SoK: Anonymous Credentials for Digital Identity Wallets

Christian Bormann[1] and Anja Lehmann[1,2]

[1] SPRIN-D - Federal Agency for Breakthrough Innovation, Germany
[2] Hasso Plattner Institute, University of Potsdam, Germany
christian.bormann@eudi.sprind.org, anja.lehmann@hpi.de

Abstract. Digital identity wallets are currently being developed around the globe, aiming to provide user-centric and secure authentication. Realizing this in a privacy-preserving manner is paramount, and even mandated in Europe which is developing the European Digital Identity Wallet with planned release in 2026. Current proposals to build these wallets are based on classic signature schemes such as ECDSA, but would benefit greatly from the use of *anonymous credentials*. Thus, there is currently a strong interest in developing the necessary standards to bring these cryptographic concepts into the real world. This work aims to inform ongoing standardization efforts by providing an overview of the most prominent solutions, and the remaining open challenges. We split our overview among two fundamental architectural approaches: (1) dedicated multi-message signature schemes that allow for efficient ZKPs, and (2) general-purpose ZKPs used on top of legacy ECDSA. We also provide a comprehensive summary of the broad feature set that anonymous credentials can provide for identity wallets, in order to demonstrate that upgrading to these systems is a worthwhile endeavor and help to design standards that can leverage the rich existing body of work.

Keywords: Anonymous Credentials · EUDI · Zero-knowledge Proofs

1 Introduction

Anonymous credentials (AC) enable privacy-preserving and user-centric authentication. Trusted issuers attest the user's attributes in form of cryptographic credentials. The user can then use these credentials to create presentations towards relying parties (RP) that reveal only the minimally necessary subset of the attested information in a verifiable manner. Despite being well-understood in the academic community with research spanning over more than two decades [27,6,13,24,89,49], they have not seen wide-spread adoption yet. However, they are currently enjoying increased attention in the context of several digital identity systems being built around the globe, e.g., in the US [91], Europe [44], Switzerland [88], and the UK [78]. Notably, both the IRTF [63,84] and ISO [56] are currently working on standards for anonymous credentials.

Particularly noteworthy are the efforts in Europe, that aims to roll out a phone-based solution – the European Digital Identity (EUDI) wallet – by the

end of 2026. What stands out in these efforts is the strong commitment to *privacy*: the underlying eIDAS regulation mandates several strong privacy properties such as selective disclosure, pseudonymity, untraceability and unlinkability, e.g. stating that "*The technical framework of the European Digital Identity Wallet shall (a) not allow providers of electronic attestations of attributes or any other party [..] to obtain data that allows transactions or user behaviour to be tracked, linked or correlated [..]; (b) enable privacy preserving techniques which ensure unlinkability*" (Article 5a, 16b of eIDAS [45]).

Despite being *the* use case anonymous credentials were invented for, the design currently proposed in the EU Architecture and Reference Framework (ARF) [41] instead relies on classical signature schemes such as ECDSA. To achieve a weak form of unlinkability and selective disclosure, it uses batch-issuance of one-time use credentials and salted hashes.

While the proposed solution achieves some of the privacy properties mandated by the regulation, it clearly cannot achieve unlinkability when the issuer and RP collude. Unlinkability across different malicious RPs requires using a fresh one-time credential in every presentation, which imposes a significant load on the issuer's infrastructure and also raises challenges in the user's secure key management. This has been highlighted in the cryptographers' feedback [10] for the ARF, advocating for the use of ACs instead.

Given recent advances and better alignment of constructions with the requirements of eIDAS, the EU is now re-considering the use of anonymous credentials and started gathering requirements for the development of a standard [41]. In fact, a core aspect that hindered the adoption of ACs so far is the lack of a unified standard, and any efforts made in the EU will also largely impact similar deployment endeavors in other areas. Thus, it is important to inform decision makers and standardization bodies on the available solution space, their individual challenges and benefits, such that this seminal step towards the real-world deployment of anonymous credentials is done in the best way possible.

Our Contributions. Our work aims to provide this informational foundation and gives an overview of the two most prominent directions to build identity systems from anonymous credentials: dedicated multi-message signatures with efficient proofs (such as BBS signatures with Schnorr proofs [6]), and general-purpose zero-knowledge (ZKP) systems build on top of ECDSA. We describe their state-of the art, status in terms of standardization and challenges they need to overcome in order to be suitable for digital identity wallets. For the first approach relying on dedicated and mature signatures such as BBS [6], we focus on the most pressing challenges: compatibility with existing user phones for the feature of device binding, as well as the support for pairing-friendly curves. When it comes to general-purpose ZKP systems, the situation is much more open and almost complementary: they can be built on top of any legacy signature scheme and even entire classical credential systems, avoiding the deployment challenges inherent in the BBS-based approach. However, they are still a highly active field of research, with a multitude of options and introduce a level of complexity

through their circuit-based design that the ZKProof work [95] aims to tame for almost 8 years already.

Before giving an overview of these technical solutions, we summarize the features of anonymous credentials that we believe are of particular interest in the digital identity context. In fact, the discussion so far focuses solely on the most basic feature of ACs – multi-show unlinkability – but does not consider the bigger picture and feature set they enable. Without being aware of the overall feature set, there is a risk that technical standards are chosen and built in a premature way, hard-coding the design choices to a bare bone solution that will be hard to extend. Apart from being relevant for properly designing the technical standards, they also demonstrate that anonymous credentials will bring significant benefits compared to the existing classical signature approach for digital identity wallets: not only in privacy but also in security and large scale deployability.

The goal to inform the ongoing standardization activities in the EU around ACs gives this SoK a more "executive-level" flavour. A more technical and full-fledged overview of anonymous credentials has been attempted over a decade ago in the ABC4Trust project, and our SoK can be seen as an updated and wallet-centric overview of such technologies [13]. Further, the focus on the use of anonymous credentials for digital identity wallets and the most pressing challenges to overcome is also what separates our work from the more general SoK on anonymous credentials by Kakvi et al. [58].

2 Features of Anonymous Credentials

This section gives an overview of the general privacy and security features that anonymous credentials (ACs) can provide in the context of digital identity systems. We focus mainly on features where they differ from the current approach of batch-issued classic certificates as described in the EUDI ARF [41]. Note that they not only provide additional privacy properties but also enable better security and scalability than the currently designed system.

Core Features. We start with the fundamental features that are supported by most anonymous credential systems. The systems entails three main entities: issuers, users and relying parties (RP). The issuers attest attributes in form of credentials to the user. The user can then derive presentations based on these credentials to transmit authenticated attribute-based information to the relying party which is able to verify the authenticity of those statements.

Selective Disclosure: With every presentation, the user can choose the subset of attributes they want to selectively disclose towards the RP.

Multi-show Unlinkability: Relying party to relying party unlinkability ensures that different presentations of the same user, that don't reveal any identifying information, cannot be linked. This prevents profiling of users in scenarios where only limited information about their credentials is revealed, e.g., in age proofs or proofs of personhood.

Untraceability: Relying party to issuer unlinkability prevents that a malicious issuer can link the credentials it issued to the derived presentations the user sends to the relying parties. This again aims to prevent any profiling and tracking through malicious relying parties and issuers.

Single Key Binding: Many or even all credentials that belong to the same user can be bound to the same user private key. Such single key binding reduces complexity for the user/device and increases security. When all of the user's credentials are bound to the same key, this disincentivizes the sharing or pooling of credentials through malicious users, and is often referred to as all-or-nothing transferability.

Presentation Freshness: Presentation of a credential never reveals the original credential, but only a proof thereof. The proof can be bound to a particular context, e.g., the session nonce or the RP's identity towards which the presentation is done. From such a presentation, no further presentation for a different context can be derived. This ensures that certified attributes cannot be used in a context other than what was intended by the user.

