

zk-creds: Flexible Anonymous Credentials from zkSNARKs and Existing Identity Infrastructure

Michael Rosenberg*
University of Maryland
micro@cs.umd.edu

Jacob White†
Purdue University
white570@purdue.edu

Christina Garman
Purdue University
clg@cs.purdue.edu

Ian Miers
University of Maryland
imiers@cs.umd.edu

Abstract—Frequently, users on the web need to show that they are, for example, not a robot, old enough to access an age restricted video, or eligible to download an ebook from their local public library without being tracked. Anonymous credentials were developed to address these concerns. However, existing schemes do not handle the realities of deployment or the complexities of real-world identity. Instead, they implicitly make assumptions such as there being an issuing authority for anonymous credentials that, for real applications, requires the local department of motor vehicles to issue sophisticated cryptographic tokens to show users are over 18. In reality, there are multiple trust sources for a given identity attribute, their credentials have distinctively different formats, and many, if not all, issuers are unwilling to adopt new protocols.

We present and build *zk-creds*, a protocol that uses general-purpose zero-knowledge proofs to 1) remove the need for credential issuers to hold signing keys: credentials can be issued to a bulletin board instantiated as a transparency log, Byzantine system, or even a blockchain; 2) convert existing identity documents into anonymous credentials without modifying documents or coordinating with their issuing authority; 3) allow for flexible, composable, and complex identity statements over multiple credentials. Concretely, identity assertions using *zk-creds* take less than 150ms in a real-world scenario of using a passport to anonymously access age-restricted videos.

1. Introduction

Privacy-preserving identification is an apparent contradiction in terms: one cannot both wish to simultaneously identify themselves and stay private. But this is increasingly necessary on today's internet. For example, Australia, the EU, and the UK age-restrict access to some video content, requiring identification via a credit card or photo of an official ID to access it [37, 36]. The tracking and data exposure risks raised by such requirements can be eliminated with privacy-preserving cryptography: anonymous credentials allow a user to assert that they meet some access criteria, e.g., are over 18, without revealing anything else about themselves, linking their viewing habits to their identity, or even linking distinct video views together. Beyond this narrow application, anonymous credentials could be extended to complex

identity statements—for example, checking residency for accessing online library resources or petitioning local elected representatives¹—and the composition of credentials such as the pairing of a vaccine card with a photo ID.

While the subject of extensive academic work [26, 21, 22, 6, 5, 35, 17, 55], anonymous credentials have thus far seen little deployment.² In large part, this is because most existing systems are designed with a number of assumptions about identity that, while suitable for advancing a body of cryptographic knowledge, produce designs that can be difficult to actually deploy in real-world identity systems.

Existing anonymous credential schemes make, at a minimum, some subset of the following simplifying assumptions: there is a single issuer for a given identity property (e.g., date of birth); when there are multiple issuers for a property, the property formats are compatible; there exist reputable authorities that are able and (more importantly) willing to be responsible for holding signing keys, verifying identity properties, and running sophisticated cryptographic protocols for issuing anonymous credentials; all attribute formats needed for a credential are known in advance; and the set of authorities for a given identity attribute or credential can be enumerated at the time one instantiates the system.

In this paper, we build *zk-creds*, a flexible, issuer-agnostic anonymous credential toolkit for complex identity statements. The general-purpose proving functionality supported by zkSNARKs gives us the flexibility to address most of these challenges. From this approach, we automatically gain support for ad-hoc composition of credentials and access criteria, and issuance by (threshold) signatures. In many cases, we can even remove the need for trusted anonymous credential issuers entirely by instantiating a *bulletin board* to track issued credentials, no longer requiring issuers to hold signing keys or other secrets. Concretely, this work contributes the design and implementation of a toolkit for flexible privacy-preserving credentials, and builds two example applications.

1. New York State provides such a platform with no privacy guarantees [50].

2. Notable exceptions for human uses of credentials are limited trials of CL-sigs with Idemix [24] and, although it is not a full-fledged anonymous credential scheme, Privacy Pass [28]. Intel also makes use of DAA [13] for device attestation, and a MAC variant of anonymous credentials is also being used for private groups in Signal [25].

* ORCID: 0000-0001-9784-125X

† ORCID: 0000-0002-6850-2133

1.1. Past work and real-world limitations

Several approaches have tried to address the limitations of anonymous credentials, focusing primarily on the problems of finding and trusting issuers.

Distributing issuance via multiple issuers. To reduce the trust needed in issuers, schemes have explored threshold issuance [55] and support for multiple issuers [20]. While this improves the situation if there are multiple willing issuers, it does not address the potential scarcity of issuers who are willing or able to deploy novel (or any) cryptography. Nor does it provide a means to reconcile the differing identity document formats or use cases multiple issuers would have.

Decentralizing issuance by removing signing keys. In Decentralized Anonymous Credentials [35], credentials are maintained in some form of transparency log which can either be centralized and audited, distributed across cooperating parties, or operated in a decentralized fashion by a Byzantine system or blockchain. While this approach removes one obstacle to credential issuance by avoiding signing keys, the concrete protocol has performance and operational limitations. For example, the protocol requires that all clients have the full list of issued credentials, and does not address any of the other complexities of real use cases.

The messy reality of identity claims. We now return to our initial example: an anonymous credential to allow access to age-restricted videos and prevent tracking of browsing habits. In theory, whichever authority issues identity cards in a country can also issue anonymous digital credentials to everyone of age. But in practice, a number of problems arise when attempting to deploy such a scheme with existing anonymous credentials.

First, there is not a single source of identity documents (e.g., the US has 50+ drivers license issuers) and few might wish to participate due to the burden of deploying new technology. Fewer still can be trusted to secure the requisite signing keys for issuing credentials.

Second, requirements will change. What started as a token for being over 18 will need to support other age checks—under 12, over 21, over 65—necessitating more complex credentials, access criteria, and potentially credential revocation and reissuance.

Third, each ID issuer will, by default, form its own anonymity set. Even for “multi-authority” schemes designed to avoid this, differences in data fields can distinguish populations:³ a foreign diplomat accessing age-restricted content in their host country may be distinguished from a resident using a local ID.

Fourth, new identity documents need to be integrated as they emerge to avoid access equity issues, and these documents may have differing formats. For example, many cities now issue IDs in part for undocumented residents [41].

Finally, even for something as conceptually simple as “of age,” identity statements are not necessarily simple: in

3. Consider something as simple as date formats: Japanese Drivers Licenses give birth year relative to eras that correspond to the reign of the emperor. However, they use the Gregorian calendar for months and days.

the event age limits differ between jurisdictions, a video platform needs to check where the viewer is located, and IP geolocation may be insufficient (e.g., in the case of Tor or a VPN). Credentials can directly encode a home address but, even for physical credentials, this does not work in practice: people move and do not update their IDs, and as a result need to provide alternative proofs of address. Supporting this privately requires composing credentials for, e.g., age and residency.

Minimizing trust when issuing credentials. Current (even non-)anonymous credential protocols assume the same party verifies claimed identity properties and signs cryptographic credentials. This requires finding a single party who is trusted for two different (and not necessarily related) tasks: one who is both capable of verifying identity attributes and competent to manage signing keys. The linking of these two roles is often unnecessary and complicates deployment.

First, many uses of anonymous credentials do not use identity attributes which must be verified by a trusted party to issue a credential. Looking ahead, we describe an issuer-less Privacy Pass-like construction [28] where Sybil-resistant anonymous tokens are issued by making a blockchain payment. This has no trusted parties—neither for verifying that the user is not a Sybil, nor for signing a credential.

Second, even when we must trust some entity to verify identity attributes (e.g., a passport issuer for a user’s date of birth), it is not necessary to trust an additional party to hold key material for a novel cryptographic scheme. Looking ahead, we offer the minimal trust assumptions in many such cases by replacing signing with a distributed bulletin board.

Third, even where there is a trusted party who both verifies identity attributes and issues credentials, trusting a party to maintain a list is safer than trusting them to secure signing keys. In existing anonymous credential schemes, compromise of issuing keys is frequently undetectable and rollback requires rekeying and reissuing. With issuance via a list, compromise is detectable and easily reversible.

All of the aforementioned issues can be addressed by a scheme that is flexible, dynamically adaptable to new use cases post-deployment, minimizes the need to find new trusted parties, and can support complex access criteria that are agnostic to the issuer or credential format.

1.2. Our contribution

We introduce `zk-creds`, a toolkit for privacy-preserving authentication protocols and anonymous credentials that offers flexible identity assertions and does not need trusted issuers. A key contribution of `zk-creds` over previous works is the usage of general-purpose zero-knowledge proofs rather than bespoke proof systems over blind signatures.

The switch to general-purpose zero-knowledge proofs as the basis for anonymous credentials, instead of blind signatures, is a paradigm shift: rather than imagining a subset of use cases and designing custom protocols for each while balancing cryptographic tradeoffs, `zk-creds` gets full privacy and full expressivity even after a protocol

is designed and deployed. A single scheme, built with `zk-creds`, is adaptable to shifting requirements without requiring the development of new custom cryptographic protocols. Moreover, application-specific logic can be defined and modified in simple programming languages via a number of publicly available tools [49, 3, 10] with the instantiation of the scheme handled by the compiler.

General-purpose zero-knowledge proofs enable `zk-creds` to support flexible and composable access criteria. `zk-creds` not only allows users to privately show that their credential(s) meet some arbitrary access criteria check, but it also allows these criteria to be defined at any time (even after system setup or credential issuance), by any party, and composed dynamically as *gadgets*. This flexibility allows `zk-creds` to meet the reality of real-world authentication mechanisms: requirements can dynamically change at any time, as can use cases and even identity issuers.

The second major contribution of `zk-creds` is its support of existing government identity infrastructure without modification or collaboration. Using general-purpose zero-knowledge proofs, we can convert the digital (non-anonymous) identity information that is increasingly included in national identity cards and passports into anonymous credentials. This *zk-supporting-documentation* allows us to provide a digital analogue of walking into a US Department of Motor Vehicles and presenting existing identity documents to get a driver’s license, *without* exposing any information to the issuer. We implement an end-to-end example of this paradigm, using a zero-knowledge proof over the data in unmodified NFC-enabled US passports to create credentials for accessing age-restricted videos.

As a third contribution, `zk-creds` supports publicly verifiable credentials. Because the issued credential list can be maintained by a public system (i.e., a transparency log or blockchain) and each credential can include *zk-supporting-documentation* justifying its issuance, the set of all issued credentials is publicly auditable. This is not possible when credentials can be surreptitiously issued via signing keys.⁴ As such, we need only trust the issuer to add credentials (and their supporting documentation) to the list, and any compromise or malfeasance is detectable and reversible.

To summarize, in this paper we design, build, and benchmark `zk-creds` which:

- drastically improves performance over existing decentralized schemes via reusable proofs where ShowCred takes < 150ms;
- supports existing physical identity documents (e.g., passports) without modification via *zk-supporting-documentation*;
- provides support for flexible and composable gadgets that can be combined to express complex access criteria checks even after system setup;
- allows for public auditability of issued credentials, without harming anonymity;

4. We note, however, that while the credentials are fully auditable, identity statements require inherent trust in something. If the identity infrastructure, e.g., US passports, is not trusted, then we can make no guarantees.

- provides (of independent interest), blind Groth16, a novel mechanism for privately linking together multiple zero-knowledge proofs in a way that enables proof rerandomization and reuse; and
- includes a full application for age-restricted video access with cloning resistance, using existing passports for issuance.

Non-goals: On-chain verification of credentials. This work constructs flexible credentials that can be *issued* without a central party holding a signing key (although we also support signature-based issuance). This should not be confused with a different area of both industrial and academic work (see e.g., [52]), which considers *verification* of existing (i.e., centrally issued) anonymous credentials by a smart contract. The question for on-chain verification of anonymous credentials is not how to remove centralized issuance, but simply how to minimize the cost of verification given the extreme cost of smart contract execution. Reducing verification costs is typically done by batching verification inside a zero-knowledge proof⁵ and is generically applicable to any anonymous credential scheme, including the ones proposed here.

2. Overview

`zk-creds` is a system for issuing credentials to users and privately showing that a credential meets access criteria.

2.1. Example application and credential lifecycle

A *credential* is a commitment to a set of attributes (e.g., name, date of birth, etc.). A credential is *issued* (see Figure 1) when it is made a leaf in a Merkle tree. We call the set of leaves the *issuance list*. Optionally, protocol designers can require *zero-knowledge supporting-documentation* that the attributes match some (existing) document without revealing any additional information. In our implemented example (see Section 7), this consists of a zero-knowledge proof that the attributes in the credential match a US passport.

We emphasize that, while our example trusts an existing identity document issuer—the passport authority—it requires no additional trusted parties, no existing parties to take on additional trusted roles, or even modification of the issuer. With standard anonymous credentials, we would need to also trust the security of credential signing keys, likely held by an additional party. In our example, we need only some way of maintaining a list of issued credentials.

Clients, once issued a credential, *show* a credential to gain access to some resource, as shown in Figure 2. The client presents a non-interactive zero-knowledge proof that: 1) they have a credential (a commitment) in the Merkle tree of issued credentials and 2) the attributes meet some *access criteria*.

