International Journal of Open Publication and Exploration (IJOPE), ISSN: 3006-2853
Volume 13, Issue 2, July-December, 2025, Available online at: https://ijope.com

Zero-Knowledge Proofs in Decentralized Identity Systems: Privacy, Interoperability, and Cross-Border Regulatory Compliance

Dr. Antonis Mavropoulos

Assistant Professor, Department of Informatics and Information Security, Aristotle University of Thessaloniki, Greece

ABSTRACT

Decentralized Identity (DID) and Self-Sovereign Identity (SSI) frameworks aim to restore user agency over digital attestations by disintermediating centralized identity providers. However, naive implementations of decentralized identity inherit substantial privacy risks, such as persistent identifier correlation, metadata leakage, and cryptographic signature tracking across verifiers. Zero-Knowledge Proofs (ZKPs) have emerged as the foundational cryptographic primitive capable of reconciling this tension by enabling verifiable attestations without revealing underlying claims or cryptographic artifacts. This comprehensive review examines the integration of modern ZKP constructions—including zk-SNARKs, zk-STARKs, and multi-message signature schemes (BBS/BBS+)—within decentralized identity ecosystems. We systematically evaluate the trilemma among mathematical privacy (unlinkability and data minimization), cross-platform technical interoperability (spanning W3C Verifiable Credentials 2.0, OpenID for Verifiable Credential Issuance/Presentations, and distributed ledgers), and international regulatory compliance (specifically the European Union's General Data Protection Regulation [GDPR], the updated eIDAS 2.0 European Digital Identity Architecture, and Anti-Money Laundering/Financial Action Task Force Travel Rules). Through critical analysis of empirical and theoretical literature, this study delineates current architectural bottlenecks, formalizes an integrated theoretical evaluation framework, synthesizes research gaps, and outlines future directions for scalable, compliant, and cryptographically sound decentralized identity systems.

Keywords: Decentralized Identity, Self-Sovereign Identity (SSI), Zero-Knowledge Proofs (ZKPs), Verifiable Credentials, BBS+ Signatures, zk-SNARKs, GDPR Compliance, eIDAS 2.0, Cross-Border Regulatory Compliance.

INTRODUCTION

Digital identity management has historically relied on centralized or federated directory architectures (e.g., SAML, OAuth 2.0, OpenID Connect). While operationally efficient for service providers, these topologies aggregate personal identifiable information (PII) into massive corporate or state honeypots, introducing systemic risks of catastrophic data breaches, mass surveillance, and single points of architectural failure. In response, the Decentralized Identity (DID) and Self-Sovereign Identity (SSI) paradigms have advanced an alternative paradigm where individuals hold and control their identity attestations locally, typically via digital credential wallets. Standardized by organizations like the World Wide Web Consortium (W3C) and the Decentralized Identity Foundation (DIF), these systems deploy Verifiable Credentials (VCs) and Decentralized Identifiers (DIDs) to construct cryptographic trust without centralized intermediaries.

Despite their decentralizing promise, early implementations of SSI revealed critical privacy regressions. Traditional digital signatures (such as ECDSA or Ed25519) calculated over identity payloads transform the issuer's signature into a global unique identifier. When a subject presents the same signed credential to multiple verifiers, colluding relying parties can correlate interactions and reconstruct user profiles across disparate domains. Furthermore, standard presentations often disclose excessive personal data (the "over-disclosure" problem)—for instance, revealing an exact residential address, birth date, and legal name merely to prove that an individual is over the legal age of majority.
Centralized Identity Model (Vulnerable to Honeypot Exploitation):

Zero-Knowledge Proofs (ZKPs) resolve this fundamental limitation. Originating in the seminal work of Goldwasser, Micali, and Rackoff (1985), ZKPs allow a prover to demonstrate to a verifier that a mathematical statement is true without revealing any information beyond the statement's validity. In decentralized identity, this allows a holder to prove predicates (e.g., "Age ≥ 18", "Citizenship $\in \text{EEA}$", "Income within tax bracket") or perform selective disclosure without releasing raw underlying attributes, plaintext claims, or linkable cryptographic signatures.

However, incorporating advanced zero-knowledge primitives into production decentralized ecosystems introduces non-trivial complexities. The field faces a structural trilemma:

1. **Cryptographic Privacy:** Achieving strict unlinkability, forward secrecy, and minimal metadata exposure.
2. **Global Interoperability:** Ensuring heterogeneous wallet implementations, verification libraries, proving systems, and distributed registries can seamlessly interoperate across administrative borders.
3. **Cross-Border Regulatory Compliance:** Reconciling zero-knowledge unlinkability and the immutability of distributed ledgers with strict legal regimes, notably the European Union's General Data Protection Regulation (GDPR) Right to Erasure (Article 17), eIDAS 2.0 European Digital Identity (EUDI) Wallet mandates, and Anti-Money Laundering (AML) / Know Your Customer (KYC) requirements.