Certified Pseudonyms: Anonymous credentials do not come with the classical notion of public keys for the user. Instead, they allow to derive verifiable yet unlinkable pseudonyms from a user's secret key embedded in her credentials. From a single key, many unlinkable pseudonyms can be verifiably derived and bound to attested user attributes. Pseudonyms can optionally be context-bound, e.g., realizing a 1-to-1 binding of a user to an RP or enforcing rate-limiting, while keeping her pseudonyms unlinkable across RPs/contexts. This introduces a controlled level of linkability back into the system.

Further Features. Beyond the basic features summarized above, ACs can provide more capabilities that are impossible with classic certificates. Most are designed to further improve privacy and security by minimizing the information that users have to expose during issuance and presentation.

Predicate Proofs: Instead of revealing an exact attribute, the presentation proves that a specific property of or statement about an attribute holds true. A prominent example is the age over 18 proof instead of revealing the date of birth. This feature allows for more flexibility during presentation (issuers do not need to know the concrete requests) and also helps to hide correlation factors, such as expiration date, in a presentation.

Composite Proofs: Composite proofs (presentations) can be constructed over several credentials, also allowing more advanced logic over the different attributes. For instance, a composite proof could demonstrate that a user possesses both, a valid university credential, and a government ID with a matching name, all without revealing the name. Composite proofs also enable to separate between user and device binding: credentials can be bound to a user secret, which is never revealed, but be used to prove that several credentials belong together. When needed, a special device-bound credential that blindly connects the user secret with the device secret can further es-

tablish a proof of device-binding. This will allow for a simpler backup and restore, as the attribute credentials are not directly bound to a user device.

Deniable Presentations: Presentations can be bound to a specific relying party, and constructed such that only this RP will be convinced of the correctness of the user's attributes (also called designated verifier proof). If the RP forwards the proof to a third party, this does not provide convincing cryptographic evidence of the validity of the user's attributes. This capability would allow for switching between presentations that offer repudiation or non-repudiation based on use-case and regulatory requirements.

(Partially-)Blind Issuance: Allows for the issuance of data that stays hidden from the issuer. This enables distributed issuance concepts, e.g., for a central issuance authority for academic credentials. The sensitive data, such as students' names and grades, is checked and vouched for by a local university, but remains hidden to the centrally provided issuance infrastructure, which only sees the university and final degree, but signs the whole document.

Issuer Hiding: Creates a presentation that only reveals membership in a certain trust ecosystem instead of revealing a specific issuer. For example, proving that one owns a master's degree from a German university, without revealing the specific university.

Conditional Disclosure: The presentation contains verified but initially protected attributes, which can be discovered by a dedicated party only when needed. For instance, making the personal address only accessible to the postal service, not the online shop, or the credit card number used in a hotel reservation that is can only be revealed in case of a no-show.

Benefits of Anonymous Credentials. Discussions about ACs are often focused solely on the privacy features, but in the case of the EUDI Wallets, they also bring deployment advantages that we summarize next.

Enhanced Privacy and Security. Solutions based on ACs achieve stronger privacy and security guarantees, which can lead to better user acceptance.
(1) Relying on ACs is the only practical way to fully achieve the unlinkability requirement specified in ğ16 of the eIDAS regulation. Realizing this requirement and incorporating modern cryptography sets a high standard for digital wallets and lessens the risk that users will face and perceive when using the technology. This holds particularly in unregulated use-cases where a certain level of trust in the issuer not acting maliciously cannot be assumed.
(2) Classic certificates reveal the original attribute credential with every presentation, which harms the actual and perceived security of the system. Average users will not understand the difference between a credential with and without a fresh proof-of-possession (PoP), but might consider the system insecure when copies of their credentials can be made easily. This is also an actual risk: RPs might accept credentials without a fresh PoP, e.g., due to false configuration or the intent to act outside of the regulated EUDI ecosystem.

With ACs, every attribute presentation has an implicit freshness proof and cannot be reused outside the original context. Further, to verify the user at-

tributes, the context string must be given as input for verification as well. This typically contains the targeted RP's identity and thus allows for immediate tracing of the faulty or malicious RP that has lost/reused the users' credentials.

Reduced System Complexity. ACs can help to reduce the system's complexity, both for issuers and users, for instance, as follows:
(3) The currently proposed EUDI solution achieves multi-show unlinkability at the cost of system complexity, via single-use credentials. In contrast, this is achieved naturally for ACs, as they allow to derive many unlinkable presentations from the same credential. By getting rid of one-time use credentials, batch issuance, and the need to manage possibly hundreds of device keys, the overall system becomes more robust, and the expected load is reduced significantly.
(4) Composite proofs enable to decouple the device-binding aspect from attribute credentials, which leads to a simpler backup as described above. They also enable an efficient way to handle short-lived credentials [13]: instead of regularly updating all of the user's credentials, they can all be bound to a dedicated revocation credential. Only the revocation credential must get refreshed regularly, and binding is proven in a privacy-preserving way.
(5) Issuance can be outsourced without the risk of leaking data or introducing traceability, e.g., for central – yet partially blind – issuance for use-cases like medical or educational credentials. This has significant benefits for settings where the actual issuance infrastructure is provided by an external service provider.

3 Anonymous Credential Options

Generally, anonymous credentials always follow the same blueprint: (1) at issuance, the user obtains a credential, which is an issuer's signature on a set of attributes; and (2) at presentation, the user proves knowledge (via a ZKP) of such an issuer's signature instead of revealing it directly. The relying party then verifies that the proof is valid for the issuer's public key. The proof system used in the presentation must be *sound* and *zero-knowledge*. Soundness ensures that a malicious prover cannot convince an honest verifier of an incorrect statement; in this context, it guarantees that users can only prove possession of a credential or attested attributes contained in a credential that was genuinely issued, preventing forgery. The zero-knowledge property guarantees that the verifier does not learn anything beyond the validity of the proven statement. For every presentation, the user creates a fresh ZKP of their credential, which yields the desired multi-show unlinkability and untraceability. Concrete protocols for such constructions can be grouped into two categories:

Multi-message signature schemes with efficient proofs: These are digital signature schemes that are designed to natively support selective disclosure and efficient proofs of knowledge of a signature. This is the most simple and efficient way to build anonymous credentials, but requires a dedicated and non-standard signature scheme for the issuers.

General-purpose ZKPs: General-purpose ZKPs are a class of ZKP protocols that allow proving statements over arbitrary computations. The core advantage of anonymous credentials using such ZKPs is that they can be built upon legacy systems and standard signatures. However, they are several orders of magnitude more complex and less efficient than systems built from dedicated signature schemes.

The following section provides an overview of both approaches and the most promising protocols therein. The description focuses on the maturity and security, complexity, and hardware support for secure device binding.

3.1 Multi-Message Signatures with Efficient Proofs

Multi-message signatures refer to a class of signature schemes with two main properties built-in: (i) the signature scheme allows to sign several messages (i.e., the user attributes) *individually*, which natively enables selective disclosure, and (ii) possess a cryptographic structure that allows to efficiently prove knowledge of a signature through a ZKP (typically by avoiding that messages get hashed).

Currently, the most popular multi-message signature scheme supporting ACs is BBS [17]. Follow-up works from the initial signature scheme have provided formal security proofs, support for multiple attributes, and improved the zero-knowledge proof of a signature [6,24,89]. Other multi-message signature schemes are the RSA-based CL-signatures [27], pairing-based CL-signatures [28], and PS signatures [77]. The ZKP used with BBS is essentially a non-interactive Schnorr-proof of the required discrete-logarithm statements, with an efficiency comparable to classical DL-based signature. That is, proof generation and verification take a few milliseconds, and the proofs have a size of several hundred bytes.

Maturity and Security. BBS signatures are currently the most studied and widely used multi-message signature scheme, with more than two decades of research and analysis, and thus were recommended in the cryptographers' feedback to the EUDI ARF [10]. BBS signatures have already been ISO standardized in ISO/IEC 20008-2:2013 [55] (the amendment from 2023 now also includes PS signatures), and the updated variant is currently undergoing standardization in the IRTF CFRG [63]. All aforementioned features from pseudonyms [60], deniability [38], to partially-blind issuance [59] and conditional disclosure [29] to issuer hiding [16] have already been proposed and are easily compatible and composable through Schnorr-type proof statements [13] – see the paragraph on further features and extendability below. Anonymous credentials have already seen real-world deployment through the TPM DAA and SGX EPID protocols, and the pairing-free version (see discussion below) is used, e.g., by Signal [34], and has recently been considered by Apple [94] and Google [86] to replace batch-issuance based token implementations such as PrivacyPass.