Crucially, the zero-knowledge proof hides which credential is used, the credential’s attributes, and the details of how the access criteria were met. In our example, the client uses a credential containing their birth date to show they are over 18 and gain access to a website. The verifier learns only that

5. This is, in essence, a specialized zk-rollup [16].




Figure 1: A credential in `zk-creds` is issued by adding it to a Merkle tree after (optionally) presenting `zk-supporting-documentation` to justify issuance.




Figure 2: To show a credential in `zk-creds`, the prover uses knowledge of their credential opening and the position of the credential in the issuance list to construct a zero-knowledge proof. The verifier need only know the issuance list root.

the client is over 18, not their identity, which credential they used, their exact birth date, or any of the other attributes in their credential.

2.2. Design features

As illustrated by our example application, our approach to building anonymous credentials has a number of important features.

Flexible access criteria. Application developers can define arbitrary access criteria at any point in the lifetime of the system. We support common features from the anonymous credential literature, such as hidden-attribute credentials, inequality or expiry checks, rate limiting, and cloning resistance where violating the rate limit (e.g., by sharing a credential with others) results in the credential’s identification and revocation. And because we support the efficient encoding of access criteria as an NP relation, we can easily support more complex criteria than existing schemes, such as a proof of residency in a given municipality to access ebooks from a local public library.

Auditable issuance. Credential issuance can be publicly auditable as well: e.g., in our passport example anyone can download the list of credentials and, with `zk-supporting-documentation`, verify both that a credential was issued and why. Even without such documentation, all issued credentials are visible and issuance can be investigated. In contrast, it is impossible to enumerate, let alone audit, every credential signed with a given key.

Flexible credential management. Because credential issuance is simply a matter of list management, credential issuance is flexible. In many cases, we need not find a trusted party at all: a simple bulletin board is sufficient, as

is a blockchain. In other cases, a central party can maintain the list without needing to be trusted to secure signing keys.

Signature-issued credentials. Separately, in cases where there is a party who is trusted to issue correct credentials without public auditability and is trusted, willing, and able to secure signing keys, our implementation of `zk-creds` also supports issuing credentials via signatures. In this case, there is no issuance list. When feasible, this leads to faster credential shows and removes the overhead of managing a witness to list membership.

Witness management. Using Merkle trees for credential issuance requires the user to maintain an up-to-date witness to the credential’s membership in an issuer’s list. Periodically, the user can ask the issuer for an updated witness. Looking ahead, this also lets any user update their witness by downloading logarithmic-sized updates from the tree’s *frontier* and a constant number of Merkle roots (see Appendix B).

Revocation. Many existing approaches require expensive asymmetric cryptographic operations for each revocation. Some schemes, like EPID [14], require each credential show to perform work linear in the number of revoked private keys. Other schemes use, e.g., RSA accumulators, which require recomputing accumulator witnesses per revocation at cost linear in the number of revocations. And, while more efficient accumulators exist [9], such techniques have not been used for revocation, to the best of our knowledge. In contrast, each revocation in `zk-creds` only requires removing the credential from its Merkle tree, incurring only logarithmic costs in the number of issued credentials. This captures revocation of the credential where the holder’s public identity is known, and so called *private key revocation* where a stolen or leaked credential is banned.

3. Preliminaries

3.1. General notation

We write $x := z$ to denote variable assignment, and $y \leftarrow S$ to denote sampling uniformly from a set S . $y := A(x; r)$ denotes the execution of a probabilistic algorithm A on input x , using randomness r . We write $\vec{x} := x_1, x_2, \dots$ to denote a variable-length list, and boldface to denote a vector. For an arbitrary, efficiently computable predicate P , we say that a *proof of knowledge of a relation* $R = \{(x; w) : P(x, w)\}$ with respect to an *instance* x is a proof of knowledge of the *witness* w such that $P(x, w)$ is satisfied. We use $\text{Com}(v; r)$ to denote a commitment to the value v with randomness r . The security parameter of our system is denoted by λ .

3.2. Merkle trees

In `zk-creds` we use Merkle trees T to represent set membership. The root of a tree T is denoted T_{root} . A Merkle *forest* F is a set of Merkle tree roots. Merkle trees have the following functionality:

- $T.\text{Insert}(v) \rightarrow T'$ Inserts the value v into the next free leaf in T and returns the modified tree.
- $T.\text{Remove}(v) \rightarrow T'$ Removes v from the tree (if present) and returns the modified tree.

$T.\text{AuthPath}(v) \rightarrow \theta$ Creates an *authentication path* θ that proves that $v \in T$. The size of θ is proportional to the height of the tree.

3.3. Cryptographic building blocks

We describe two non-interactive zero-knowledge (NIZK) proof systems we use to build `zk-creds`. Both systems operate within a *type-3 non-degenerate bilinear group* which we denote `bg`.

Groth16. Groth16 [39] is a trusted-setup zkSNARK scheme. To describe `zk-creds`, it suffices to specify the functionality of Groth16 that any general-purpose NIZK proof scheme supports:

$\text{G16.Setup}(\text{bg}, \text{desc}) \rightarrow \text{crs}$ Generates a common reference string (`crs`) for the given arithmetic circuit description and bilinear group.

$\text{G16.Prove}(\text{crs}, x, w) \rightarrow \pi$ Proves the circuit described by `crs` is satisfied, where x are the public inputs and w are the witnesses.

$\text{G16.Verify}(\text{crs}, \pi, x) \rightarrow \{0, 1\}$ Verifies the proof π with respect to the given public inputs.

Groth16 Linkage. We describe a high-level interface that allows us to construct a (blinded) *linkage* proof over Groth16 proofs. This allows one to show that a hidden collection of Groth16 proofs π_1, \dots, π_k all share some subset of *hidden common inputs* \mathbf{x}^* , not known to the verifier. Concretely, this proof system proves that

$$\bigwedge_{i=1}^k \text{G16.Verify}(\text{crs}_i, \pi_i, (\mathbf{x}^*, \mathbf{x}_i))$$

where \mathbf{x}_i are the non-hidden public inputs (i.e., public inputs known to the verifier). See Appendix C for the full description and security proofs of LinkG16. For `zk-creds`, however, it suffices to specify the functionality:

$\text{LinkG16.Link}(\mathbf{x}^*, \{\text{crs}_i, \pi_i\}_{i=1}^k) \rightarrow \pi_{\text{link}}$ Constructs a zero-knowledge proof of the above relation, with respect to hidden common inputs \mathbf{x}^* .

$\text{LinkG16.LinkVerify}(\pi_{\text{link}}, \{\text{crs}_i, \mathbf{x}_i\}_{i=1}^k) \rightarrow \{0, 1\}$ Verifies the above statement with respect to the given public inputs and Groth16 CRSs.

3.4. Cryptographic assumptions

We state the security properties of the above schemes and the cryptographic assumptions necessary to achieve them. For brevity, we defer the definitions of the specific assumptions to the cited references.

Groth16 is perfectly zero-knowledge and weak white-box simulation-extractable against algebraic adversaries under q -dlog and a linear independence assumption [4]. We also assume this result holds under the common Groth16 substitution $\gamma = 1$. We use Poseidon [38] to instantiate a hash for Merkle trees, as well as commitments. Finally, we assume that the key-prefixed Poseidon hash function, used to instantiate the gadgets in Section 5.3, is a PRF. As Poseidon is a sponge construction, prefixing is secure. Separately, see Appendix F for an alternate instantiation using Pedersen hashes and perfectly hiding commitments.

4. Definitions

4.1. Security definitions

Security definitions are given by an ideal functionality in Figure 6 of Appendix E, corresponding to the usual security properties of anonymous credentials: unforgeability, correctness and unlinkability. It also implies an additional security property, session binding: shows of a credential are inherently bound to the channel or session in which they are presented, thus preventing replay attacks.

Threat model. The ideal functionality corresponds to the following general threat model. We assume all issuers, verifiers, and (almost all) users are malicious and can collude. We inherit the standard requirement that, for anonymity, there must be at least two honest users with valid issued credentials. We also assume that there exists a reliable mechanism for parties to agree on the list of issued credentials.

4.2. Anonymous credentials

We give a generic overview of the data structures and algorithms which our scheme instantiates.

Let a *credential* be the commitment $\text{cred} := \text{Com}(\text{nk}, \text{rk}, \text{attrs}; r)$ where `nk` is the *pseudonym key*, a private random value used to generate persistent pseudonyms; `rk` is the *rate key*, a private random value used to generate rate-limit tokens; `attrs` $\in \mathcal{A}$ is an arbitrary set of public and hidden *attributes*; and r is the commitment randomness. Note that the values within the credential remain private by the hiding property of commitment schemes, with the exception of attribute information revealed by the user.

We say that an issuer I *issues* a credential if it appears on I 's credential list `CL`. Looking ahead, while the credential list may be instantiated in many ways (e.g. an accumulator, or Merkle tree), we later instantiate this as a Merkle forest and a list of corresponding Merkle trees $\text{CL} := (F, \overline{T})$, with an *authentication path* θ providing membership attestation. Every issuer has some *issuance criteria* ι that the requester must meet in order to have their `cred` issued, e.g., that the birth date in `cred` matches a signed digital passport. Over an issuer-authenticated channel, the requester (running `IssueReq`) sends `cred` to the issuer with some *zk-supporting-documentation* `sd`, e.g., a Groth16 proof or a digital signature, that convinces the issuer of the criteria. The issuer runs `IssueGrant` and, upon success, adds `cred` to their list and returns an authentication path θ attesting to its issuance.

Next, let the list of all possible *access criteria* be $\Phi := \{\phi \mid \phi : \mathcal{A} \rightarrow \{0, 1\}\}$ which can be defined dynamically by users, verifiers, or even third-parties for an application using `zk-creds`, even after system instantiation. Over an anonymous channel, a user *shows* a credential (running `ShowCred`) by presenting a zero-knowledge proof that they have a valid issued credential (with θ as a private witness) whose attributes in witness w satisfy the verifier's access criteria. The proof must also be bound to some session context, `aux`. The verifier runs `VerifyShow` and, upon success, grants access to the user.

An anonymous credential system with *zk-supporting-documentation* can then be defined by the following algo-

rithms (where the subscript U , I , or V denotes that the user, issuer, or verifier runs the algorithm, respectively):

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$ Generates the system parameters.

$\text{IssueSetup}(\text{pp}) \rightarrow \iota$ Establishes the *public* attribute fields and issuance criterion ι for obtaining a credential.

$\text{ShowSetup}(\text{pp}) \rightarrow \phi$ Establishes a new access criterion $\phi \in \Phi$ for showing a credential.

$\text{IssueReq}_U(\text{pp}, \iota, \text{attrs}, w_{\text{sd}}, \text{iaux}_{\text{sd}}) \rightarrow (\text{cred}, \text{sd})$ Creates and requests a credential cred with zk-supporting-documentation sd under issuance criteria ι .

$\text{IssueGrant}_I(\text{pp}, \iota, \text{CL}, \text{cred}, \text{sd}) \rightarrow (\text{CL}', \theta)$ Decides whether to grant a user the requested credential cred ; if so, adds it to the list CL and returns its issuance attestation θ .

$\text{ShowCred}_U(\text{pp}, \phi, \text{CL}, \text{cred}, \theta, w, \text{aux}) \rightarrow (\pi_{\text{link}}, \text{aux})$

Shows that an issued cred satisfies access criteria ϕ .

$\text{VerifyShow}_V(\text{pp}, \phi, \text{CL}, \pi_{\text{link}}, \text{aux}) \rightarrow b$ Validates a credential show, including the current session context in aux .

$\text{RevokeCred}_I(\text{pp}, \text{CL}, \text{cred}) \rightarrow \text{CL}'$ Revokes a credential.

5. Construction

`zk-creds` assumes there is a list of issued credentials maintained by either a trusted party, some Byzantine system, or a blockchain. Our scheme provides three sets of functionalities: `issue`, `show`, and `revoke`. Through the issuance process, a user, Alice, convinces the issuance mechanism she should be given a credential. Once her credential is put on the issuance list, Alice can then use the credential to show she meets some access criteria. Conceptually, using a credential involves two steps: (1) a membership proof that the credential is on the issuance list, and (2) a proof that the committed attributes meet some access criteria. Finally, an issuer is able to revoke a credential if need be by simply removing it from the list.

We can realize this paradigm in different ways and using different set-membership techniques such as an RSA accumulator, purpose built zero-knowledge schemes [61], or even using signatures of issuance.

In our construction of `zk-creds`, we realize membership proofs using Merkle forests, a new approach that allows developers to trade a slight increase in verification time and witness data for a large reduction in proving time. For access criteria checks, we provide developers with a set of *gadgets*. Gadgets can be composed to form complex access criteria checks. Finally, we tie these components together with a new blind Groth16 proof, of potentially independent interest, that lets us show multiple Groth16 proofs shared the same blinded input without—as in `commit-and-prove`—creating a persistent identifier. This allows us to reuse the membership proof across multiple credential shows without the reused proofs being tracked.

We now give details on our specific instantiation and describe the full construction in Figure 3.

