01. Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 566–582. Springer, Heidelberg, December 2001.
- BDCET21. Dmytro Bogatov, Angelo De Caro, Kaoutar Elkhiyaoui, and Björn Tackmann. Anonymous transactions with revocation and auditing in hyperledger fabric. In Mauro Conti, Marc Stevens, and Stephan Krenn, editors, *Cryptology and Network Security*, pages 435–459, Cham, 2021. Springer International Publishing.
- BEK⁺21. Jan Bobolz, Fabian Eidens, Stephan Krenn, Sebastian Ramacher, and Kai Samelin. Issuer-Hiding Attribute-Based Credentials. In Mauro Conti, Marc Stevens, and Stephan Krenn, editors, *Cryptology and Network Security*, pages 158–178, Cham, 2021. Springer International Publishing.
- BN06. Paulo S. L. M. Barreto and Michael Naehrig. Pairing-friendly elliptic curves of prime order. In Bart Preneel and Stafford Tavares, editors, *SAC 2005*, volume 3897 of *LNCS*, pages 319–331. Springer, Heidelberg, August 2006.
- Bur20. Michael Burkhart. Mercurial signatures implementation. Github, 2020. <https://github.com/burkh4rt/Mercurial-Signatures>.
- CDD17. Jan Camenisch, Manu Drijvers, and Maria Dubovitskaya. Practical UC-secure delegatable credentials with attributes and their application to blockchain. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 683–699. ACM Press, October / November 2017.
- CL04. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Heidelberg, August 2004.
- CL19. Elizabeth C. Crites and Anna Lysyanskaya. Delegatable anonymous credentials from mercurial signatures. In Mitsuru Matsui, editor, *CT-RSA 2019*, volume 11405 of *LNCS*, pages 535–555. Springer, Heidelberg, March 2019.
- CLPK22. Aisling Connolly, Pascal Lafourcade, and Octavio Perez Kempner. Improved Constructions of Anonymous Credentials from Structure-Preserving Signatures on Equivalence Classes. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *Public-Key Cryptography – PKC 2022*, pages 409–438, Cham, 2022. Springer International Publishing.
- CV02. Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 21–30. ACM Press, November 2002.

- DHS15. David Derler, Christian Hanser, and Daniel Slamanig. A new approach to efficient revocable attribute-based anonymous credentials. In Jens Groth, editor, *15th IMA International Conference on Cryptography and Coding*, volume 9496 of *LNCS*, pages 57–74. Springer, Heidelberg, December 2015.
- FHS19. Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Structure-preserving signatures on equivalence classes and constant-size anonymous credentials. *Journal of Cryptology*, 32(2):498–546, April 2019.
- KDJL⁺19. Hui Kang, Ting Dai, Nerla Jean-Louis, Shu Tao, and Xiaohui Gu. Fabzk: Supporting privacy-preserving, auditable smart contracts in hyperledger fabric. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 543–555, 2019.
- MR19. Subhra Mazumdar and Sushmita Ruj. Design of anonymous endorsement system in hyperledger fabric. *IEEE Transactions on Emerging Topics in Computing*, pages 1–1, 2019.
- NVV18. Neha Narula, Willy Vasquez, and Madars Virza. Zkledger: Privacy-preserving auditing for distributed ledgers. In *Proceedings of the 15th USENIX Conference on Networked Systems Design and Implementation, NSDI’18*, page 65–80, USA, 2018. USENIX Association.
- Pro. Implementation. Available upon request.
- 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.
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997.
- Zur13. IBM Research Zurich. Specification of the identity mixer cryptographic library v2.3.0., 2013.