There are two main challenges related to the real-world deployment of BBS: the reliance on **pairing-friendly elliptic curves** and how to use Secure Elements (SE) for **device binding**. We sketch possible solutions for both.

Device Binding. For device binding, the BBS credential includes a device public key, and every presentation must include a fresh proof-of-possession (PoP) under that key. Device binding aims at non-transferability of credentials as the device secret key is stored at a particularly secured hardware component – the Secure Element (SE), and has been raised as the main open challenge to overcome in the response to the cryptographers' feedback [82]. There are different options for how this PoP can be realized, with different trade-offs.

BBS with Native Device Binding: The most established way for native BBS device binding is the Direct Anonymous Attestation (DAA) protocol [21,23], which essentially requires a simple Schnorr signature (in $\mathbb{G}_1$ of the pairing-friendly curve used for BBS) from the SE and only a single exponentiation per presentation. That is, SE does *not* have to compute the full BBS ZKP, but only the PoP part of it, which takes a few milliseconds. The DAA protocol was ISO standardized in 2013 [55], and is supported by the TPM2.0 standard and Intel's SGX attestation [20]. A recent work by Hesse et al. [54] proposes a slightly different solution that requires only a single call – instead of two as in DAA – to the SE, by relying on an additional secret key that is shared by the SE and host. The SE's contribution is again a variation of a Schnorr signature (in $\mathbb{G}_1$). Another alternative is described in the BBS♯ protocol [40], which also requires only a single call to the SE (for either an ECDSA or a Schnorr signature, both in $\mathbb{G}_1$) but does so without a shared secret. The protocol has no formal security analysis yet, though. The most recent work on native device binding for BBS credentials [48] further improves on that, as it requires only access to a plain BLS-signature interface on the SE, i.e., requires only a single call and no state. It is also the first protocol for BBS signatures that does not require trusting the SE for unlinkability and comes with a privacy-preserving variant for cloud-based settings.

While being highly efficient, the main hurdle of such native device binding is that it requires the SE to perform computations in $\mathbb{G}_1$ (or $\mathbb{G}_2$) of a pairing-friendly curve, whereas current SEs only support P-256. While the necessary API for this approach is very simple and lightweight, it *does* require changes in the SEs and proper re-certification, and thus does not work for immediate deployment. However, several countries, e.g., Germany [80], the Netherlands [79], and Sweden [81], aim for a cloud-based solution for device binding for credentials that require a high level of assurance, i.e., the PoP is executed by an external cloud HSM. In this setting, there are fewer limitations regarding the supported cryptography, and native device-binding protocols such as [48] can be integrated with less hurdles.

Note on BBS♯ in the pairing-free setting. If the device binding protocol from BBS♯ is used with the *pairing-free* variant (see discussion below), then their device binding is compatible with existing SEs through a standard ECDSA interface. However, switching to a pairing-free version inherently comes with additional complexity, e.g., in the form of batch-issuance.

BBS + ECDSA: The main challenge with the native device binding is that it is not compatible with existing SEs on the market. Essentially, the only interface available thereon is to create an ECDSA signature over P-256. However, using more advanced cryptography, it is possible to combine a BBS credential (in a pairing-friendly curve) to a standard ECDSA PoP (in P-256). The BBS credential then – in addition to the user's attributes – contains the ECDSA public key *dpk* as an attribute, encoded through the x/y-coordinates of the public key. The SE holds the standard ECDSA secret key *dsk* and is used to compute a standard ECDSA signature σ_{PoP} for a fresh session nonce. However, neither the (full) PoP signature nor the device public key get revealed in a presentation. Instead, the client computes a zero-knowledge proof π_{PoP} that the user has a valid ECDSA signature σ_{PoP} for the session nonce and under the public key certified *dpk* in the BBS credential. The bridging to the BBS credential is done through commitments: the main BBS ZKP proves that it contains a *dpk* that is only revealed in freshly committed form com_{dpk} and π_{PoP} proves that the user knows a signature under that committed key.

The proof of π_{PoP} can be done through either a dedicated proof protocol such as CDLS [30] (building upon ZKAttest [46]) – relying entirely on classical Schnorr proofs – or an optimized zkSNARK, e.g., based on Google's Longfellow [49], Microsoft's Crescent [74] or the protocol from Woo et al. [93]. A BBS-ECDSA hybrid based on ZKAttest/CDLS has already been implemented by Ubique [2] in the context of an innovation challenge run by the German EUDI project. Their implementation currently takes around 800ms to generate proofs, which have a size of 150KB. Thus, while it is more costly than native device binding, it is practical enough for online authentication.

Comparison to Full-fledged SNARK. Compared to computing a zkSNARK over an ECDSA attribute-credential (and ECDSA-PoP), e.g., [49,74,93] the BBS-ECDSA approach has several advantages: zkSNARKs that are efficient enough for real-world deployment currently require highly-optimized circuits and proof systems. Generic approaches are at least 1-2 magnitudes slower. This makes zkSNARKs in the short/midterm mostly applicable for scenarios where hardcoding the circuit and proof statement does not limit the flexibility and extendability of the system. This is the case here, as the statement to be proven via the zkSNARK is independent of the attribute statements proven on top, which is done via a simple and flexible Schnorr proof. Further, as the ECDSA PoP signature is an *one-time* signature computed for a fresh nonce, the signed message as well as parts of the ECDSA signature (the r value) can be revealed in the zero-knowledge proof, which significantly reduces the size and complexity of the proof. In particular, no statement over a hash function must be proven, which is often the most complex part within the circuit.

Status of Pairings. BBS, as well as all other highly efficient multi-show unlinkable anonymous credential schemes, relies on a pairing-friendly curve (with a type-3 pairing). An exception are the RSA-based CL signatures [27], which

are less efficient than their pairing-based variants and thus have seen only little attention since the introduction of pairing-based cryptography in 2001.

A pairing-friendly curve is a special elliptic curve that enables to compute bilinear maps among different groups. While the cryptography is well understood [67], pairing-friendly curves are not widely deployed. Some pairing curves, such as BN254 have been standardized through ISO and the TCG (in the context of DAA), but the currently most popular curve BLS12-381 lacks a formal standard yet. The curve is fully specified through the IRTF drafts [84] and [47] though. There also exists a NIST report from 2015 on pairing-based cryptography [67] providing an overview, application scenarios, and challenges.

Pairing-based cryptography has seen some real-world adoption, mostly in the context of BLS-signatures [18], e.g., by Ethereum and DFINITY. Industry-grade implementations of BLS12-381 are available, e.g., BLST (https://github.com/supranational/blst), arkworks (https://arkworks.rs/), Cloudflare CIRCL (https://github.com/cloudflare/circl).

While BLS12-381 is currently the de facto standard for (roughly) 128-bit security, options for higher security levels exist too, see e.g., [84]. A recent overview of 192-bit pairing-based curves was given by Aranha et al. [3], recommending BLS24-509 as the most efficient curve for that security level.

Pairing-Free Solutions. If supporting a pairing-friendly curve is considered to be infeasible, one can switch to a pairing-free variant of ACs. This essentially trades each computation that requires a pairing with an interaction with the issuer (and computation relying on the issuer's secret key). Thus, while it will provide significant privacy improvements over currently planned solutions, it will not solve the main deployment bottleneck introduced through batch-issuance.

The most prominent solution for pairing-free ACs are *one-time use anonymous credentials* such as U-Prove [75] or Anonymous Credentials Light [8], which are based on blind signatures extended to include attributes. Blind signatures are inherently *single-show* credentials, and thus require batch issuance of credentials for multi-show unlinkability.