This review provides a comprehensive synthesis of contemporary research on ZKPs within decentralized identity systems. We trace foundational cryptographic mechanics, critically examine published empirical and architectural studies, formalize an evaluation framework, analyze cross-border regulatory intersections, and map persistent research gaps and open trajectories.

LITERATURE REVIEW

To assess the state of the art, this section critically analyzes ten seminal and contemporary published works spanning cryptographic primitives, decentralized identity frameworks, client-side efficiency, and regulatory compliance.

Review of Published Studies

1. **Camenisch and Lysyanskaya (2001, 2004):** Founded attribute-based anonymous credentials using dynamic accumulators and bilinear maps. Their work demonstrated that identity presentations could achieve mathematical unlinkability. However, initial implementations incurred excessive computational complexity, limiting early smartphones from computing proofs in real-time.
2. **Boneh, Boyen, and Shacham (2004); Looker et al. (2023):** Standardized short group signatures and multi-message digital signature suites. Looker et al. formalized the W3C and IETF BBS Cryptosuite specifications, establishing that pairing-friendly elliptic curves (notably BLS12-381) allow blinding and selective disclosure on mobile devices with sub-100ms execution times. However, their security models presuppose static attribute lists, creating engineering overhead when applied to deeply nested JSON-LD hierarchies.
3. **Groth (2016):** Introduced the Groth16 zk-SNARK protocol, minimizing proof sizes to three group elements ($2 \in G_1, 1 \in G_2$) with constant-time verification. While Groth16 remains the gold standard for on-chain identity verifiers due to negligible gas consumption, the scheme requires a per-circuit trusted setup ceremony. If the toxic waste parameters from the ceremony are compromised, malicious actors can forge valid identity proofs without detection.
4. **Ben-Sasson et al. (2018):** Introduced transparent zk-STARKs, completely avoiding trusted setups and demonstrating asymptotic prover scalability. In decentralized identity evaluations, STARKs resolve long-term trust assumptions and guard against prospective quantum decryption. However, empirical evaluations demonstrate that STARK proof sizes exceed Bluetooth Low Energy (BLE) and Near Field Communication (NFC) single-frame transmission limits, posing usability challenges in offline retail identity scenarios.
5. **Windley et al. (2019):** Analyzed Hyperledger Indy and the AnonCreds protocol, providing foundational operational evidence of production-scale decentralized anonymous credentials. While proving the real-world viability of predicate evaluations (e.g., verifying employment status without revealing employee IDs), the architecture revealed significant vendor lock-in, as AnonCreds was deeply coupled to Hyperledger Indy's native ledger methods, preventing seamless cross-chain interoperability.
6. **Dunphy and Fabrega (2018):** Published a comprehensive architectural survey on distributed identity systems, dissecting the trade-offs between ledger immutability and user privacy. Their critical insight was that storing DIDs or document state roots on public ledgers creates systemic traffic analysis vectors, where correlating transaction frequencies can de-anonymize subjects even when identity payloads are fully hashed.
7. **Bünz et al. (2018):** Formulated Bulletproofs—short non-interactive zero-knowledge proofs requiring no trusted setup. While primarily adopted for confidential crypto-asset transactions, several identity architectures adapted Bulletproofs for logarithmic range proofs (e.g., verifying credit thresholds or age brackets). However, verification time scales linearly with statement complexity, making batch verification across multiple decentralized credentials prohibitively slow compared to SNARKs.
8. **Glöckler et al. (2023, 2025):** Investigated the intersection of eIDAS 2.0 Architectural Reference Framework (ARF) and cryptographic unlinkability. They exposed that early European Digital Identity (EUDI) specifications relied on conventional X.509 signatures over SD-JWT (Selective Disclosure JSON Web Tokens), failing to

prevent issuer or verifier collusion. Their empirical findings confirmed that without zero-knowledge cryptographic blinding (such as BBS+ or AnonCreds), digital identity wallets cannot strictly fulfill the data minimization mandates of GDPR Article 5(1)(c).

9. **Zhou et al. (2024):** Surveyed ZKPs for blockchain-based identity sharing. They categorized computational bottlenecks in decentralized identity, concluding that mobile wallet prover time and client-side memory exhaustion represent the primary obstacles to adopting universal recursive zk-SNARKs (like Nova or Halo2) across consumer smartphones.
10. **Končius et al. (2025/2026):** Provided benchmark evaluations comparing client-side execution libraries (ZoKrates, SnarkJS, Arkworks) across mobile and web runtime environments. Their work demonstrated that WebAssembly (WASM) execution environments in digital wallets suffer a 3.5× to 5× latency degradation compared to native Rust implementations, reinforcing the need for hardware-accelerated cryptographic primitives in decentralized identity architectures.