5.1. Merkle forests

Rather than using a single Merkle tree to accumulate credentials and prove membership, `zk-creds` uses a forest of Merkle trees. The membership proof attests to two parts: $\text{cred} \in T$ for some Merkle tree T , and $T_{\text{root}} \in F$ where F

is the forest of Merkle trees containing issued credentials. Compared to a single Merkle tree, the Merkle forest approach gives us a tunable tradeoff between proving time and verification time. Shorter Merkle trees can drastically reduce proving costs—up to 50%, or 143ms. Furthermore, since forest membership is a simple OR-proof over Merkle tree roots, the cost of a larger forest is negligible to the prover and allows for a much larger list. Also, we note that, for the size of the forests we consider, the additional verification cost is trivial (137 μ s).

Witness management. Showing a credential in `zk-creds` requires knowing the witness (a.k.a., authentication path) θ attesting to issuance in its Merkle tree. θ must be updated as credentials are added or removed from the tree. In a naïve construction, a user might download newly-issued credentials to update the tree. However, this requires the user to construct and maintain a local copy of their entire Merkle tree, which is often impractical. On the other extreme, users could periodically query an issuer or list manager to provide the updated θ . However, this uniquely identifies the credential and strongly correlates with subsequent shows, posing a large privacy risk (especially in low-use deployments). We present two expedient approaches in Appendix B, and introduce better constructions leveraging Merkle forests in the full version [54].

5.2. Blind Groth16

The membership proof is the most costly part of `ShowCred`. Looking ahead, it takes 460ms to complete. While the access criteria check must be redone for every show in many cases—for example, rate-limited shows include a token that uniquely identifies reuse—the membership proof does not change unless more credentials are issued.

We use a blind Groth16 linkage proof to combine a membership proof with (perhaps multiple) access proofs. Blind Groth16 lets us reuse an already computed membership proof in multiple shows without breaking privacy. Furthermore, it expands the functionality of the system by supporting the easy composition of access criteria: without this ability to compose access criteria, system designers would need to either: 1) dynamically generate circuit parameters for gadgets as they are needed, 2) determine in advance all the gadgets they will support, or 3) generate the circuit parameters for every combination of gadgets that could be used.

Concretely, blind Groth16 lets us prove that a number of Groth16 proofs are all made with respect to the same credential *without* revealing the credential. At a high level, the algorithm works as follows. First, it prepares the public inputs (here, cred and its Merkle root) shared by the underlying proofs. For each proof, it then blinds a copy of the prepared input, and blinds the proof in a way that cancels with the blinded input. Finally, it proves that all the blinded inputs are consistent with each other. After canceling the blinding factors, the verification equation is identical to the typical Groth16 verification equation. For more detail, see the description of `LinkG16` in Appendix C.

<p>Setup($1^\lambda, h, n$):</p> <ol style="list-style-type: none"> 1. Choose bilinear group bg w/ large prime order p 2. Let $desc_T$ be a circuit w/ public inputs $(T_{root}, cred)$, where T is a Merkle tree of height h. It asserts that there is an auth path θ attesting to $cred \in T$ 3. Let $desc_F$ be a circuit w/ public inputs $(T_{root}, cred, F)$, where F is a forest of n Merkle trees. It asserts that the tree is in the forest: $(T_{root} = F_1) \vee \dots \vee (T_{root} = F_n)$ 4. Compute $crs_T := G16.Setup(bg, desc_T)$ 5. Compute $crs_F := G16.Setup(bg, desc_F)$ 6. Let $pp := (bg, crs_T, crs_F)$ 7. Return pp 	<p>IssueSetup(pp):</p> <ol style="list-style-type: none"> 1. Decide on issuance criteria $(\iota_{zk}, \iota_{pub})$ for supporting docs 2. Let $desc$ be a circuit w/ public input $iaux_{zk}$ asserting: $\iota_{zk}(attrs, iaux_{zk}) = 1 \wedge cred$ opens to $(nk, rk, attrs)$ 3. Compute $crs_\iota := G16.Setup(bg, desc)$ 4. Return (crs_ι, ι_{pub}) <p>ShowSetup(pp):</p> <ol style="list-style-type: none"> 1. Decide on access criteria ϕ for satisfying attributes 2. Let $desc$ be a circuit w/ public $(T_{root}, cred, aux)$ asserting: $\phi(nk, rk, attrs, aux) = 1 \wedge cred$ opens to $(nk, rk, attrs)$ 3. Compute $crs_\phi := G16.Setup(bg, desc)$ 4. Return crs_ϕ
<p>IssueReq($pp, crs_\iota, attrs, w_{sd}, (iaux_{zk}, iaux_{pub})$):</p> <ol style="list-style-type: none"> 1. Sample r, nk, rk // commitment nonce, pseudonym key, rate key 2. Commit $cred := Com(nk, rk, attrs; r)$ 3. Let $w := (w_{sd}, r, nk, rk, attrs)$ // collect the witnesses 4. Prove $\pi_\iota := G16.Prove(crs_\iota, (cred, iaux_{zk}), w)$ 5. Let $sd_{zk} := (\pi_\iota, iaux_{zk})$ // collect the supporting docs 6. Let $sd_{pub} := iaux_{pub}$ 7. Send $(cred, sd_{zk}, sd_{pub})$ to issuer 8. Receive the root of the modified credential tree T' and a Merkle auth path θ attesting to $cred \in T'$ 9. Store $(nk, rk, r, attrs, \theta)$ <p>IssueGrant($pp, (crs_\iota, \iota_{pub}), (F, \bar{T}), cred, (sd_{zk}, sd_{pub})$):</p> <ol style="list-style-type: none"> 1. Parse $(\pi_\iota, iaux_{zk}) := sd_{zk}$ 2. Check $\iota_{pub}(sd_{pub})$ // check public supporting docs 3. Check $G16.Verify(crs_\iota, \pi_\iota, (cred, iaux_{zk}))$ // check ZK SD 4. If either check fails, reject issuance and abort 5. Else, choose T from forest F and let $T' := T.Insert(cred)$ 6. Let $\theta := T'.AuthPath(cred)$ // θ attests to $cred \in T'$ 7. Store T' and update the forest F 8. Send θ to user U 	<p>ShowCred($pp, \bar{crs}, (F, T_{root}), cred, \theta, \{w_i, aux_i\}_{i=1}^k$):</p> <ol style="list-style-type: none"> 1. Parse $(\{crs_{\phi_i}\}_{i=1}^k, crs_T, crs_F) := \bar{crs}$ and $cred$'s $T_{root} \in F$ 2. For all $i = 1, \dots, k$, compute the access criteria proof: $\pi_{\phi_i} := G16.Prove(crs_{\phi_i}, (T_{root}, cred, aux_i), w_i)$ 3. Prove $\pi_T := G16.Prove(crs_T, (T_{root}, cred), \theta)$ // tree 4. Prove $\pi_F := G16.Prove(crs_F, (T_{root}, cred, F), nil)$ // forest 5. Let $\bar{\pi} := (\{\pi_{\phi_i}\}_{i=1}^k, \pi_T, \pi_F)$ // collect the proofs 6. Prove $\pi_{link} := LinkG16.Link((T_{root}, cred), \{crs_i, \pi_i\}_{i=1}^{k+2})$ 7. Send $(\pi_{link}, \{aux_i\}_{i=1}^k)$ to verifier V <p>VerifyShow($pp, \bar{crs}, F, \pi_{link}, \{aux_i\}_{i=1}^k$):</p> <ol style="list-style-type: none"> 1. Parse $(\{crs_{\phi_i}\}_{i=1}^k, crs_T, crs_F) := \bar{crs}$ 2. Let $(aux_{k+1}, aux_{k+2}) := (nil, F)$ // collect auxiliary inputs 3. Check $LinkG16.LinkVerify(\pi_{link}, \{crs_i, aux_i\}_{i=1}^{k+2})$ 4. Upon success, accept. Else, reject <p>RevokeCred($pp, (F, \bar{T}), cred$):</p> <ol style="list-style-type: none"> 1. Find $cred \in T$ within forest F; if $cred$ not found, abort 2. Let $T' := T.Remove(cred)$ 3. Store T' and update the forest F 4. Return T'

Figure 3: zk-cred Construction. NB: Although the inputs $(T_{root}, cred)$ are public in all Groth16 proofs in ShowCred, they are hidden from the verifier by LinkG16. Also note that any necessary updates to the auth path θ or credential list (F, \bar{T}) are handled out-of-band.

5.3. Gadgets

Since ShowCred supports arbitrary statements, verifiers have the flexibility to add and remove helpful subcircuits, or *gadgets*, from their protocol. In fact, rather than embedding gadgets in existing circuits, verifiers can make use of the structure of ShowCred to create a separate proof for each gadget and link them together. The benefit to this kind of customization is twofold: users can precompute and cache standalone gadget proofs separately from other access criteria proofs, and verifiers are freed from having to define custom circuits and generate the CRSs.

Gadgets are arbitrary NP relations which can capture nearly any conceivable identity check. We now describe some gadgets that serve as building blocks for zk-creds-based systems. Recall that rk denotes the rate key, used for generating rate-limit tokens, and nk denotes the pseudonym key, used for deriving uniform but linkable tokens.

Linkable show Reveals a pseudonym $PRF_{nk}(ctx)$ that persists across interactions in a given context ctx , but is unlinkable to any use of the credential in other contexts.

For example, a single Sybil-resistant credential could be used for creating unlinkable accounts across sub-forums within a single site, such as Discord servers or subreddits.

Rate limiting Limits users to performing ShowCred only N times per epoch, for some verifier-chosen rate limit N . Every ShowCred, the user produces a pseudorandom token $tok = PRF_{rk}(\text{epoch}||\text{ctr})$, reveals epoch, and proves that ctr is less than N .

Cloning resistance Performs the same function as rate limiting, but deanonymizes rate violators. The technique was introduced by Camenisch et al. [19] (Section 5.2). Every run of ShowCred, the user receives a nonce from the verifier and sends two tokens:

$$\begin{aligned} tok_1 &= PRF_{rk}(\text{epoch}||\text{ctr}) \\ tok_2 &= id + H(\text{nonce}) \cdot PRF'_{rk}(\text{epoch}||\text{ctr}) \end{aligned}$$

where id is an identifying attribute (e.g., credential hash). As above, ShowCred proves the tokens are constructed correctly. If one of these shows is reused, tok_1 will be repeated, but tok_2 will be distinct, giving the verifier

two instances of the tok_2 equation and two unknown variables: id and $\text{PRF}'_{\text{rk}}(\text{epoch}||\text{ctr})$. Solving the equation for id identifies the credential holder. Note that if id is the credential (or its hash), then the cloned credential can be revoked immediately by removing it from the issuance list.

Expiry Opens an attribute e in the credential and proves that $e > \text{today}$, i.e., that the credential is not yet expired.

Session binding Gives the verifier the ability to reject replayed ShowCred proofs by binding a verifier-chosen nonce, or *session context*, to every ShowCred. This can be done by including an empty proof that takes the nonce as the public input⁶.

Join Allows credentials to be composed by *joining* them along some common attribute(s), such as full name or address. This is either done inside a single ShowCred, or between separate ShowCredits by publicly committing to the common attribute(s).

5.4. Signature-issued credentials

One of the major advantages of `zk-creds` is that we do not need issuers who are trusted to hold signing keys and can instead use a transparency log or blockchain to issue credentials. However, if this feature is not needed, we can build a traditionally issued anonymous credential scheme that retains `zk-creds`' other features and flexibility. Because our gadget-based approach is flexible, we can replace the membership-check gadget with one that instead verifies a signature. This means `zk-creds`, like Coconut [55], also supports credentials that are issued by signing under a standard signature (e.g., ECDSA or Schnorr) either by a single party or by a threshold of parties via a threshold signature scheme such as FROST [44]. Moreover, we can compose credentials issued via a list with ones issued via signatures.

5.5. Additional features

Notably, our construction of `zk-creds` also allows for the construction of protocols with several features previously only available in dedicated schemes.

Hidden issuer We can completely hide the identity of a credential issuer, e.g., in situations where leaking where a credential came from can cause significant harm to privacy. While this is not a new notion in the literature, very few existing schemes support this hidden issuer property. `zk-creds` supports this inherently, as ShowCred can be performed with respect to synthetic lists created by concatenating lists maintained by different issuers, thus hiding the issuer. One drawback though is that multiple issuers will likely issue different credential formats.

Hidden credential type `zk-creds` can also be configured to hide the credential type which is both independently useful as well as necessary to fully support hidden issuer. In `zk-creds`, credentials with different attributes can be padded to the same size post-issuance, and then used

interchangeably for and efficiently verified over an access criterion. This is accomplished by constructing a new circuit that is the OR of the criteria on the individual credentials.

Delegation Issuance authority is delegatable in `zk-creds`. Authority to issue a credential can be shown via `zk-supporting-documentation` that is itself the show of another credential. Moreover, because the proof in `zk-supporting-documentation` is general-purpose, the delegation process can constrain attributes in the credential being issued. For example, we could define a credential for an authority that can only be delegated three layers deep by having a hidden attribute of delegation level decremented each time. Credential attributes can be selectively delegated as well.

5.6. NIZK setup

Groth16 requires a one-time trusted setup to generate a set of parameters called a common reference string (or CRS) for each statement (a.k.a., circuit). Once this CRS is generated, it can be used throughout the lifetime of the system to prove different instances of the statement.