Another approach being explored is to rely on *keyed-verification anonymous credentials* (KVAC), which are used for deployments where issuer and verifier are the same entity. KVACs can be seen as symmetric (MAC) variants of ACs, first introduced as CMZ MAC [33,71] (PS signatures are the publicly-verifiable version thereof), and later also being adapted for BBS MAC [9,71]. Such KVACs have already been deployed by Signal [34], and more recently, KVAC-based protocols have been proposed by Apple [94] and Google [86], in order to overcome deployment challenges in systems such as PrivacyPass.

While KVACs themselves are not suitable for user-centric and decentralized identity systems, they have recently been adopted for this application through the introduction of *Server-Aided Anonymous Credentials (SAAC)*. The idea was initially put forward by Desmoulins et al. [40] as BBS♯, and was later refined and formally analyzed by Chairattana-Apirom et al. (SAAC) [31]. The main idea is that all attribute-related statements in a presentation are verified directly by the

relying party, and only the one equation that requires the issuer's secret key for validation is checked with the help of the issuer. The underlying observation here is that this extra check does not reveal information about the user's attribute or identity to the issuer. There are two main ways how this can be implemented:

RP-Issuer Interaction: While the relying party verifies the bulk parts of the presentation, it asks the issuer for the verification of a sub-part of every proof. This requires high availability of the issuer, as every relying party relies on the issuer and its secret key based verification.

Batch Issuance: The user gets a batch of one-time use sub-tokens from the issuer in advance, and can spend them independently later on. This would require a similar deployment setting as the current batch issuance of salted hash based credentials with ECDSA, but provide the full feature set of anonymous credentials, and stronger privacy guarantees. In particular, Chairattana-Apirom et al. (SAAC) [31] show how to ensure that the issuer's sub-tokens can be used in an untraceable manner.

Most of the computations in the server-aided anonymous credential schemes based on BBS or PS signatures are equivalent to their pairing-based and multi-show unlinkable versions: the entire ZKP part is identical, only the pairing check gets replaced through a different computation. Thus, a server-aided and pairing-free version can also be used as a migration path towards full-fledged ACs.

Further Features & Extendability. Apart from efficiency, another core advantage of (DL-based) multi-message signature schemes is their simplicity and extendability. The proof system used on top of such signature schemes is simply a Schnorr proof. There exist general descriptions of how credential systems with an extensive feature set can be composed from multi-message signatures and Schnorr proofs [26,13], and adapting this to BBS or PS signatures is straightforward. There are also IETF specifications that extend the BBS core scheme [63], with pseudonyms [60], and blind issuance [59]. A promising approach is the IETF draft on Sigma protocols (which is the core of Schnorr proofs) [73] that would enable a more modular approach to extend DL-based multi-message signature with various Schnorr-based proof statements.

For some features, PS signatures [77] have advantages over BBS signatures, in particular when aiming at features involving the issuer secret key, e.g., credential aggregation [85,66] or distributed issuance [87,25]. If these features are important and cannot be rectified through other means, PS signatures might be the better choice. Note that the Schnorr-based proofs over the user's attributes built on top of the signature proof, as well as the different options for device binding work exactly the same for both BBS and PS, i.e., switching the signature scheme, or using both in parallel, is easily possible.

Quantum-Safeness. The most efficient multi-message schemes currently all rely on assumptions that are *not* quantum-safe, such as the DL or RSA assump-

tion. This only affects unforgeability, privacy can hold perfectly and unconditionally, which implies quantum-safeness. The latter depends on the concrete protocol and feature set, though, e.g., extensions for pseudonyms or conditional disclosure are, so far, often based on the DL assumption too. This is not an inherent limitation, and realizing these extensions in a quantum-safe manner for privacy protection is possible – albeit at higher costs.

Quantum-safe variants for such multi-message signatures with efficient proofs exist as well, e.g., [19,4,43]. As they rely on new number-theoretic lattice assumptions, more time to study the hardness of these problems is desirable.

3.2 General-Purpose Zero-Knowledge Proofs

The second path to construct anonymous credentials is to use general-purpose ZKPs, sometimes also referred to as circuit-based ZKPs. They allow a prover to convince a verifier that it correctly executed a particular algorithm, while possibly hiding all in- and outputs. The construction of most general-purpose ZKPs follow the same blueprint: first, the algorithm or function that the statement shall be proven over is transformed into an *arithmetic circuit*. The proof system then takes the circuit (or rather its translation into a constraint system) as well as private and public input, and proves correct evaluation of that circuit.

The main benefit of general-purpose ZKPs is that they allow to build anonymous credential systems based on *any* signature scheme, and in particular on legacy ECDSA – not only for the device, but also issuer signature. For a long time, these proof systems have been purely of theoretical interest, as their efficiency was far from being practical. This has changed with the introduction of zkSNARKs [14], which stands for *Succinct Non-Interactive Arguments of Knowledge*, which have seen their first practical implementation with Pinocchio [76] in 2013. Technically, zkSNARKs are only a particular type of general-purpose ZKPs, but this term is often used synonymously for the entire class of such ZKPs, and we also use the latter understanding here. The idea of using zkSNARKs to build privacy-preserving authentication on top of legacy systems was first explored in Cinderella by Delignat-Lavaud et al. [39] for the verification of X.509 Certificate Chains, and zk-creds by Rosenberg et al. [83] built the first anonymous credential system upon legacy systems.

The most recent advances entail the protocols by Frigo and Shelat [49], Paquin et al. [74], and Woo et al. [93], all enabling ZKPs over ECDSA-signed data in a few seconds with varying proof sizes. In fact, the work by Frigo and Shelat [49] not only proves knowledge of a plain ECDSA signature, but over a standardized credential format (mdoc), keeping the entire issuance infrastructure unchanged. From a deployment perspective, this approach introduces a separate privacy layer on top of existing standards and implementations, allowing for a clear separation of concerns. Also, device binding is solved naturally, as no changes to the Secure Elements are needed.

The main drawbacks of this approach are its complexity, since the proof systems are several magnitudes more complex than Schnorr proofs over BBS, and its maturity, since it is a field that is still evolving actively. And while our

work can not give a comprehensive overview of the state of the art on zkSNARKs and their individual trade-offs (we refer to [62] for such an overview), we want to highlight the main challenges that the adoption of this generic approach in the eID context needs to address and overcome.

Maturity and Security. There has been a significant amount of research in this area over the past decade, with Groth16 [53], Plonk [51], Bulletproofs [22] or Ligero [1] belonging to the most prominent systems. The recent systematization of knowledge paper on zkSNARKs lists over 40 constructions, and several new proof systems with significant performance improvements are published every year [62]. They differ significantly in their design choices, cryptographic assumptions and efficiency – often aiming at particular aspects for efficiency, such as either prover or verifier time, or proof sizes. For instance, schemes from pairings and DL-based assumptions, such as Groth16 [53] and Bulleproofs [22] allow for highly compact proofs, but have modest prover time. Proof systems such as Ligero can be built entirely from hash functions, which are faster and can provide quantum-safeness, but come with much larger proof sizes.

While these proof systems attract a significant amount of implementations and adoption, especially from the blockchain space, this ongoing development hinders a solid security analysis and understanding. It also poses a challenge for standardization, as any standardization effort that is started now will most likely not represent the state of the art by the time it is finished. Further, applications in the blockchain context (e.g., zk-Rollups) often focus on the succinctness property of zkSNARKs and, albeit using the term "zk", do not require or achieve privacy [90]. Thus, the results must be analyzed carefully to evaluate their suitability in the EUDI context, which does require the zero-knowledge property.

Multitude of Instantiations: zkSNARKs is the general term, and many different instantiations of that concept exist. This can be compared to the general concept of multi-message signatures, with BBS signatures being a concrete instantiation that has evolved over time as the most promising candidate (together with PS signatures). Different zkSNARK proof systems result in varying properties in terms of prover speed, verifier speed, proof size, and security assumptions. To the best of our knowledge, there is no consensus in the community yet, which is the best or most promising zkSNARK proof system to date and such an evaluation likely heavily depend on the intended usage of the system. To avoid a lock-in situation, a modular specification that allows to re-use common parts of these protocols is desirable. The ZKProofs community works on a similar approach for more than 7 years already [95], indicating that this is an endeavor requiring significant effort.