Architectural Paradigms and Cross-Platform Interoperability

Decentralized identity architecture relies on the foundational "Trust Triangle": the Issuer, the Holder (Subject), and the Verifier (Relying Party). Integrating ZKPs fundamentally reconfigures the trust interactions across this triangle.

Reconfiguring the Trust Triangle

1. **Issuance:** The Issuer binds claims to a holder's public key or secret commitment and signs the payload (e.g., using BBS+ or a SNARK-friendly algebraic commitment like Poseidon). Crucially, the issuer does not participate in future verifications, eliminating identity tracking by the issuing entity.
2. **Proof Generation (Holder):** The holder transforms the stored credential into a Zero-Knowledge Verifiable Presentation (VP). Using client-side witness generation, the wallet blinds the signature, selectively hides attributes, and calculates mathematical proofs over hidden attributes (e.g., evaluating range predicates).
3. **Verification (Verifier):** The verifier inspects the cryptographic proof π using the issuer's public key published on a trust registry (e.g., a distributed ledger, decentralized file system, or DNS/DID Document). The verifier confirms validity without seeing hidden claims or obtaining a correlating identifier.

Standards Convergence and Technical Impediments

Interoperability across disparate systems requires standardization across data models, proof formats, and transport protocols.

- **Data Models (W3C VC 2.0 vs. SD-JWT):** The W3C Verifiable Credentials Data Model v2.0 formally incorporates Data Integrity Cryptosuites supporting BBS+ signatures (bbs-2023) and zero-knowledge proofs. Concurrently, the IETF has standardized Selective Disclosure JSON Web Tokens (SD-JWT / RFC 9528), which relies on salted hashes for selective attribute disclosure. However, SD-JWT uses conventional asymmetric signatures over digest arrays; revealing any single salt-value pair exposes the underlying signature to verifier tracking. Truly unlinkable SD-JWT requires layering BBS+ or SNARK wrappers over the JWT claims.
- **Transport and Exchange Protocols:** Interoperability requires standard presentation exchange mechanisms. OpenID for Verifiable Credential Issuance (OIDC4VCI) and OpenID for Verifiable Presentations (OIDC4VP) have emerged as the dominant transport layers. Aligning OIDC4VP with ZKP payloads requires defining standardized cryptographic query mechanisms (e.g., DIF Presentation Exchange and W3C Digital Credentials API) capable of parsing zero-knowledge circuit inputs without protocol serialization failures.
- **Revocation Architectures:** Revocation represents a major vulnerability for identity privacy. Traditional Online Certificate Status Protocol (OCSP) or simple Certificate Revocation Lists (CRLs) leak operational metadata, allowing issuers to trace every time and place a credential is used. In privacy-preserving decentralized architectures, revocation is managed via **Dynamic Cryptographic Accumulators** (e.g., Merkle trees, RSA accumulators, or Bilinear accumulators). The holder provides a zero-knowledge non-membership proof demonstrating that their credential identifier is not present within the issuer's revocation accumulator set, without revealing which credential in the set is theirs.

Comparative Analytical Synthesis

To assess operational performance, we evaluate proving systems across structural dimensions, followed by an analysis of decentralized identity frameworks.

Table 1: Cryptographic Proving Systems in Decentralized Identity

Scheme	Prover Time (Mobile)	Verifier Time	Proof Size	Trusted Setup	Post-Quantum Security	Primary Identity Application
BBS+ Cryptosuite	< 100 ms	< 15 ms	~150–350 bytes	No (Requires pairing curves)	No (Vulnerable to Shor's algorithm)	Selective disclosure, anonymous attestations
Groth16 (SNARK)	1.5 – 5.0 s	< 5 ms	128–192 bytes	Yes (Per-circuit ceremony required)	No (Discrete logarithm dependent)	Complex predicate proofs, on-chain verification
PLONK / Halo2 (SNARK)	3.0 – 8.0 s	< 10 ms	500–1000 bytes	Universal / Transparent (No per-circuit setup)	No (Elliptic curve pairings)	Generalized recursive identity circuits
zk-STARKs	5.0 – 15.0 s	20 – 80 ms	40–150 KB	None (Transparent)	Yes (Hash-function based)	Quantum-resistant state machine proofs
Bulletproofs	2.0 – 6.0 s	Linear in size ($O(n)$)	1–2 KB	None (Transparent)	No (Discrete logarithm dependent)	Confidential numerical range proofs (Age, Income)