Distributed setup for Groth16 CRSs is a solved problem via multiparty computation setup *ceremonies* [12, 8, 43] that need only two honest parties. These have been run with hundreds of users and used to secure billions of dollars in cryptocurrency. These protocols are efficient and produce *subversion-resistant zero-knowledge proof systems* [33]—systems which ensure that, even if all parties in a setup are malicious, the proofs are still zero-knowledge and user privacy is unaffected.

Independently, as mentioned previously, we have also sought to minimize the impact of this CRS on the flexibility of `zk-creds`. By utilizing blind Groth16, a system designer does not need to decide on and pre-generate CRSs for all possible combinations of access criteria they wish to support and can instead just generate CRSs on a per-gadget basis.

5.7. Security argument

The security of `zk-creds` rests on the assumption that Groth16 proofs are sound, perfectly zero-knowledge, and simulation-extractable (in the Algebraic Group Model), that the Poseidon hash function is collision resistant and a PRF, and, for other variants of our scheme, that Schnorr signatures are secure under the discrete log assumption, and Pedersen commitments are computationally binding under the same assumption and perfectly hiding. Informally, because all messages between parties in our security game are “authenticated” by zero-knowledge proofs, the simulator can extract on all adversarially generated messages and proxy them to the ideal functionality. Similarly, for any honest interactions in the ideal functionality, the simulator can model the adversary's view of the real-world protocol by simulating the zero-knowledge proofs with respect to random commitments and random outputs. Finally, the simulator maintains a list mapping real and ideal world credentials to allow it to handle revocation events. We defer a full

6. This binds the nonce to the Groth16 proof, assuming some basic properties about the circuit [4].

description of the simulator to Appendix E and provide a proof of security in the full version [54].

6. Implementation and evaluation

We now detail the evaluation of `zk-creds`.

6.1. Code and setup

Hardware. All benchmarks were performed on a desktop computer with a 2021 Intel i9-11900KB CPU with 8 physical cores and 64GiB RAM running Ubuntu 20.04 with kernel 5.11.0-40-generic.

Code. `zk-creds` consists of 7.6k lines of Rust code⁷ and relies on the Arkworks [3] zkSNARK framework. For benchmarks and statistics, we used the Criterion-rs crate. In addition, we modified an existing Android passport scanner app to extract intermediate cryptographic values from the passport and dump them to a JSON file.⁸ The Rust code uses the dump file format for all its passport proofs.

Statistics. Each figure and plot shows the median runtime of 100 executions. Over all experiments, the maximum estimated relative standard error of the median is 1.8%. For completeness, our plots include error bars indicating the 95% confidence interval, though they are not visible due to the low error.

Instantiating cryptographic primitives. We set $\lambda = 128$. We compute Groth16 proofs over the BLS12-381 curve [11]. The collision-resistant hash function used in all Merkle trees are domain-separated instances of the Poseidon hash function [38], configured to be compatible with BLS12-381 and a 128-bit security level ($\alpha = 3$ and capacity = 1). We instantiate Com using key-prefixed Poseidon as well. Separately, in Appendix F, we give benchmarks for an instantiation with provably secure Pedersen hashes and commitments.

6.2. Microbenchmarks

In this section we measure performance of various common gadgets, with reasonable parameter choices. Recall the performance of `zk-creds` depends on credential list size, attribute size, criteria complexity, and number of standalone gadget proofs. We measure the effects of these parameters in greater detail in Appendix A. We also measure the sizes of the associated proving and verifying keys in Appendix G.

Table 1 gives a summary of `zk-creds`'s performance for common usage scenarios. We assume a setting where 2^{31} credentials have been issued. The first variant (C) minimizes client-side proving cost by separating the Merkle forest proof, allowing for reuse across shows. A basic show takes 5ms to produce and 3ms to verify, assuming precomputation of the Merkle tree and forest membership proofs. More complex statements like rate-limited credentials with clone resistance take 90ms to show and 5ms to verify. If the Merkle tree membership proof is not precomputed, then the full show takes an additional 460ms but verification time is unchanged.

7. Code repository: <https://github.com/rozbb/zkcreds-rs>

8. Code repository: <https://github.com/rozbb/zkcreds-passport-dumper>

The Simple Possession benchmark shows the prover has an attribute-less credential on a list and proves no predicates. This maps to a use case such as possessing a valid access card, as that is often sufficient to enter a building. The remaining benchmarks build on Simple Possession, adding their own predicate to the set of linked proofs. For example, Expiry proves possession, but also bears a single attribute and proves that it has a value less than some timestamp.

Separately, we give an alternate variant of `zk-creds` in Table 1 (S) which is optimized for verification latency and server throughput. The server-optimized construction combines the Merkle membership and criteria check circuits into a single monolithic zkSNARK without proof reuse. As a result, clients pay approximately the full ShowCred cost every time, but since proofs are a single Groth16 proof rather than a linkage proof, they can be batch-verified by the server at 1.8 verifications per millisecond per core. We note it may be possible to batch verify the non-optimized scheme as well, but throughput would be lower as there are at least three times as many proofs.

Finally, to demonstrate the full flexibility of our approach, we also provide a signature-based variant of `zk-creds` in Table 2, where credentials are issued in a more traditional signature-based manner by a trusted issuer or threshold quorum of issuers. We implemented a signature gadget for checking Schnorr (and therefore FROST threshold [44]) signatures and measured the proving time to be 129ms (compare to 460ms for tree-based issuance).

7. Case studies: `zk-creds` as a toolkit

We design, implement, and benchmark two full scenarios for `zk-creds` using credentials derived from existing government identity infrastructure without any modifications or coordination. Many government identity documents now include the ability to perform various authentication protocols (e.g, the German and Estonian [32, 29] smart-card enabled national IDs). For our applications, we use US passports, which contain a signed digital copy of the passport's basic data in an NFC-readable chip. Our applications validate that `zk-creds` can be used as a toolkit by application developers to support privacy-preserving identity in realistic applications with complex, compound access criteria.

7.1. Digital passport data

Over 150 countries issue passports with an NFC-readable chip which is standardized in ICAO Doc. 9303 [53, 40]. We are interested in the first two *data groups* on the chip. Data group 1 (DG₁) contains the textual info available on the passport's data page: name, issuing state, date of birth, and passport expiry. Data group 2 (DG₂) contains a JPEG-encoded image of the passport holder's face. The ICAO standard also requires the immutable part of the chip's contents to be signed by the issuing state. For example, every US passport has an RSA PKCS#1 v1.5 signature under a known US State Department public key.

zk-supporting-documentation for passports. While we could just reveal the signed passport to the credential

Client-opt. (C) Server-opt. (S)	ShowCred			VerifyShow			Proof Size	
	C	C (full)	S	C	S	S (batch)	C	S
Simple Possession	5ms	465ms	450ms	3ms			744B	
Expiry	53ms	526ms	461ms	} 5ms	} 1.5ms	} 1.8 verifs/ms	744B	192B
Linkable Show	43ms	494ms	457ms					
Rate Limiting	60ms	507ms	461ms					
Clone Resistance	90ms	542ms	530ms					

TABLE 1: Gadget microbenchmarks using Poseidon hashes for two versions of `zk-creds`. Membership proofs are done on a Merkle forest of size 2^{31} (tree height = 24, #trees = 2^8). The first configuration (C) minimizes client-side proving cost; the second configuration (S) maximizes server throughput. ShowCred (full) gives the cost of including membership recomputation while showing a credential. VerifyShow (batch) gives throughput for verifying a set of 100 proofs. We emphasize that all verification numbers are *single-threaded*, allowing for efficient concurrent processing.

Client-opt.	ShowCred	VerifyShow	Proof Size
Simple Possession	2ms	2ms	424B
Expiry	50ms	} 5ms	744B
Linkable Show	41ms		
Rate Limiting	58ms		
Clone Resistance	88ms		

TABLE 2: Benchmarks for `zk-creds` configured for minimizing client side proving cost, using signature-based issuance. Unlike the client-optimized benchmarks in Table 1, which uses list-based membership, there is no distinction between a partial or full ShowCred, since the membership proof (a signature verification proof) never has to be recomputed.

issuer, this 1) requires the issuer be trusted to maintain the confidentiality of sensitive information, and 2) is wholly incompatible with issuance via a bulletin board or blockchain. Instead, we design and implement a zero-knowledge proof that the attributes of a credential commitment match the signed contents of DG_1 and DG_2 .

Parsing the contents of an e-passport in zero-knowledge is non-trivial: the signature is not just over DG_1 and DG_2 , but the *econtent hash*, which is calculated over the mostly variable-length data groups DG_1, \dots, DG_{16} . Variable length inputs are particularly challenging to parse with a fixed-size zero-knowledge circuit. Our zero-knowledge proof is made possible by realizing that the econtent hash is actually a tree hash, roughly of the form $H(H(H(DG_1) \| H(DG_2)) \dots \| H(DG_{16})) \dots$. We do not care about the contents of DG_3 through DG_{16} , so their hashes can be used directly as witnesses to the proof. $H(DG_2)$ is the image hash, which can either be hidden as a witness or revealed by showing all of DG_2 .⁹ Since DG_1 contains the attributes we care about, we must parse DG_1 inside the zero-knowledge proof. Luckily, since DG_1 mirrors the content of the passport’s Machine Readable Zone (MRZ), it is fixed-length. The proof then hashes the data group digests along with other fixed-length values until it has computed the econtent hash. We avoid the cost of checking the RSA

9. In reality, DG_2 contains slightly more than the bare image; it also has the image’s creation and expiry dates.

signature by simply revealing it and the econtent hash.¹⁰ Computing this proof takes less than 2 seconds.

Why zk-proofs over passports are insufficient. We cannot just use the zero-knowledge passport proof as a credential for accessing age-restricted content: to prevent credential sharing, we need cloning resistance, which requires that credentials include a secret random seed that a passport lacks.

7.2. Instantiating an issuance bulletin board

An instantiation of `zk-creds` requires a publicly accessible bulletin board to distribute the credential list, as well as parties running our software. We stress that `zk-creds` can be deployed either on a blockchain or (at much lower cost) a more centralized transparency log; we choose to instantiate the bulletin board here as a smart contract on an Ethereum Virtual Machine (EVM)-compatible blockchain to demonstrate that full decentralization is feasible and because, unlike many other consensus systems in development or deployment, the EVM has comparatively robust development tooling and documentation.

The smart contract stores a list of credentials and corresponding zk-supporting-documentation—together, referred to as an *issuance request*—as well as the current roots of the Merkle forest. An issuer issues a credential by posting the full issuance request to the smart contract. This allows any external auditor to download the full list, reconstruct a local copy of the Merkle forest, then verify in zero-knowledge that each credential was validly issued. While any party (including the user) can audit the full list themselves if desired, they need not do so if they trust another party (e.g., an issuer or auditor) to promptly update and verify the credential list for them. Periodically, users request their credential’s authentication path and the updated root from the bulletin board or auditing party. To perform VerifyShow, a verifier only needs to retrieve the current Merkle tree roots F from the bulletin board.

Furthermore, we must prevent DoS attacks from blocking or flooding additions to the bulletin board. Our prototype instantiation assumes a smart contract operator who is

10. Revealing the econtent hash reveals the identity of the requester to the passport authority during issuance. However, since ShowCred proofs are unlinkable to issuance, this is not a problem. Moreover, it is possible to hide the hash and verify the RSA signature in the zkSNARK, at a cost of about 868k constraints using the approach from [45], or 3.6s in our benchmarking environment.

	IssueReq	IssueGrant	ShowCred	ShowCred (full)	VerifyShow
Age-restricted vid.	1.97s	2ms	143ms	602ms	5ms
Entering a bar	1.97s	2ms	98ms	557ms	5ms

TABLE 3: `zk-creds` case study benchmarks. IssueReq is the time to convert a passport into a credential using `zk-supporting-documentation`. ShowCred is the time it takes a user to prove they are over 18. All other parameters are the same as in Table 1.

authorized to add to the bulletin board contract on behalf of the issuer. An alternative solution, allowing for operator-free setup, is to verify all `zk-supporting-documentation` and Merkle tree root computations within the EVM smart contract itself. While this is feasible [58], verifying proofs and Merkle tree updates on-chain without extensive optimizations is expensive. Another solution is to instead support on-chain proof verification with an optimistic rollup [31]: bulletin board additions include a deposit which is burned if, after the smart contract evaluates a challenge, it determines that the proofs or computed Merkle tree roots are invalid. While the challenge still requires costly computation, this is not paid for during issuance.

We implement our smart contract in Solidity and the requisite client-side scripts to post and retrieve data in `web3.js`. Note that we can deploy our contract to any EVM-compatible chain, such as Avalanche’s C-Chain or Ethereum itself. EVM contract operations are measured in *gas*, which roughly acts as a complexity-weighted count of the EVM instructions used. Posting each issuance request costs 576,808 gas, while posting a Merkle root costs 78,355 gas. The price of gas and the underlying currency it is priced in is highly volatile. As of early June 2022, Ethereum gas prices range from about 20 to 50 Gwei (a Giga-wei is 10^{-9} ETH), and 1 ETH is about \$1800 USD; as such, posting an issuance request costs about \$20–\$50 USD. Gas prices are proportionally similar for Avalanche but, at \$25 USD per token, actual costs would be 70 times smaller.