Trusted Setup: A particular differentiator among proof systems is whether they require a trusted setup or not. A trusted setup is a ceremony in which one or more parties create a common reference string (CRS) using secret and randomly generated inputs. The random inputs, sometimes also called "toxic waste" [70],

must be destroyed afterwards to guarantee the soundness of the protocol. This trusted setup can exist in different forms: (i) a circuit-specific trusted setup, where every change to the circuit requires an update to the trusted setup, or (ii) a universal (and updatable) trusted setup where the CRS can be used by all circuits within certain size boundaries and the setup can be updated by different parties over the time. For the updatable setup, to break soundness an adversary would need to get the inputs of all parties, thus a single party not leaking its input guarantees the soundness of the system. Groth16 is an example that falls in the first category, whereas Plonk requires the relaxed setup of the second category. So-called *transparent* proof systems, or transparent setups do not have that limitation, and the public parameters are derived from public randomness or are not necessary at all. Examples in that category are Bulletproofs, zkSTARKs, and hash-based proof systems such as Ligero [1]. Clearly, systems with transparent setups are easier to deploy than systems that require a trusted setup. However, systems with trusted setups have been successfully deployed in real-world applications, e.g., ZCash [70]. Similar setups are also not uncommon in existing internet infrastructure, with an example being the DNSSEC root signing ceremony that happens every 3 months [36].

Security Assumptions: As previously mentioned, the different proof systems are based on different security assumptions. Generally speaking, pairing-based constructions usually rely on strong non-falsifiable assumptions [69] or idealized models like AGM [50], whereas most transparent constructions make use of the random oracle model (ROM) which in turn can cause some challenges in the analysis and instantiations of security parameters [15]. Further, security is often traded for efficiency, e.g., the scheme Frigo and Shelat [49] achieves practical efficiency with prover time of 1.2s only for a reduced security level of 86-bit; performance numbers for a higher security level are unknown.

Complexity & Efficiency. The construction and security proofs of general-purpose ZKPs are significantly more complex than those of special-purpose constructions like multi-message signature schemes. For instance, the aforementioned protocol by Frigo et al. runs requires roughly 17.000 lines of code (without unit tests or benchmarks) [49]. Apart from the proof system itself, the complexity and efficiency are determined by the size of the circuit to be proven. The circuit is translated into a constraint system over which the proof is then computed. Small circuits have a few thousand constraints, whereas larger ones reach millions of constraints, and the same function can be represented in different circuits targeting different optimizations. To illustrate, the overview by Liang et al. [62] shows that to prove knowledge of a hash preimage x, such that $\mathsf{SHA256}(x) = y$, the necessary circuit can be represented in 27k to 600k constraints and prover time varies between 0.28s and 556s depending on proof system and circuit. Zero-knowledge proofs over ECDSA-signed credentials contain several such sub-proofs.

This complexity has already led to several attacks on proof systems and their implementations [32,35,61], which allowed malicious users to forge proofs for ran-

dom statements. In fact, the security review of the Google ZKP also revealed such vulnerabilities, including a " high-severity finding that allows arbitrary forgery of mdoc attribute claims"[42]. Thus, tool support to verify the correctness of the implementation and circuits become an essential part for the security properties of the overarching system. There is on-going work on the formal verification of proof systems, circuits, and implementations [37,7], but more work is necessary to reach a point where formal verification of constructions and implementations for complex systems needed in the context of ACs becomes feasible.

Further attention also needs to be given to optimizing the prover time, and to trade-offs made by tailoring the proof system for efficiency:
Prover performance: Most zkSNARKs have been designed for blockchain use-cases, and often aim at small proof sizes and verifier speed. In the context of digital credential systems, prover time is more important, though. Recent works aim to overcome those prover performance challenges by optimizing the building blocks for ECDSA-based signatures [49] or by splitting the proof generation into a computationally expensive pre-computation step and efficient presentation step [74]. Leveraging the performance gains from these recent advances, implementations on smartphones become feasible. While those novel optimizations build upon established techniques such as Ligero [1] and Groth16 [53], the optimizations themselves are relatively new and their constructions have not yet undergone thorough review by the scientific community.
General-purpose vs. tailored proof systems: While zkSNARKs are general-purpose proof systems by promise, the practical variants of these proof systems typically target a particular application and then develop a tailored protocol for this target. These optimizations usually identify costly parts in circuits and provide optimized components for those subparts/sub-proofs, or even move them out of the circuit as recently proposed by [57,72]. Those optimizations reduce prover time, but introduce additional complexity into the system, which needs to be formally analyzed, and make the protocol suitable only for a specific class of problems or arithmetic structures.

An interesting trend in the context of blockchain applications has been the research on so-called zero-knowledge Virtual Machines (zkVM) that allow for any program compiled with a specific compiler backend (or architecture) to be executed in a specially constructed virtual machine [52,5,92]. However, in contrast to their name, not all of these systems feature zero-knowledge properties, and the focus seems to be often more on verifiable computation and succinctness. The benefit of these constructions is that it simplifies the creation and extension of applications that use zkSNARKs for developers, but adds additional complexity in the proof system and does not provide practical efficiency yet.

Further Features & Extendability. Apart from the compatibility with legacy systems, another theoretical advantage of general-purpose ZKPs is the ability to express *any* logic and policy on the users' attributes in the signed credential. Thus, general-purpose ZKPs can support more expressive statements over attributes than what is possible through plain Schnorr-proofs. The latter is es-

sentially limited to statements that prove the equality, or and/or statements over DL-represented attributes. However, as mentioned above, this general-purpose feature is typically not (fully) achieved by the practical instantiations yet, which achieve the necessary efficiency by hard-coding a particular policy, credential format, and signature scheme into the proof system. This leads to a situation where general-purpose ZKPs can be less extendable than systems based on simple Schnorr proofs: adding practical extensions that go beyond the capabilities of the original protocol often require expert cryptographic knowledge, and significant resources for secure and efficient development.

Quantum-Safeness. The most widely deployed zkSNARKs e.g., Groth16 [53], Plonk [51], and Bulletproofs [22] rely on DL-based assumptions, and thus are not quantum-safe. Similar to BBS, the zero-knowledge property is often unconditional, which implies the quantum-safeness for zero-knowledge. Further, Groth16 and Plonk also rely on pairing-friendly curves, as do BBS signatures. Hash-based proof systems such as FRI [11] and Ligero [1], which is used by Frigo and shelat [49], have the advantage of quantum-resistance for soundness, if sufficiently large hash functions are used. There have also been significant advances in the area of lattice-based proof systems, e.g., [64,12,65], which will be the natural candidate when switching to lattice-based signatures such as ML-DSA.

Finally, it is important to note that while a proof system might provide quantum-safe soundness, this does not extend to the anonymous credentials scheme when it is build on top of classic signatures, such as ECDSA [49]. In fact, the practical efficiency of these proof systems often stems from highly-optimized and scheme-specific circuits. These circuits and techniques cannot be immediately re-used when switching to quantum-safe signature schemes such as ML-DSA, and in fact are likely to require a significant amount of research to develop targeted optimizations (and possibly new proof systems too).

4 Conclusion

Overall, several practical instantiations for a ZKP-based credential system exist that can improve the privacy and security in the EUDI ecosystem. They branch into two main approaches: 1) using a dedicated signature scheme that allows for simple Schnorr-based ZKPs, or 2) using a more powerful general-purpose ZKP system that can be used with any signature scheme, such as ECDSA. Both approaches differ significantly in their maturity, complexity, and legacy aspects.