7.3. Credentials from existing identity infrastructure

Given a construction of `zk-supporting-documentation` for a passport and a choice of bulletin board in an Ethereum smart contract, application developers can now readily build access control schemes. Once the credentials are issued, multiple developers can independently rely on them by either composing existing identity gadgets or defining new ones.

Issuance. We provide a toolchain to convert a passport into an anonymous credential. An Android app extracts the NFC passport data and a separate program converts it into a credential containing the holder’s nationality, full name, and date of birth (*dob*); a rate key (*rk*); the passport expiry date (*expiry*); and the hash of the image of the holder’s face (*facehash*). Separately the program computes the `zk-supporting-documentation` that this credential is correct with respect to the signature and *econtent* hash.

The IssueReq payload sent to the issuer consists of the signed *econtent* hash, credential, and `zk-supporting-documentation` proof. IssueGrant verifies the proof and *econtent* hash signature. Upon success, the issuer adds the

credential to their list and returns a Merkle authentication path.

Scenario 1: Viewing age-restricted content on the internet. Age-restricted content is common on the internet. For example, in Switzerland, the EU, and the UK, YouTube requires users to upload an image of their ID or credit card in order to prove their age [36]. In this scenario we demonstrate the feasibility of `zk-creds` for proving age without revealing any personally identifying information.

Our `zk-creds` toolkit has three features that are crucial to building a real-world feasible age verification credential. First, it can be used without coordination with existing identity infrastructure. Second, it can readily support other identity credentials, provided they indicate date of birth and are signed. Third and, most crucially, it can create credentials that are clone-resistant (via the gadget in Section 5.3) with easy revocation of cloned credentials. This last point is essential: while a zero-knowledge proof over a passport is itself an anonymous credential,¹¹ practical usage demands cloning protections. And cloning resistance requires a rate key to be bound to the credential and kept secret from the issuer. Existing identity documents (such as a passport) lack such a key. `zk-creds` allows composition of identity without coordinating with the passport issuer or any trusted party to add such information.

Given issued credentials via passports, building a privacy-preserving age verification scheme with `zk-creds` is straightforward and requires no new cryptography: website developers need simply define the issuers they will accept¹² and construct the access criteria they need using gadgets. For this scenario, the only issuer is our passport-based issuer, and the access criteria being proved are age, expiry, and non-cloning. Concretely, the access predicate is:

$$\begin{aligned} & \text{dob} \leq \text{today} - 18\text{yrs} \wedge \text{expiry} > \text{today} \\ & \wedge \text{CloneResistance}(\text{rk}, \text{nonce}, \text{tok}_1, \text{tok}_2, \dots) \end{aligned}$$

where `CloneResistance`, `nonce`, `tok1`, and `tok2` are as described in Section 5.3.

Table 3 gives performance numbers. Given a credential, it takes Alice 143ms to show a website she is over 18 (and 602ms when she must recompute her membership proof). The server can verify her assertion in 5ms. If we wish to optimize for server verification time or throughput, we can switch to `zk-cred`’s server-optimized construction and achieve 1.5ms verification times and 1.8 verifications per ms per core. Extrapolating from the server-optimized benchmarks in Table 1, proving times would increase to approximately 595ms.

11. When augmented to hide the passport signature.

12. This defines which credential list they use or defines a new list as some subset of existing ones.

Scenario 2: A cryptographer walks into a bar. To purchase alcohol in the United States, one needs to show photo ID and proof they are at least 21 years old. But showing a driver’s license reveals name, sex, weight, and date of birth. And if the license’s barcode is scanned [1], additional information is revealed, potentially including whether the holder is insulin-dependent, hearing-impaired, developmentally disabled, or, surprisingly, a sexual predator [59].

We now build a system for in-person age verification coupled with photographic verification. Importantly, our goal in this scenario is *selective disclosure* and not anonymity. Anonymity in an in-person setting is often not possible or even desired. Rather, what we want is privacy: the ability to control what information is revealed and limit it to what is necessary—the patron’s photo and the fact that they have an unexpired ID with a birth date making them of drinking age.

We envision a hypothetical system where bar patrons have an ID-wallet application on their phone. The app, using ShowCred, presents an identity assertion (e.g., via a QR code or NFC) to an app on a bouncer’s phone which checks the assertion with VerifyShow and displays the user’s photo along with whether they are over 21. In contrast to scanning the user’s driver’s license, this reveals only the minimal information necessary. The necessary access predicate is:

$$\text{dob} \leq \text{today} - 21\text{yrs} \wedge \text{expiry} > \text{today} \wedge \text{facehash} = H(\text{face})$$

Table 3 gives concrete costs for local computations. To show that a patron is over 21 takes 98ms in the common case and 557ms if membership proofs must be recomputed. As before, verification is 5ms in either case.

8. Extensions and applications

We now discuss extensions to `zk-creds` and additional applications of our approach.

Other signed identity sources. As shown by our e-passport example, `zk-creds` can transform legacy identity sources into anonymous credentials if there is a digitally signed component. This raises an interesting question about what parts of existing identity and credential infrastructure include such signatures. For example, digital diplomas for many US universities include digital signatures over the diploma holder’s name, degree, and institution. Many emails are signed with DKIM which, while problematic in many contexts [56], could be a source of identity or membership in an organization. New York’s Excelsior Pass for COVID vaccination contains the holder’s name, birth date, and a signature.¹³ Other existing digital protocols may contain a signature that establishes ownership of a resource (e.g., a phone number in an eSIM or virtual SIM card) or identity (e.g., Apple’s digital driver’s license features).

Complex access criteria. We have discussed conceptually simple access control criteria such as “my credential is not expired,” or “I am of age,” perhaps with a cryptographically

complex mechanism for clone resistance. However, real-world access criteria can be far more complex. `zk-creds` provides a way to deal with such criteria without requiring coordination with identity issuing authorities for every custom access check that must be implemented.

An example of this comes in the form of online petitions and discussions. New York State has an online portal for discussion and petition which asks a user for their address to match them with the appropriate state senator [50]. While this check does not seem to be enforced, one could imagine both wanting to enforce this constituency check and allowing constituents to leave non-identifying comments. Similarly, some online resources, such as ebooks from the New York Public Library, are limited to city residents; enforcing this currently requires in-person registration for a library card to present proof of address, and opens up a (hypothetical) risk of tracking online reading habits [2].

Geolocating an address to the bounds of, e.g., a city council district, however, is not simple. The computation is, by the standards of credential schemes, complex, and involves converting the address to a location and then performing a point-in-polygon check.¹⁴ For a small number of fixed boundaries (e.g., federal congressional districts), one might imagine avoiding the problem by issuing identity documents with this information included. But even in cases where the identity issuer would cooperate, coordinating all geocoding restrictions one might want to realize (e.g., anonymous discussion boards for a school zone, a neighborhood, or even specific apartment building) is impractical and may cause credential sizes to blow up.

Because `zk-creds` supports general purpose zero-knowledge proofs, geocoding restrictions are made more feasible with Groth16 gadgets: even if the Groth16 proof for the gadget is expensive, the resident or an outsourced prover avoids recomputing it every show. After the first time, the proof can be reused arbitrarily, until the user’s Merkle tree is updated by a new issuance.

Sybil-resistant IDs from email or money. Internet services currently prevent multiple account registration (Sybils) by requiring the user to consume a (hopefully) scarce resource, such as money (via a micropayment), attention (via a CAPTCHA), and possession of, e.g., phone number or email address.

`zk-creds` provides several avenues for Sybil-resistant credentials with anonymity. Credentials can be issued based on signed identity documents (e.g., a passport, as demonstrated in Section 6) with the signature as a uniqueness check. Similarly, `zk-creds` can thwart Sybils via cryptocurrency: a simple smart contract issues credentials if and only if a small fee is paid.

Finally, and perhaps most surprisingly, we can use possession of a valid email address as a Sybil-resistance mechanism without the use of a trusted third party or cooperation with the email provider. A DomainKeys Identified Mail (DKIM)

13. This was obtained by scanning the pass’s QR code, whose contents are a W3C Verifiable Credential [60].

14. Indeed, geolocation is the correct way to prove residency. For example, the borders of New York City’s 10 city council districts are defined by 1.8 megabytes of geospatial vector bounds [51].

header, which appears on all outgoing mail of most modern email providers, contains a signature from the email provider. By embedding the credential issuance request in the email body, we get an externally verifiable proof of possession of an email address that can be used to issue Sybil-resistant accounts. This allows us to leverage the Sybil resistance mechanisms used by Gmail, for example.

These short- and long-lived IDs can be reused and rate-limited across actions, similar to the functionality of Privacy Pass [28], which issues one-time-use anonymous access tokens for every CAPTCHA a user completes.

Oracle- and self-issued credentials. A number of academic works and industrial systems have emerged to address the so-called “oracle problem”: how does a consensus scheme such as a blockchain come to agreement about real world events?

One class of solutions [30] relies on incentive systems and the ability to challenge the veracity of data. These approaches, if viable, could be used to issue anonymous credentials based on public online reputation data (e.g., Twitter follower count, Reddit karma, etc.). Crucially, because *zk-creds* forces all issued credentials to be on a public list, any malfeasance by an issuer could be detected and punished.

An orthogonal approach is the creation of notaries who attest to data on third party servers. DECO [62] proposes a 2PC protocol between a client and a *notary* that authenticates data retrieved over a TLS connection from a third-party server. DECO also allows users to construct zero-knowledge proofs to only selectively disclose HTTP response contents. These features would allow a user to obtain a credential for their name and address by, e.g., logging into their utility provider and retrieving the bill. Moreover, in the event the user has trustworthy secure hardware, they can *self-issue* the credential by attesting the hardware ran this notary check itself.

We note, however, that such schemes inherently rely on assumptions (either non-collusion or the security of trusted hardware) that become increasingly infeasible as the value of the credential goes up. But, for example, it may be viable as a simple Sybil prevention or anti-spam mechanism.

Composable credentials. When new use cases for existing credentials emerge, they often require the combination of two different credentials. Take, for example, US-issued vaccine cards. Because these contain a person’s name, but not a photo, COVID-19 vaccine mandates frequently required restaurants and bars to ask customers for a vaccine card *and* a photo ID with a matching name. *zk-creds* supports this type of post-hoc composition: two credentials both containing a field for, e.g., full name, can be jointly shown using the Join gadget described in Section 5.3.

9. Related work

Anonymous credentials derive from a long line of work, starting with Chaum [26], and subsequently seeing numerous extensions [26, 20, 21, 22, 19, 6, 18, 5, 35, 17, 55]. While showing a credential initially allowed little more than (unlinkably) presenting a signed token connected to a user’s

pseudonym, the schemes were generalized and extended to provide more sophisticated properties such as issuance of hidden attributes, rate-limiting, *k*-show, and efficient selective disclosure of attributes. Because it uses general-purpose zero-knowledge proofs, *zk-creds* can (and does) capture all of these properties as implemented.

One drawback to deploying the majority of these schemes is the requirement of a single, trusted issuer. As such, existing work has sought to solve this issuance problem. We briefly compare and contrast other approaches to addressing this.

Distributed issuance. In 2014, Garman et al. [35] were the first to propose the notion of decentralized anonymous credentials. Our approach is directly inspired by their work.

While the approach of Garman et al. removes the assumptions of a single issuer and the need for issuers to hold keys, it both leaves open a number of essential questions for operating a real system and introduces new ones which we address. First, showing a credential requires users to 1) locally store the full credential list and 2) compute proofs which take time linear in the number of issued credentials for every show (even for static credential lists). In contrast, *zk-creds* develops a new approach and new cryptography (blind Groth16) to allow proof reuse and fast credential showing. And, via the use of Merkle trees, *zk-creds* requires users to store only logarithmic-sized witness data to compute credential shows. Second, while Garman et al. suggest the possibility of more complex features, they do not implement them. More importantly, the set of bespoke sigma protocols they use does not provide for the composition of credentials and identity attributes or support complex identity statements. By developing a protocol based on general-purpose zkSNARKs, we do. Lastly and most significantly, the work does not answer the question of how the decision to issue a credential could be made without disclosing sensitive information to the issuer (i.e., the very problem zk-supporting-documentation addresses). This also makes it impossible to audit credentials that are issued via Garman et al.’s construction, in marked contrast to *zk-creds*.

Threshold issuance. The Coconut [55] anonymous credential scheme addresses the issuance challenge via threshold signatures. In Coconut, credentials are issued by *n* static parties under the assumption $t > 1/2$ of them are honest. The scheme is clever and achieves efficiencies on par with single-issuer schemes. However, while threshold issuance increases the security of a scheme by requiring an attacker to corrupt more parties, it only addresses the scarcity of issuers if we have an abundance of parties who are willing to issue credentials but, for whatever reason, no individual one is trusted. In many settings, it is a challenge to find even a single party who both 1) is empowered to make identity statements and 2) is willing and able to run cryptographic infrastructure even if, by fiat, we trust them. Moreover, Coconut only supports selective disclosure of attributes in a credential, not complex zero-knowledge proofs over attributes. It does not meet our design goals of flexibility or dynamic generation of access criteria.

Finally, we note that *zk-creds*, as shown by the version

given in Table 2, supports credential issuance via threshold Schnorr schemes (e.g., FROST [44]), so we do capture the functionality of Coconut.

Decoupling issuance from identity verification. More broadly, another line of work, starting with TLSNotary [57], considers convincing third parties of the correctness of data. DECO [62] extended this protocol to support TLS 1.3 and used zero-knowledge proofs for selective disclosure (e.g., the party learns a bank account balance is over a threshold, but not the balance itself or the account holder’s identity). Applying this to the anonymous credential setting, one could use it to separate finding a cryptographic issuer for credentials from the process of verifying entitlement to a credential. Such approaches are complementary to `zk-creds` and we consider them as an extension in Section 8. Without such integrations, however, such an approach would still require at least one trustworthy party to be willing to run a highly available web service that holds live signing keys.

Decentralized identity (DID). We note another area of research on decentralized identity which has been the subject of both academic and industrial work, including some standardization. While `zk-creds` is compatible with these goals inasmuch as it allows issuance without a party holding keys, such works are largely orthogonal: decentralized identity as an area considers who makes decentralized identity assertions, while our work considers how to transform some identity assertion (centralized or otherwise) into an anonymous credential without introducing additional trusted parties.

10. Conclusion and future work

The approach we develop here—a switch from blind signatures to zero-knowledge proofs as the foundation for anonymous credentials—implies several avenues for further work. In this section, we enumerate a few immediate consequences of this paradigm and future research areas.

Instantiating `zk-creds` with improved zero-knowledge proofs. We have instantiated `zk-creds` with Groth16. The `zk-creds` approach we develop, however, is proof system-agnostic. As such, instantiating the `zk-creds` paradigm with other proof systems, such as ones without trusted setup (e.g., [7, 39, 15]) or with universal setup (e.g., [48, 34, 27]), is entirely possible. For the monolithic construction, such a change is a drop-in replacement. If we wish to support precomputation of separate membership proofs, as in our client-optimized scheme, we must either adapt blind Groth16 to the new proof system or take an alternative approach, e.g., recursive proofs. The choice of proof system is also tied to the choice of accumulator scheme.

Instantiating `zk-creds` with alternative accumulator schemes or primitives. We have instantiated `zk-creds` using Merkle trees. However, as with proof schemes, the same approach generalizes to other accumulator mechanisms such as RSA accumulators [35], polynomial commitments [42], Verkle trees [46], or perhaps special-purpose schemes for zkSNARKs [61]. Again, for many such accumulators, this is a simple black-box replacement of the membership zkSNARK. Instantiating `zk-creds` with alternative accumulators will

offer different tradeoffs for witness size, witness update requirements, witness computation cost, and accumulator verification cost (and hence zkSNARK proving time). A particularly exciting prospect is the co-design of accumulator schemes and zkSNARKs to achieve drastically improved performance.

Co-design or co-selection of zero-knowledge proofs and proving systems. Similarly, one could instantiate either our existing version of `zk-creds`, or a different construction, with different cryptographic primitives. For example, one might replace Poseidon with a newer, circuit-optimized hash function, perhaps making use of low-degree gates in proving systems like Plonk [34].

11. Acknowledgements

We would like to thank Mary Maller for the idea and sketch of the LinkG16 proof system, and Daniel Benarroch Guenun for starting the discussion of Merkle Forests. We would like to thank the anonymous reviewers and shepherd for their helpful comments.

Michael Rosenberg’s work was supported by the National Defense and Engineering Graduate (NDSEG) Fellowship. Jacob White and Christina Garman’s work was partially supported by NSF grants CNS-1816422 and CNS-2047991. Ian Miers’s work was supported in part by a Facebook Research Award. Image credits [47].

References

- [1] ACLU Blog - Jay Stanley, “A Creeping Private-Sector ‘Checkpoint Society’—and a Small Step to Protect Your Privacy,” 2012, <https://www.aclu.org/blog/national-security/creeping-private-sector-checkpoint-society-and-small-step-protect-your>.
- [2] American Civil Liberties Union, “Librarians Speak Out for First Time After Being Gagged by Patriot Act,” 2006, <https://www.aclu.org/press-releases/librarians-speak-out-first-time-after-being-gagged-patriot-act>.
- [3] Arkworks-rs, *Arkworks Ecosystem Homepage*, 2022, <https://arkworks.rs/>.
- [4] K. Bagheri, M. Kohlweiss, J. Siim, and M. Volkhov, “Another look at extraction and randomization of groth’s zk-snark,” in *Financial Cryptography and Data Security*.
- [5] F. Baldimtsi and A. Lysyanskaya, “Anonymous credentials light,” in *ACM CCS 2013*.
- [6] M. Belenkiy, M. Chase, M. Kohlweiss, and A. Lysyanskaya, “P-signatures and noninteractive anonymous credentials,” in *TCC 2008*.
- [7] E. Ben-Sasson, I. Bentov, Y. Horeish, and M. Riabzev, “Scalable, transparent, and post-quantum secure computational integrity,” *Cryptology ePrint Archive*, Report 2018/046, 2018, <https://eprint.iacr.org/2018/046>.
- [8] E. Ben-Sasson, A. Chiesa, M. Green, E. Tromer, and M. Virza, “Secure sampling of public parameters for succinct zero knowledge proofs,” in *2015 IEEE Symposium on Security and Privacy*.
- [9] D. Boneh, B. Bünz, and B. Fisch, “Batching techniques for accumulators with applications to IOPs and stateless blockchains,” in *CRYPTO 2019, Part I*.
- [10] S. Bowe, *Bellman: zk-SNARKs in Rust*, 2017, <https://electriccoin.co/blog/bellman-zksnarks-in-rust/>.
- [11] —, *BLS12-381: New zk-SNARK Elliptic Curve Construction*, 2017, <https://electriccoin.co/blog/new-snark-curve/>.

- [12] S. Bowe, A. Gabizon, and I. Miers, “Scalable multi-party computation for zk-SNARK parameters in the random beacon model,” Cryptology ePrint Archive, Report 2017/1050, 2017, <https://eprint.iacr.org/2017/1050>.
- [13] E. F. Brickell, J. Camenisch, and L. Chen, “Direct anonymous attestation,” in *ACM CCS 2004*.
- [14] E. Brickell and J. Li, “Enhanced privacy id from bilinear pairing,” Cryptology ePrint Archive, Report 2009/095, 2009, <https://ia.cr/2009/095>.
- [15] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, “Bulletproofs: Short proofs for confidential transactions and more,” in *2018 IEEE Symposium on Security and Privacy*.
- [16] V. Buterin, “An Incomplete Guide to Rollups.”
- [17] J. Camenisch, M. Dubovitskaya, K. Haralambiev, and M. Kohlweiss, “Composable and modular anonymous credentials: Definitions and practical constructions,” in *ASIACRYPT 2015, Part II*.
- [18] J. Camenisch and T. Groß, “Efficient attributes for anonymous credentials,” in *ACM CCS 2008*.
- [19] J. Camenisch, S. Hohenberger, M. Kohlweiss, A. Lysyanskaya, and M. Meyerovich, “How to win the clonewars: Efficient periodic n-times anonymous authentication,” in *ACM CCS 2006*.
- [20] J. Camenisch and A. Lysyanskaya, “An efficient system for non-transferable anonymous credentials with optional anonymity revocation,” in *EUROCRYPT 2001*.
- [21] —, “A signature scheme with efficient protocols,” in *SCN 02*.
- [22] —, “Signature schemes and anonymous credentials from bilinear maps,” in *CRYPTO 2004*.
- [23] J. Camenisch and M. Stadler, “Proof systems for general statements about discrete logarithms,” Tech. Rep., 1997.
- [24] J. Camenisch and E. Van Herreweghen, “Design and implementation of the idemix anonymous credential system,” in *ACM CCS 2002*.
- [25] M. Chase, T. Perrin, and G. Zaverucha, “The signal private group system and anonymous credentials supporting efficient verifiable encryption,” in *ACM CCS 2020*.
- [26] D. Chaum, “Security without identification: Transaction systems to make big brother obsolete,” *Communications of the ACM*, 1985.
- [27] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. P. Ward, “Marlin: Preprocessing zkSNARKs with universal and updatable SRS,” in *EUROCRYPT 2020, Part I*.
- [28] A. Davidson, I. Goldberg, N. Sullivan, G. Tankersley, and F. Valsorda, “Privacy pass: Bypassing internet challenges anonymously,” *PoPETS*, no. 3, Jul. 2018.
- [29] e-Estonia, “ID-card,” <https://e-estonia.com/solutions/e-identity/id-card/>.
- [30] S. Ellis, A. Juels, and S. Nazarov, “Chainlink: A decentralized oracle network,” <https://research.chain.link/whitepaper-v1.pdf>.
- [31] Ethereum Development Docs, “Optimistic rollups,” <https://ethereum.org/en/developers/docs/scaling/optimistic-rollups/>.
- [32] Federal Office for Information Security, “German eID,” <https://www.bsi.bund.de/EN/Topics/ElectrIDDDocuments/German-eID/german-eID.html>.
- [33] G. Fuchsbauer, “Subversion-zero-knowledge SNARKs,” in *PKC 2018, Part I*.
- [34] A. Gabizon, Z. J. Williamson, and O. Ciobotaru, “PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge,” Cryptology ePrint Archive, Report 2019/953, 2019, <https://eprint.iacr.org/2019/953>.
- [35] C. Garman, M. Green, and I. Miers, “Decentralized anonymous credentials,” in *NDSS 2014*.
- [36] Google, “Access age-restricted content & features - Google Account Help,” <https://support.google.com/accounts/answer/10071085?p=age-verify>.
- [37] —, “Watch age-restricted videos,” <https://support.google.com/youtube/answer/10070779>.
- [38] L. Grassi, D. Kales, D. Khovratovich, A. Roy, C. Rechberger, and M. Schofnegger, “Starkad and Poseidon: New hash functions for zero knowledge proof systems,” Cryptology ePrint Archive, Report 2019/458, 2019, <https://eprint.iacr.org/2019/458>.
- [39] J. Groth, “On the size of pairing-based non-interactive arguments,” in *EUROCRYPT 2016, Part II*.
- [40] ICAO, “Machine Readable Travel Documents. Part 10: Logical Data Structure (LDS) for Storage of Biometrics and Other Data in the Contactless Integrated Circuit (IC),” International Civil Aviation Organization (ICAO), Tech. Rep., https://www.icao.int/publications/Documents/9303_p10_cons_en.pdf.
- [41] IDNYC, “About IDNYC,” <https://www1.nyc.gov/site/idnyc/about/about.page>.
- [42] A. Kate, G. M. Zaverucha, and I. Goldberg, “Constant-size commitments to polynomials and their applications,” in *ASIACRYPT 2010*.
- [43] M. Kohlweiss, M. Maller, J. Siim, and M. Volkhov, “Snarky ceremonies,” Cryptology ePrint Archive, Report 2021/219, 2021, <https://eprint.iacr.org/2021/219>.
- [44] C. Komlo and I. Goldberg, “Frost: Flexible round-optimized schnorr threshold signatures,” in *Selected Areas in Cryptography*.
- [45] A. E. Kosba, Z. Zhao, A. K. Miller, Y. Qian, T.-H. H. Chan, C. Papamanthou, and E. Shi, “C0c0: A framework for building composable zero-knowledge proofs,” <https://eprint.iacr.org/2015/1093>.
- [46] J. Kuszmaul, “Verkle trees,” <https://math.mit.edu/research/highschool/primes/materials/2018/Kuszmaul.pdf>, MIT, Tech. Rep., 2018.
- [47] Litviniuk, Anna, “Supportfemale 2 SVG,” CC Attribution License <https://creativecommons.org/licenses/by/3.0/>.
- [48] M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn, “Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings,” in *ACM CCS 2019*.
- [49] J. L. Muñoz-Tapia, M. Belles, M. Isabel, A. Rubio, and J. Baylina, “CIRCOM: A Robust and Scalable Language for Building Complex Zero-Knowledge Circuits,” 2022, https://www.techrxiv.org/articles/preprint/CIRCOM_A_Robust_and_Scalable_Language_for_Building_Complex_Zero-Knowledge_Circuits/19374986.
- [50] New York State Senate, “Find your senator,” <https://www.nysenate.gov/registration/nojs/form/start/find-my-senator>.
- [51] NYC Planning, “Political and Administrative Districts — Download and Metadata,” <https://www1.nyc.gov/site/planning/data-maps/open-data/districts-download-metadata.page>.
- [52] D. Rathee, G. V. Policharla, T. Xie, R. Cottone, and D. Song, “Zebra: Anonymous credentials with practical on-chain verification and applications to kyc in defi,” Cryptology ePrint Archive, Paper 2022/1286, 2022, <https://eprint.iacr.org/2022/1286>.
- [53] ReadID, “Which countries have ePassports?” <https://www.readid.com/blog/countries-epassports>.
- [54] M. Rosenberg, J. White, C. Garman, and I. Miers, “zk-creds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure,” Cryptology ePrint Archive, Paper 2022/878, 2022, <https://eprint.iacr.org/2022/878>.
- [55] A. Sonnino, M. Al-Bassam, S. Bano, S. Meiklejohn, and G. Danezis, “Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers,” in *NDSS 2019*.
- [56] M. A. Specter, S. Park, and M. Green, “KeyForge: Non-attributable email from forward-forgable signatures,” in *USENIX Security 2021*.
- [57] TLSNotary, “Prove an HTTPS page was in your browser,” <https://tlsnotary.org/>.