BBS-based credential systems provide the best efficiency, maturity, and ease of extendability (for a certain set of features), and therefore have been recommended by the Cryptographers' Feedback [10]. However, they require additional support for standardization, depending on the device binding possibly additional hardware support, and the approval of pairing-friendly curves. The latter can be omitted through the use or BBS♯/Server-aided credentials, but this inherently requires additional communication with the issuer, similar to batch-issued credentials. Also, BBS-based signatures will require changes to all issuers, holders,

and verifiers that already support ECDSA-based attribute credentials. Device-binding from existing Secure Elements is feasible without any changes though.

In terms of compatibility with legacy systems, specialized general-purpose ZKPs have a compelling advantage as they can be built on top of the existing infrastructure. It will require changes to all holders, and verifiers, but issuers and especially issuer infrastructure can remain (mostly) unchanged. The latter only holds if credential formats are kept the same, whereas specialized formats would allow for much better performance. These protocols are significantly more complex and less efficient than BBS-based credential systems. Practical protocols have only been developed very recently, and thus do not enjoy the same level of scrutiny and maturity when it comes to their security analysis. Further, the complexity of the proof system and the speed with which these proof systems are currently developing will pose a challenge for standardization and formal verification. Currently, efficiency is typically gained by tailoring the proof system to a particular application and instantiation, and extensions that go beyond that initial scope will require significant resources and expert knowledge. In theory, however, general-purpose ZKPs can express arbitrary statements over the user's attributes, whereas Schnorr-proofs are more limited to DL-based statements.

Further, while none of the proposed solutions are fully quantum-safe yet, their further development and integration will provide valuable insights into the concrete constraints and actual feature sets that are needed by the real-world application. These insights then allow a more focused research to re-design such schemes and protocols, e.g., from lattices or hash-based signatures. The research in these areas in the last two years already demonstrates that practical quantum-safe anonymous credential systems are within reach. In terms of upgradeability, we expect similar resources to be required when migrating from either option. Even when parts rely on a quantum-safe proof system, the main work of designing efficient circuits and *practical* anonymous credential systems through tailored optimizations needs to be redone. Finally, we like to stress that the quantum risk for signatures is substantially different than for encryption, for which immediate PQC-transition to prevent harvest-now decrypt-later attacks is adviced. In contrast, anonymous credential schemes with quantum-vulnerable soundness, but perfect/statistical privacy can be used until practical quantum-computers capable of breaking their soundness become a reality, which is in line with the NIST recommendation on the use of quantum-vulnerable algorithms for authentication systems [68].

Acknowledgments. We thank Paul Bastian, Andrey Sidorenko, and Alexandros Zacharakis for their helpful feedback on an earlier draft of this work. The authors are solely responsible for the content of this publication; the positions presented here do not reflect the views of SPRIND.

References

1. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligero: Lightweight sublinear arguments without a trusted setup. In: Thuraisingham, B.M., Evans, D.,

Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 2087–2104. ACM Press (Oct / Nov 2017). https://doi.org/10.1145/3133956.3134104
2. Amrein, P.: https://github.com/UbiqueInnovation/zkattest-rs (2025)
3. Aranha, D.F., Fotiadis, G., Guillevic, A.: A short-list of pairing-friendly curves resistant to the special TNFS algorithm at the 192-bit security level. CiC **1**(3), 3 (2024). https://doi.org/10.62056/angyl86bm
4. Argo, S., Güneysu, T., Jeudy, C., Land, G., Roux-Langlois, A., Sanders, O.: Practical post-quantum signatures for privacy. In: Luo, B., Liao, X., Xu, J., Kirda, E., Lie, D. (eds.) ACM CCS 2024. pp. 1523–1537. ACM Press (Oct 2024). https://doi.org/10.1145/3658644.3670297
5. Arun, A., Setty, S.T.V., Thaler, J.: Jolt: SNARKs for virtual machines via lookups. In: Joye, M., Leander, G. (eds.) EUROCRYPT 2024, Part VI. LNCS, vol. 14656, pp. 3–33. Springer, Cham (May 2024). https://doi.org/10.1007/978-3-031-58751-1_1
6. Au, M.H., Susilo, W., Mu, Y.: Constant-size dynamic k-TAA. In: De Prisco, R., Yung, M. (eds.) SCN 06. LNCS, vol. 4116, pp. 111–125. Springer, Berlin, Heidelberg (Sep 2006). https://doi.org/10.1007/11832072_8
7. Bailey, B., Miller, A.: Formalizing soundness proofs of SNARKs. Cryptology ePrint Archive, Report 2023/656 (2023), https://eprint.iacr.org/2023/656
8. Baldimtsi, F., Lysyanskaya, A.: Anonymous credentials light. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 1087–1098. ACM Press (Nov 2013). https://doi.org/10.1145/2508859.2516687
9. Barki, A., Brunet, S., Desmoulins, N., Traoré, J.: Improved algebraic MACs and practical keyed-verification anonymous credentials. In: Avanzi, R., Heys, H.M. (eds.) SAC 2016. LNCS, vol. 10532, pp. 360–380. Springer, Cham (Aug 2016). https://doi.org/10.1007/978-3-319-69453-5_20
10. Baum et al.: Cryptographers' feedback on the eu digital identity's ARF. https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/issues/200 (2024)
11. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Fast reed-solomon interactive oracle proofs of proximity. In: Chatzigiannakis, I., Kaklamanis, C., Marx, D., Sannella, D. (eds.) ICALP 2018. LIPIcs, vol. 107, pp. 14:1–14:17. Schloss Dagstuhl (Jul 2018). https://doi.org/10.4230/LIPIcs.ICALP.2018.14
12. Beullens, W., Seiler, G.: LaBRADOR: Compact proofs for R1CS from module-SIS. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part V. LNCS, vol. 14085, pp. 518–548. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38554-4_17
13. Bichsel, P., Camenisch, J., Dubovitskaya, M., Enderlein, R., Krenn, S., Krontiris, I., Lehmann, A., Neven, G., Nielsen, J.D., Paquin, C., et al.: D2. 2 architecture for attribute-based credential technologies-final version. ABC4TRUST project deliverable (2014)
14. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Goldwasser, S. (ed.) ITCS 2012. pp. 326–349. ACM (Jan 2012). https://doi.org/10.1145/2090236.2090263
15. Block, A.R., Tiwari, P.R.: On the concrete security of non-interactive FRI. In: Galdi, C., Phan, D.H. (eds.) SCN 24, Part I. LNCS, vol. 14973, pp. 275–296. Springer, Cham (Sep 2024). https://doi.org/10.1007/978-3-031-71070-4_13
16. Bobolz, J., Eidens, F., Krenn, S., Ramacher, S., Samelin, K.: Issuer-hiding attribute-based credentials. In: Conti, M., Stevens, M., Krenn, S. (eds.) CANS

21. LNCS, vol. 13099, pp. 158–178. Springer, Cham (Dec 2021). https://doi.org/10.1007/978-3-030-92548-2_9
17. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Berlin, Heidelberg (Aug 2004). https://doi.org/10.1007/978-3-540-28628-8_3
18. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. Journal of Cryptology **17**(4), 297–319 (Sep 2004). https://doi.org/10.1007/s00145-004-0314-9
19. Bootle, J., Lyubashevsky, V., Nguyen, N.K., Sorniotti, A.: A framework for practical anonymous credentials from lattices. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part II. LNCS, vol. 14082, pp. 384–417. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38545-2_13
20. Brickell, E., Li, J.: Enhanced privacy id from bilinear pairing for hardware authentication and attestation. In: IEEE International Conference on Social Computing (2010) (2010)
21. Brickell, E., Li, J.: A pairing-based daa scheme further reducing tpm resources. In: International Conference on Trust and Trustworthy Computing (2010). pp. 181–195. Springer (2010)
22. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy. pp. 315–334. IEEE Computer Society Press (May 2018). https://doi.org/10.1109/SP.2018.00020
23. Camenisch, J., Drijvers, M., Lehmann, A.: Anonymous attestation using the strong Diffie Hellman assumption revisited. Cryptology ePrint Archive, Report 2016/663 (2016), https://eprint.iacr.org/2016/663
24. Camenisch, J., Drijvers, M., Lehmann, A.: Universally composable direct anonymous attestation. In: Cheng, C.M., Chung, K.M., Persiano, G., Yang, B.Y. (eds.) PKC 2016, Part II. LNCS, vol. 9615, pp. 234–264. Springer, Berlin, Heidelberg (Mar 2016). https://doi.org/10.1007/978-3-662-49387-8_10
25. Camenisch, J., Drijvers, M., Lehmann, A., Neven, G., Towa, P.: Short threshold dynamic group signatures. In: Galdi, C., Kolesnikov, V. (eds.) SCN 20. LNCS, vol. 12238, pp. 401–423. Springer, Cham (Sep 2020). https://doi.org/10.1007/978-3-030-57990-6_20
26. Camenisch, J., Krenn, S., Lehmann, A., Mikkelsen, G.L., Neven, G., Pedersen, M.Ø.: Formal treatment of privacy-enhancing credential systems. In: Dunkelman, O., Keliher, L. (eds.) SAC 2015. LNCS, vol. 9566, pp. 3–24. Springer, Cham (Aug 2016). https://doi.org/10.1007/978-3-319-31301-6_1
27. Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Berlin, Heidelberg (May 2001). https://doi.org/10.1007/3-540-44987-6_7
28. Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 56–72. Springer, Berlin, Heidelberg (Aug 2004). https://doi.org/10.1007/978-3-540-28628-8_4
29. Camenisch, J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 126–144. Springer, Berlin, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_8