#leaves	2^{14}	2^{30}	2^{46}	2^{62}
Issuance proof time	141ms	403ms	670ms	962ms

TABLE 4: Cost of a proof of credential list membership as the size of the list varies.

- [58] Tornado Cash, “Core deposit circuit,” <https://docs.tornado.cash/tornado-cash-classic/circuits/core-deposit-circuit>.
- [59] F. H. S. a. M. Vehicles, “2D barcode reader calibration sheet — 2016 AAMVA standard,” https://www.flhsmv.gov/pdf/newdl/fl_2dbarcodecalibration.pdf.
- [60] W3C Recommendation, “Verifiable Credentials Data Model v1.1,” <https://www.w3.org/TR/2022/REC-vc-data-model-20220303/>.
- [61] A. Zapico, V. Buterin, D. Khovratovich, M. Maller, A. Nitulescu, and M. Simkin, “Caulk: Lookup arguments in sublinear time,” Cryptology ePrint Archive, Paper 2022/621, 2022, <https://eprint.iacr.org/2022/621>.
- [62] F. Zhang, D. Maram, H. Malvai, S. Goldfeder, and A. Juels, “Deco: Liberating web data using decentralized oracles for tls,” in *ACM CCS 2020*.

Appendix A. Performance across parameter choices

We expand on the microbenchmarks in Section 6.2 by investigating how `zk-creds` scales with respect to various parameters.

Recall that showing a credential requires proving: 1) the credential is in the list of issued credentials, 2) the relevant attributes are part of the credential, 3) the attributes meet the access criteria being shown, 4) each of the above is about the same data (linkage proof). Thus, the performance of `zk-creds` depends on the number of issued credentials (via 1 above), the size of the attribute (via 2), and the complexity of the access criteria (via 3).

Membership benchmarks. Recall that a membership proof consists of a proof of credential membership in a tree, followed by a proof of membership of that tree in a forest. In Figure 5, we show the performance of proving membership as the shape of the forest changes. For a fixed number of total leaves, we find the size of the forest (and, consequently, height of its trees) that minimizes membership proving time. This results in a 50% (143ms) speedup in the best case. Further, the verifier pays nothing for this optimization, since all public inputs are prepared in advance and reused for all verifications.

In Table 4, we show the time to compute a proof of membership in the credential list as the list size varies. This represents the baseline cost of the issuance portion of any `ShowCred` call. The benchmark consists of one Groth16 proof of tree membership plus one Groth16 proof of forest membership. For a fixed number of leaves, the tree height is chosen using the optimal parameters from the experiments in Figure 5.

Linkage benchmarks. In Figure 4, we plot the size, proving time, and verification time of linkage proofs as the number of standalone gadget proofs varies. Every additional gadget adds 330B to the proof size.

Appendix B. Merkle witness management

A `zk-creds` implementation which prioritizes both efficiency and privacy must carefully consider how it in-

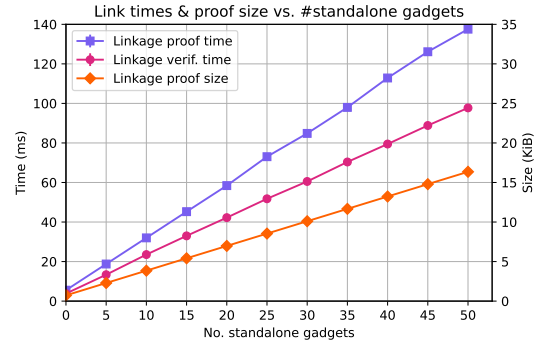


Figure 4: Costs of linkage proofs as #standalone gadgets varies.

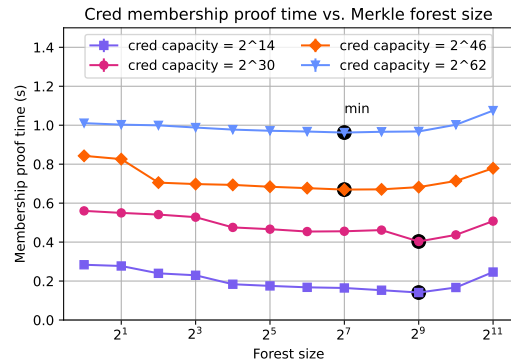


Figure 5: Costs of proving tree then forest membership, as number of trees in forest varies.

stantiates the credential list and, consequently, how it wants its users to practically go about witness management. We introduce two approaches and their tradeoffs for forests of fixed-height Merkle trees. See the full version for more in-depth discussion and a better optimized Merkle forest construction using these techniques.

Assume that credentials are added to a Merkle forest from left to right, and that issued credentials are never removed or modified once added to a tree in the list. Also, suppose there is a *list manager*, perhaps distinct from the issuer, who distributes the issuer’s credential list to users and verifiers.¹⁵

Observe that, if a user has a valid Merkle authentication path θ attesting to their credential’s issuance at time t , not all nodes in θ will usually need updating by time $t' > t$. Let the *frontier* be the subset of information about the list that is necessary for a given user to update their θ . As a first approach, consider that a user can request only the subset of *frontier nodes* in θ which will have updated from time t to t' , instead of the full θ at time t' . However, requesting specific nodes to update θ still reveals identifying information about the credential (reducing the anonymity set) and, by necessity of requests, likely correlates in time with the subsequent usage of that credential.

As a second approach, consider the following techniques: 1) Instead of the user *requesting* specific updates for θ , the list

15. Assuming public auditability, the list manager is only trusted to not DoS the system. If a party suspects that e.g. the authentication path is invalid, they or an auditor can reliably obtain and audit the credential list themselves.

manager proactively *broadcasts* all relevant tree updates to users. This requires $O(\log N)$ nodes per addition to Merkle tree (capacity N), or $O(N \log N)$ bandwidth per user across the Merkle tree's lifetime. With this, any user can parse the streamed nodes to update θ with only $O(\log N)$ space. And, since users no longer need to reveal when they do so, this incurs little to no privacy risk. With no update requests, adversaries can no longer correlate updates with subsequent shows by observing network interactions.

2) Issuers *batch* newly-issued credentials by some epoch. While this works for any list instantiation, for Merkle trees this still only requires $O(\log N)$ frontier nodes to represent an *entire batch's* changes to any previously-issued credential's θ .

Both approaches require (competitive) small-constant logarithmic storage to update a credential's authentication path, with fewer privacy risks than the approaches in Section 5.1.

Appendix C. LinkG16

C.1. Groth16

We describe a trusted-setup zkSNARK scheme, due to Groth¹⁶ [39], which operates over a non-degenerate type-3 bilinear group. We use \mathbb{F} to denote the scalar field of the bilinear group. At a high level, a Groth16 proof proves that an arithmetic circuit over \mathbb{F} is satisfied by a set of public inputs (values known to the verifier) and private inputs (values not known to the verifier).

G16.Setup(bg, desc) \rightarrow crs Generates a common reference string for the given arithmetic circuit description. crs contains the elements from the bilinear group bg necessary to compute the expressions in G16.Prove below.

G16.Prove(crs, $\{a_i\}_{i=0}^\ell, \{a_i\}_{i=\ell+1}^m$) \rightarrow π Proves the circuit described by crs is satisfied, where $a_0, \dots, a_\ell \in \mathbb{F}$ represent the circuit's public input wires and $a_{\ell+1}, \dots, a_m \in \mathbb{F}$ represent the private wires. π is of the form (A, B, C) where $A, C \in \mathbb{G}_1$ and $B \in \mathbb{G}_2$.

G16.Verify(crs, π, \hat{S}) \rightarrow $\{0, 1\}$ Verifies the proof $\pi = (A, B, C)$ with respect to the *prepared public input* $\hat{S} = \sum_0^\ell a_i W_i$ by checking the relation

$$e(A, B) \stackrel{?}{=} e([\alpha]_1, [\beta]_2) \cdot e(C, [\delta]_2) \cdot e(\hat{S}, H),$$

where $[\alpha]_1, [\beta]_2$, and $[\delta]_2$ come from crs, and W_i is the crs value whose coefficient represents the value of the i -th wire in the circuit. As shorthand, verification can also be written as G16.Verify(crs, π, \mathbf{a}).

G16.Rerand(crs, π) \rightarrow π' Rerandomizes the proof $\pi = (A, B, C)$ by sampling $\zeta, \omega \leftarrow \mathbb{F}$ and computing

$$\pi' := (\zeta^{-1}A, \zeta B + \zeta\omega[\delta]_2, C + \omega A).$$

By Theorem 3 in [4], the output of Rerand is perfectly indistinguishable from a fresh proof of the same underlying statement.

16. To support linkage, we diverge slightly from Groth's original construction by setting one of the trapdoor values γ to 1. This does not affect security; zero-knowledge and knowledge soundness were proven by Kohlweiss et al. [43] for a strictly larger CRS which also sets $\gamma := 1$.

We now describe and prove the security of the LinkG16 proof system.

C.2. Goal

The purpose of LinkG16 is to show that k Groth16 proofs over heterogeneous circuits $\text{crs}_1, \dots, \text{crs}_k$ all share the same first t public inputs $\{a_0, \dots, a_{t-1}\}$ *without* revealing the inputs. That is, given k Groth16 proofs π_1, \dots, π_k , we wish to construct a zero-knowledge proof of the following relation:

$$R_{\text{linkG16}} = \left\{ \begin{array}{l} (\{\text{crs}_i, \hat{S}_i\}_{i=1}^k; \{a_j\}_{j=0}^{t-1}, \{\pi_i\}_{i=1}^k) : \\ \bigwedge_{i=1}^k \text{G16.Verify}(\text{crs}_i, \pi_i, \hat{S}_i + \sum_{j=0}^{t-1} a_j W_j^{(i)}) \end{array} \right\}$$

where $W_j^{(i)}$ represents the wire W_j in crs_i , and $\hat{S}_i = \sum_{j=t}^\ell a_j W_j^{(i)}$ is the verifier-known prepared input for the i -th proof.

C.3. Construction

We define the new proof system below.

LinkG16.Link($\{a_j\}_{j=0}^{t-1}, \{\text{crs}_i, \pi_i\}_{i=1}^k$) \rightarrow π_{link} Sample values $z_1, \dots, z_k \leftarrow \mathbb{F}$ for blinding. For each i , commit to the shared inputs, $U_i := z_i[\delta]_1^{(i)} + \sum_{j=0}^{t-1} a_j W_j^{(i)}$. Let π_{eqwire} be an EqWire discrete-log equality proof (described in Appendix D) that the U_i commit to the same a_j values,

$$R_{\text{eqwire}} = \left\{ \begin{array}{l} (\{U_i, \text{crs}_i\}_{i=1}^k; \{a_j\}_{j=0}^{t-1}, \{z_i\}_{i=1}^k) : \\ \bigwedge_{i=1}^k U_i = z_i[\delta]_1^{(i)} + \sum_{j=0}^{t-1} a_j W_j^{(i)} \end{array} \right\}.$$

Rerandomize the underlying proofs in place, $\pi_i :=$ G16.Rerand(crs_i, π_i), then blind the proofs,

$$\pi'_i := (A_i, B_i, C_i - [z_i]_1).$$

The final output is

$$\pi_{\text{link}} := (\pi_{\text{eqwire}}, \{U_i, \pi'_i\}_{i=1}^k).$$

LinkG16.LinkVerify($\pi_{\text{link}}, \{\text{crs}_i, \hat{S}_i\}_{i=1}^k$) \rightarrow $\{0, 1\}$ Check π_{eqwire} using EqWire.Verify. Then unpack each π'_i into (A'_i, B'_i, C'_i) . For each $i = 1, \dots, k$, check

$$e(A'_i, B'_i) \stackrel{?}{=} e([\alpha]_1^{(i)}, [\beta]_2^{(i)}) \cdot e(C'_i, [\delta]_2^{(i)}) \cdot e(U_i + \hat{S}_i, H).$$

where \hat{S}_i is the Groth16 prepared public input for circuit i .

C.4. Proofs

Theorem 1 (Correctness). *If G16.Prove and LinkG16.Link are honestly computed, then LinkG16.LinkVerify succeeds.*

Proof. We show that the LinkVerify equation above holds for all i . For legibility, we omit the index i in the proof. Suppose π_{link} is computed honestly, i.e., that all U' and (A', B', C') are well-formed and that the underlying Groth16 verification

equations holds on the corresponding (A, B, C) . First, we note that, since C' and U were computed honestly,

$$\begin{aligned} & e(C', [\delta]_2) \cdot e(U, H) \\ &= e(C - [z]_1, [\delta]_2) \cdot e(\sum a_j W_j + z[\delta]_1, H) \\ &= e(C, [\delta]_2) \cdot e(-[z]_1, [\delta]_2) \cdot e(z[\delta]_1, H) \cdot e(\sum a_j W_j, H) \\ &= e(C, [\delta]_2) \cdot e(\sum a_j W_j, H). \end{aligned}$$

Using this and the fact that $A' = A$ and $B' = B$, we see that the LinkVerify equation

$$e(A', B') = e([\alpha]_1, [\beta]_2) \cdot e(C', [\delta]_2) \cdot e(U + \hat{S}, H)$$

holds if and only if

$$\begin{aligned} & e(A, B) \\ &= e([\alpha]_1, [\beta]_2) \cdot e(C', [\delta]_2) \cdot e(U + \hat{S}, H) \\ &= e([\alpha]_1, [\beta]_2) \cdot e(C', [\delta]_2) \cdot e(U, H) \cdot e(\hat{S}, H) \\ &= e([\alpha]_1, [\beta]_2) \cdot e(C, [\delta]_2) \cdot e(\sum a_j W_j, H) \cdot e(\hat{S}, H) \\ &= e([\alpha]_1, [\beta]_2) \cdot e(C, [\delta]_2) \cdot e(\hat{S} + \sum a_j W_j, H) \end{aligned}$$

which is precisely the verification equation for the i -th underlying Groth16 instance. Since this equation holds by assumption, LinkVerify succeeds. \square

Theorem 2. LinkG16 is perfect HVZK.

Proof. We define a simulator $\text{Sim}_{\text{linkg16}}$ with access to Groth16 trapdoors τ_1, \dots, τ_k as follows. For each i , sample $\bar{A}'_i, \bar{B}'_i, \bar{U}_i \leftarrow \mathbb{F}$ uniformly. Use the trapdoor τ_i to compute $C'_i \in \mathbb{G}_1$ as the unique group element which satisfies the i -th LinkVerify equation with respect to crs_i and public inputs¹⁷ $\{a_j\}_{j=t}^\ell$. Concretely,

$$\bar{C}'_i := \frac{\bar{A}'_i \bar{B}'_i - \alpha^{(i)} \beta^{(i)} - \bar{U}_i - \sum_{j=t}^\ell a_j W_j^{(i)}}{\delta^{(i)}}.$$

For all i , let $\pi'_i := ([\bar{A}'_i]_1, [\bar{B}'_i]_2, [\bar{C}'_i]_1)$ and $U_i := [\bar{U}_i]_1$. Finally, let

$$\pi_{\text{eqwire}} \leftarrow \text{Sim}_{\text{eqwire}}(\{\text{crs}_i, U_i\}_{i=1}^k).$$

The output of $\text{Sim}_{\text{linkg16}}$ is $(\pi_{\text{eqwire}}, \{U_i, \pi'_i\}_{i=1}^k)$.

This is indistinguishable from the real world protocol. In the real world: each U_i is uniformly distributed due to the blinding values z_i ; A'_i, B'_i are uniformly distributed by the Groth16 proof procedure; and each C'_i is the unique group element which satisfies the i -th LinkVerify equation. Lastly, the simulated π_{eqwire} is indistinguishable from an honestly generated one due to the perfect HVZK of the EqWire protocol. \square

Theorem 3. LinkG16 is knowledge-sound.

Proof. We define an extractor E_{linkg16} , aborting on verification error, as follows. By knowledge soundness of EqWire there exists an extractor E_{eqwire} which extracts

17. How does the verifier know a_j (for $j \geq t$) if it was only given the prepared input \hat{S} ? It is merely for brevity that LinkVerify is written to take \hat{S} . The verifier always knows the prepared input's constituent a_j values.

$\{a_j\}_{j=0}^{t-1}, \{z_i\}_{i=1}^k$ such that $U_i = z_i[\delta]_1^{(i)} + \sum a_j W_j^{(i)}$. For each $i = 1, \dots, k$, E_{linkg16} then reconstructs the underlying Groth16 proof

$$\pi_i = (A'_i, B'_i, C'_i + [z_i]_1).$$

E_{linkg16} outputs $(\{a_j\}_{j=0}^{t-1}, \{\pi_i\}_{i=1}^k)$. Since E_{linkg16} did not abort, it is the case that, for each i ,

$$\begin{aligned} & e(A'_i, B'_i) \\ &= e([\alpha]_1^{(i)}, [\beta]_2^{(i)}) \cdot e(C'_i, [\delta]_2^{(i)}) \cdot e(U_i + \hat{S}_i, H) \\ &= e([\alpha]_1^{(i)}, [\beta]_2^{(i)}) \cdot e(C'_i, [\delta]_2^{(i)}) \cdot e(z_i[\delta]_1^{(i)} + \sum a_j W_j^{(i)} + \hat{S}_i, H) \\ &= e([\alpha]_1^{(i)}, [\beta]_2^{(i)}) \cdot e(C'_i + [z_i]_1, [\delta]_2^{(i)}) \cdot e(\hat{S}_i + \sum a_j W_j^{(i)}, H) \end{aligned}$$

which is precisely the verification equation for π_i . \square

Appendix D. EqWire

We now define and prove secure a proof system for the discrete-logarithm equality relation,

$$R_{\text{eqwire}} = \left\{ \left(\{U_i, \text{crs}_i\}_{i=1}^k; \{a_j\}_{j=0}^{t-1}, \{z_i\}_{i=1}^k \right) : \bigwedge_{i=1}^k U_i = z_i[\delta]_1^{(i)} + \sum_{j=0}^{t-1} a_j W_j^{(i)} \right\}.$$

The proof system is instantiated using a proof framework due to Camenisch and Stadler [23]. Concretely, it is the Fiat-Shamir transform of the protocol described in Figure 7.

Proofs

Theorem 4. The EqWire protocol is knowledge-sound.

Proof. We define extraction in the usual way for Camenisch-Stadler sigma protocols. Let E_{eqwire} be our extractor, aborting on verification failure. The extractor receives the commitments, and then picks challenge $c \leftarrow \mathbb{F}$. It sends c and receives $\{\rho_j, \sigma_i\}_{i,j}$. The extractor then rewinds to pick a fresh $c' \leftarrow \mathbb{F}$. It sends c' and receives $\{\rho'_j, \sigma'_i\}_{i,j}$. For all i and j , the extractor computes

$$a_j := \frac{\rho_j - \rho'_j}{c' - c} \quad z_i := \frac{\sigma_i - \sigma'_i}{c' - c}$$

and outputs $(\{a_j\}_{j=0}^{t-1}, \{z_i\}_{i=1}^k)$. Since the extractor did not abort, i.e., both runs passed verification, and the commitments did not change, it is the case that $U_i = z_i[\delta]_1^{(i)} + \sum_j a_j W_j^{(i)}$ for all i . \square

Theorem 5. The EqWire protocol is perfect HVZK.

Proof. We define a simulator as follows: sample c and all σ_i, ρ_j uniformly from \mathbb{F} ; for all i , compute $\text{com}_i := \sum_j \rho_j W_j^{(i)} + \sigma_i[\delta]_1^{(i)} + cU_i$; output $(c, \{\text{com}_i, \rho_j, \sigma_i\}_{i,j})$.

This is perfectly indistinguishable from a real transcript: all σ_i and ρ_j are uniform iid since they are blinded by s_i and r_j , respectively; c is uniform and independent by definition of honest-verifier; and all com_i are uniquely determined by these values. In the simulator, each σ_i and ρ_j is uniform iid by construction, and com_i is uniquely determined by these values. \square

<p>\mathcal{F}.IssueSetup():</p> <ol style="list-style-type: none"> 1. Decide on relevant issuance criteria ι. 2. Let $\text{IssueCriteria}[I] := \text{IssueCriteria}[I] \cup \{\iota\}$ 3. Publish (I, ι) <p>\mathcal{F}.IssueReq$_U(\iota, \text{attrs}, w_{sd}, \text{iaux}_{sd})$:</p> <ol style="list-style-type: none"> 1. Sample random cred, r // commitment and its opening 2. $\text{UserCreds}_U[\text{cred}] := (\iota, r, \text{attrs})$ // construct credential 3. Let $\text{HiddenZKSD}_U[r] := w_{sd}$ 4. Send $(\text{cred}, \text{iaux}_{sd})$ to \mathcal{F}.IssueGrant 5. If the previous step aborts or returns \perp, abort 6. Else, it returns θ: send (cred, r, θ) to U <p>\mathcal{F}.IssueGrant$_I(\iota, \text{cred}, \text{iaux}_{sd})$:</p> <ol style="list-style-type: none"> 1. If issuer I is honest and user U is corrupted: <ol style="list-style-type: none"> (a) Let $(\iota', r, \text{attrs}) := \text{UserCreds}_U[\text{cred}]$ (b) Let $w_{sd} := \text{HiddenZKSD}_U[r]$ (c) Check that $\iota = \iota'$ and $\iota \in \text{IssueCriteria}[I]$ (d) Check that $\iota(\text{attrs}, \text{iaux}_{sd}, w_{sd}) = 1$ (e) If any check fails: send \perp to U and abort 2. Arbitrarily, I may choose to deny cred; if so, abort 3. Sample random θ 4. Let $\text{IssuedCreds}[\theta] := \text{cred}$ 5. Send updated IssuedCreds to I 6. Send θ to U 7. Notify all parties that cred has been issued 	<p>\mathcal{F}.ShowSetup():</p> <ol style="list-style-type: none"> 1. Decide on relevant access criteria ϕ. 2. Let $\text{AccessCriteria}[V] := \text{AccessCriteria}[V] \cup \{\phi\}$ 3. Publish (V, ϕ) <p>\mathcal{F}.ShowCred$_U(\phi, \text{cred}, \theta, r, \text{aux})$:</p> <ol style="list-style-type: none"> 1. Sample random s 2. Let $(\cdot, r', \text{attrs}) := \text{UserCreds}_U[\text{cred}]$ 3. Check $r = r'$ and abort if it fails 4. If user U is honest: let $\phi' := \phi$ and $\theta' := \theta$ 5. Else, if U is corrupted: let ϕ', θ' be arbitrary 6. Let $\text{ShowProofs}_U[s] := (\phi', \text{cred}, \theta', r, \text{aux})$ 7. Send (s, aux) to user U <p>\mathcal{F}.VerifyShow$_V(\phi, s, \text{aux})$:</p> <ol style="list-style-type: none"> 1. Let $(\phi', \text{cred}, \theta, r, \text{aux}') := \text{ShowProofs}_U[s]$ 2. Let $\text{cred}' := \text{IssuedCreds}[\theta]$ 3. Let $(\cdot, r', \text{attrs}) := \text{UserCreds}_U[\text{cred}]$ 4. Check that $\phi = \phi'$ and $\phi \in \text{AccessCriteria}[V]$ 5. Check that $\text{cred} = \text{cred}' \neq \text{nil}$, $r = r'$, and $\text{aux} = \text{aux}'$ 6. If verifier V is honest and user U is corrupted: <ol style="list-style-type: none"> check $\phi(\text{attrs}, \text{aux}) = 1$ 7. If any check fails: send false to V and abort 8. Arbitrarily, V may choose to deny cred; if so, abort 9. Else, send true to V <p>\mathcal{F}.RevokeCred$_I(\text{cred})$:</p> <ol style="list-style-type: none"> 1. Find index θ such that $\text{IssuedCreds}[\theta] = \text{cred}$; else, abort 2. Let $\text{IssuedCreds}[\theta] := \text{nil}$ 3. Notify all parties that cred has been revoked
---	---

Figure 6: An ideal functionality \mathcal{F} for zk-creds.

	IssueReq	IssueGrant	ShowCred	ShowCred (full)	VerifyShow
Age-restricted vid.	2.36s	2ms	258ms	1.05s	8ms
Entering a bar	2.36s	2ms	228ms	1.01s	6ms

TABLE 5: zk-creds case study benchmarks using Pedersen hashes. IssueReq is the time to convert a passport into a credential using zk-supporting-documentation. ShowCred is the time it takes a user to prove they are over 18.

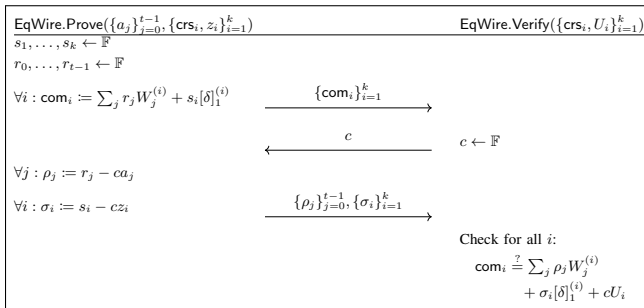


Figure 7: The EqWire protocol

Appendix G. CRS sizes of evaluated circuits

We list in Table 6 the sizes of the proving and verifying keys for all the circuits evaluated in Sections 6 and 7.

	Proving key (MB)	Verifying key (KB)
Forest membership	0.3	12.8
Tree membership	10.1	0.6
Expiry	1.5	0.6
Linkable show	0.4	0.6
Rate limiting	1.3	0.8
Clone resistance	1.2	0.8
Age-restricted vid.	2.8	0.9
Entering a bar	2.5	0.6

TABLE 6: CRS sizes for our evaluated circuits

Appendix E. Security Definition

Our security definitions are given as an ideal functionality in Figure 6.

Appendix F. Instantiation with Pedersen hashing

In Table 5, we give benchmarks for an instantiation using Pedersen (rather than Poseidon) hashes and commitments. These increase proving times but, atypically for hash functions, are provably secure (in this case, under the discrete log assumption).