30. Celi, S., Levin, S., Rowell, J.: CDLS: Proving knowledge of committed discrete logarithms with soundness. In: Vaudenay, S., Petit, C. (eds.) AFRICACRYPT 24. LNCS, vol. 14861, pp. 69–93. Springer, Cham (Jul 2024). https://doi.org/10.1007/978-3-031-64381-1_4
31. Chairattana-Apirom, R., Harding, F., Lysyanskaya, A., Tessaro, S.: Server-aided anonymous credentials. Cryptology ePrint Archive, Report 2025/513 (2025), https://eprint.iacr.org/2025/513
32. Chaliasos, S., Ernstberger, J., Theodore, D., Wong, D., Jahanara, M., Livshits, B.: SoK: What don't we know? Understanding security vulnerabilities in SNARKs. In: Balzarotti, D., Xu, W. (eds.) USENIX Security 2024. USENIX Association (Aug 2024), https://www.usenix.org/conference/usenixsecurity24/presentation/chaliasos
33. Chase, M., Meiklejohn, S., Zaverucha, G.: Algebraic MACs and keyed-verification anonymous credentials. In: Ahn, G.J., Yung, M., Li, N. (eds.) ACM CCS 2014. pp. 1205–1216. ACM Press (Nov 2014). https://doi.org/10.1145/2660267.2660328
34. Chase, M., Perrin, T., Zaverucha, G.: The Signal private group system and anonymous credentials supporting efficient verifiable encryption. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 1445–1459. ACM Press (Nov 2020). https://doi.org/10.1145/3372297.3417887
35. Ciobotaru, O., Peter, M., Velichkov, V.: The last challenge attack: Exploiting a vulnerable implementation of the fiat-shamir transform in a KZG-based SNARK. Cryptology ePrint Archive, Report 2024/398 (2024), https://eprint.iacr.org/2024/398
36. cloudflare.com: The dnssec root signing ceremony. https://www.cloudflare.com/en-gb/learning/dns/dnssec/root-signing-ceremony/ (2025)
37. Coglio, A., McCarthy, E., Smith, E., Chin, C., Gaddamadugu, P., Dellepere, M.: Compositional formal verification of zero-knowledge circuits. Cryptology ePrint Archive, Report 2023/1278 (2023), https://eprint.iacr.org/2023/1278
38. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y. (ed.) CRYPTO'94. LNCS, vol. 839, pp. 174–187. Springer, Berlin, Heidelberg (Aug 1994). https://doi.org/10.1007/3-540-48658-5_19
39. Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Parno, B.: Cinderella: Turning shabby X.509 certificates into elegant anonymous credentials with the magic of verifiable computation. In: 2016 IEEE Symposium on Security and Privacy. pp. 235–254. IEEE Computer Society Press (May 2016). https://doi.org/10.1109/SP.2016.22
40. Desmoulins, N., Dumanois, A., Kane, S., Traoré, J.: Making BBS anonymous credentials eIDAS 2.0 compliant (2025)
41. DG Connect: The european digital identity wallet architecture and reference framework. Retrieved January, https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework (2025)
42. Doyle, J., Ilunga, M.: Google longfellow security assessment. https://github.com/google/longfellow-zk/blob/main/docs/static/reviews/Longfellow_report_2025_08_18.pdf
43. Dubois, A., Klooß, M., Lai, R.W.F., Woo, I.K.Y.: Lattice-based proof-friendly signatures from vanishing short integer solutions. In: Jager, T., Pan, J. (eds.) PKC 2025, Part I. LNCS, vol. 15674, pp. 452–486. Springer, Cham (May 2025). https://doi.org/10.1007/978-3-031-91820-9_15
44. European Commission: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en (2024)

45. European Commission: Regulation on electronic identification and trust services for electronic transactions. https://eur-lex.europa.eu/eli/reg/2024/1183/oj/eng (2024)
46. Faz-Hernández, A., Ladd, W., Maram, D.: ZKAttest: Ring and group signatures for existing ECDSA keys. In: AlTawy, R., Hülsing, A. (eds.) SAC 2021. LNCS, vol. 13203, pp. 68–83. Springer, Cham (Sep / Oct 2022). https://doi.org/10.1007/978-3-030-99277-4_4
47. Faz-Hernandez, A., Scott, S., Sullivan, N., Wahby, R.S., Wood, C.A.: Hashing to Elliptic Curves. Request for Comments RFC 9380, Internet Engineering Task Force (Aug 2023). https://doi.org/10.17487/RFC9380
48. Friedrichs, K., Harding, F., Lehmann, A., Lysyanskaya, A.: Device-bound anonymous credentials with(out) trusted hardware. Cryptology ePrint Archive, Report 2025/1995 (2025), https://eprint.iacr.org/2025/1995
49. Frigo, M., shelat, a.: Anonymous credentials from ECDSA. Cryptology ePrint Archive, Report 2024/2010 (2024), https://eprint.iacr.org/2024/2010
50. Fuchsbauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 33–62. Springer, Cham (Aug 2018). https://doi.org/10.1007/978-3-319-96881-0_2
51. Gabizon, A., Williamson, Z.J., Ciobotaru, O.: PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953 (2019), https://eprint.iacr.org/2019/953
52. Goldberg, L., Papini, S., Riabzev, M.: Cairo – a Turing-complete STARK-friendly CPU architecture. Cryptology ePrint Archive, Report 2021/1063 (2021), https://eprint.iacr.org/2021/1063
53. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 305–326. Springer, Berlin, Heidelberg (May 2016). https://doi.org/10.1007/978-3-662-49896-5_11
54. Hesse, J., Singh, N., Sorniotti, A.: How to bind anonymous credentials to humans. In: Calandrino, J.A., Troncoso, C. (eds.) USENIX Security 2023. pp. 3047–3064. USENIX Association (Aug 2023), https://www.usenix.org/conference/usenixsecurity23/presentation/hesse
55. ISO: Iso 20008-2:2013: Information technology - security techniques - anonymous digital signatures - part 2: Mechanisms using a group public key. Standard, International Organization for Standardization, Geneva, CH (Nov 2013)
56. ISO: Iso/iec 24843: Information security - attribute-based credentials. Working draft, International Organization for Standardization, Geneva, CH (Nov 2025)
57. Kadianakis, G., Maller, M., Novakovic, A.: Sigmabus: Binding sigmas in circuits for fast curve operations. Cryptology ePrint Archive, Report 2023/1406 (2023), https://eprint.iacr.org/2023/1406
58. Kakvi, S., Martin, K., Putman, C., Quaglia, E.: Sok: Anonymous credentials. In: SSR 2023. https://doi.org/10.1007/978-3-031-30731-7
59. Kalos, V., Bernstein, G.: Blind bbs signatures. Internet Draft draft-irtf-cfrg-bbs-blind-signatures-01, Internet Research Task Force (Mar 2025)
60. Kalos, V., Bernstein, G.M.: Bbs per verifier linkability. Internet Draft draft-irtf-cfrg-bbs-per-verifier-linkability-01, Internet Research Task Force (Mar 2025)
61. Khovratovich, D., Rothblum, R.D., Soukhanov, L.: How to prove false statements: Practical attacks on fiat-shamir. Cryptology ePrint Archive, Report 2025/118 (2025), https://eprint.iacr.org/2025/118

62. Liang, J., Hu, D., Wu, P., Yang, Y., Shen, Q., Wu, Z.: SoK: Understanding zk-SNARKs: The gap between research and practice. Cryptology ePrint Archive, Report 2025/172 (2025), https://eprint.iacr.org/2025/172
63. Looker, T., Kalos, V., Whitehead, A., Lodder, M.: The BBS Signature Scheme. Internet Draft draft-irtf-cfrg-bbs-signatures-09, Internet Research Task Force (Jul 2025)
64. Lyubashevsky, V., Nguyen, N.K., Plançon, M.: Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 71–101. Springer, Cham (Aug 2022). https://doi.org/10.1007/978-3-031-15979-4_3
65. Lyubashevsky, V., Seiler, G., Steuer, P.: The LaZer library: Lattice-based zero knowledge and succinct proofs for quantum-safe privacy. In: Luo, B., Liao, X., Xu, J., Kirda, E., Lie, D. (eds.) ACM CCS 2024. pp. 3125–3137. ACM Press (Oct 2024). https://doi.org/10.1145/3658644.3690330
66. Mir, O., Bauer, B., Griffy, S., Lysyanskaya, A., Slamanig, D.: Aggregate signatures with versatile randomization and issuer-hiding multi-authority anonymous credentials. In: Meng, W., Jensen, C.D., Cremers, C., Kirda, E. (eds.) ACM CCS 2023. pp. 30–44. ACM Press (Nov 2023). https://doi.org/10.1145/3576915.3623203
67. Moody, D., Peralta, R., Perlner, R., Regenscheid, A., Roginsky, A., Chen, L.: Report on Pairing-based Cryptography. Journal of Research of the National Institute of Standards and Technology **120**, 11–27 (Feb 2015). https://doi.org/10.6028/jres.120.002
68. Moody, D., Perlner, R., Regenscheid, A., Robinson, A., Cooper, D.: Nist ir 8547: Transition to post-quantum cryptography standards (2024)
69. Naor, M.: On cryptographic assumptions and challenges (invited talk). In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 96–109. Springer, Berlin, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_6
70. Nikolaenko, V., Ragsdale, S., Bonneau, J., Boneh, D.: Powers-of-tau to the people: Decentralizing setup ceremonies. In: Pöpper, C., Batina, L. (eds.) ACNS 2024, Part III. LNCS, vol. 14585, pp. 105–134. Springer, Cham (Mar 2024). https://doi.org/10.1007/978-3-031-54776-8_5
71. Orrù, M.: Revisiting keyed-verification anonymous credentials. Cryptology ePrint Archive, Report 2024/1552 (2024), https://eprint.iacr.org/2024/1552
72. Orrù, M., Kadianakis, G., Maller, M., Zaverucha, G.: Beyond the circuit: How to minimize foreign arithmetic in ZKP circuits. Cryptology ePrint Archive, Report 2024/265 (2024), https://eprint.iacr.org/2024/265
73. Orrù, M., Yun, C.: Sigma protocols. Internet Draft draft-irtf-cfrg-sigma-protocols-01, Internet Research Task Force (Feb 2025)
74. Paquin, C., Policharla, G.V., Zaverucha, G.: Crescent: Stronger privacy for existing credentials. Cryptology ePrint Archive, Report 2024/2013 (2024), https://eprint.iacr.org/2024/2013
75. Paquin, C., Zaverucha, G.: U-prove cryptographic specification v1. 1. Technical Report, Microsoft Corporation (2011)
76. Parno, B., Howell, J., Gentry, C., Raykova, M.: Pinocchio: Nearly practical verifiable computation. In: 2013 IEEE Symposium on Security and Privacy. pp. 238–252. IEEE Computer Society Press (May 2013). https://doi.org/10.1109/SP.2013.47
77. Pointcheval, D., Sanders, O.: Short randomizable signatures. In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 111–126. Springer, Cham (Feb / Mar 2016). https://doi.org/10.1007/978-3-319-29485-8_7
78. Prime Minister's Office: https://www.gov.uk/government/news/new-digital-id-scheme-to-be-rolled-out-across-uk (2025)

79. Project, D.E.W.: Nl wallet: Software system overview. https://minbzk.github.io/nl-wallet/main/architecture/c4/software-system.html (2025)
80. Project, G.E.W.: German national eudi wallet: Architecture documentation. https://bmi.usercontent.opencode.de/eudi-wallet/wallet-development-documentation-public/architecture-concept/01-scope/ (2025)
81. Project, S.E.W.: Remote pake-protected services protocol (r2ps). https://expg.eidasweb.se/s/03jK698rE (2025)
82. Rosa, P.D.: Response to cryptographer's feedback. https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/discussions/211#discussioncomment-9882388
83. Rosenberg, M., White, J.D., Garman, C., Miers, I.: zk-creds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure. In: 2023 IEEE Symposium on Security and Privacy. pp. 790–808. IEEE Computer Society Press (May 2023). https://doi.org/10.1109/SP46215.2023.10179430
84. Sakemi, Y., Kobayashi, T., Saito, T., Wahby, R.S.: Pairing-Friendly Curves. Internet Draft draft-irtf-cfrg-pairing-friendly-curves-12, Internet Engineering Task Force (Nov 2025)
85. Sanders, O., Traoré, J.: Compact issuer-hiding authentication, application to anonymous credential. PoPETs **2024**(3), 645–658 (Jul 2024). https://doi.org/10.56553/popets-2024-0097
86. Schlesinger, S., Katz, J.: Anonymous credit tokens. Internet Draft draft-schlesinger-cfrg-act, Internet Engineering Task Force (Aug 2025)
87. Sonnino, A., Al-Bassam, M., Bano, S., Meiklejohn, S., Danezis, G.: Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers. In: NDSS 2019. The Internet Society (Feb 2019). https://doi.org/10.14722/ndss.2019.23272
88. Swiss federal authorities: e-id law approved at the ballot box. https://www.eid.admin.ch/en/e-id-gesetz-an-der-urne-angenommen-e (2025)
89. Tessaro, S., Zhu, C.: Revisiting BBS signatures. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 691–721. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30589-4_24
90. Tortola, D., Lisi, A., Mori, P., Ricci, L.: Tethering layer 2 solutions to the blockchain: A survey on proving schemes. Computer Communications **225**, 289–310 (2024)
91. Transportation Security Administration: https://www.tsa.gov/real-id/real-id-mobile-drivers-license-mdls (2025)
92. Wang, R., Hazay, C., Venkitasubramaniam, M.: Ligetron: Lightweight scalable end-to-end zero-knowledge proofs post-quantum ZK-SNARKs on a browser. In: 2024 IEEE Symposium on Security and Privacy. pp. 1760–1776. IEEE Computer Society Press (May 2024). https://doi.org/10.1109/SP54263.2024.00086
93. Woo, A.P.Y., Ozdemir, A., Sharp, C., Pornin, T., Grubbs, P.: Efficient proofs of possession for legacy signatures. Cryptology ePrint Archive, Report 2025/538 (2025), https://eprint.iacr.org/2025/538
94. Wood, C.A., Yun, C.: Anonymous rate-limited credentials. Internet Draft draft-yun-privacypass-crypto-arc-00, Internet Engineering Task Force (Oct 2025)
95. ZKProof.org: Zkproof charter. https://docs.zkproof.org/general (2018)