Table 2: Decentralized Identity Framework Implementations

Framework / Standard	Underlying Cryptography	Unlinkability Guarantee	Interoperability Level	Regulatory Fit (GDPR / eIDAS 2.0)	Deployment Constraints
W3C Data Integrity BBS+	Asymmetric Pairing (BLS12-381)	Strong (Full cryptographic re-randomization)	High (Native JSON-LD / W3C VC 2.0 mapping)	High (Native Article 5(1)(c) data minimization)	Requires specialized pairing libraries in client runtimes
Hyperledger AnonCreds	CL-Signatures / Camenisch-Lysyanskaya	Strong (Blinded signature presentation)	Moderate (Historical coupling to Hyperledger Indy)	Substantial (Compliant data minimization architecture)	Legacy format migration required for W3C VC 2.0
SD-JWT (IETF RFC 9528)	Conventional Signatures (ECDSA / Ed25519) + Hashes	Weak (Issuer signature linkable across verifiers)	Very High (Standard JSON Web Signature tooling)	Moderate (Selective disclosure without unlinkability)	Colluding verifiers can correlate presentation signatures

Framework / Standard	Underlying Cryptography	Unlinkability Guarantee	Interoperability Level	Regulatory Fit (GDPR / eIDAS 2.0)	Deployment Constraints
Polygon ID / Iden3	zk-SNARKs (Groth16) over BabyJubjub / Circom	Strong (Witness evaluation inside zk-circuit)	Moderate (EVM and off-chain relay focus)	High (Zero PII processed or stored on-chain)	Significant circuit compilation and mobile prover overhead

Cross-Border Regulatory Compliance and Legal Dimensions
The deployment of ZKP-enabled decentralized identity intersects directly with international data protection and financial supervisory regulations. Far from operating in a legal vacuum, cryptographic primitives must align with statutory mandates across multiple jurisdictions.

General Data Protection Regulation (GDPR) Compliance
The European Union's General Data Protection Regulation (Regulation (EU) 2016/679) establishes strict compliance obligations that historically conflicted with public distributed ledgers:

- **Data Minimization (Article 5(1)(c)):** Mandates that personal data must be adequate, relevant, and limited to what is strictly necessary in relation to processing purposes. Standard identity verification regularly breaches this principle by over-disclosing unneeded attributes. BBS+ and zk-SNARKs fulfill data minimization *by design and by default* (Article 25). When a relying party verifies that an individual is qualified to drive or over a required age, no collateral PII (name, residential address, precise birth date) is transferred or stored.
- **The Right to Erasure / Right to Be Forgotten (Article 17):** Imposes the duty to delete personal data upon request. Writing identity hashes or DIDs to an immutable, decentralized ledger creates severe legal non-compliance, as cryptographic hashes of personal data remain classified as pseudonymous personal data under EU case law (e.g., *Breyer v. Bundesrepublik Deutschland*). ZKP-based decentralized identity resolves this conflict: distributed ledgers are utilized exclusively to anchor public verification keys, cryptographic accumulator state roots, and circuit schema definitions. No user identifiers, encrypted PII, or personal transaction hashes ever touch the immutable ledger.
- **Cross-Border Transfers (Chapter V):** Under Articles 44–49, transferring personal data outside the European Economic Area (EEA) requires an adequacy decision or standard contractual clauses (SCCs). Because a zero-knowledge presentation constitutes a mathematical proof of claim validity without transmitting personal data, the transfer of a pure ZKP across international borders does not trigger the statutory restrictions of international PII transmission, dramatically reducing cross-border enterprise compliance friction.

eIDAS 2.0 and the European Digital Identity (EUDI) Wallet
Regulation (EU) 2024/1183 (eIDAS 2.0) mandates that European Member States provide an EU Digital Identity Wallet to citizens, obligating qualified relying parties (such as banks, utilities, and telecommunications providers) to accept qualified electronic attestations of attributes (QEAAs).

The primary operational challenge within the eIDAS 2.0 Architecture Reference Framework (ARF) has been the prevention of user tracking. Under legacy federated models, the state identity provider tracks citizen log-in sessions across commercial platforms. If EUDI Wallets adopt standard digital signatures, colluding relying parties could aggregate user presentation profiles across systems.

Recent revisions to the ARF explicitly recognize zero-knowledge primitives and anonymous credentials (notably BBS+ cryptosuites and AnonCreds) as mechanisms that satisfy both Level of Assurance (LoA) "High" and the stringent unobservability requirements demanded by European privacy watchdogs. The credential holder proves authenticity from an authorized governmental issuer without allowing the verifier to identify the underlying credential certificate or correlate presentations.

Financial Surveillance, AML, and FATF Travel Rule Reconciliations

Financial integrity frameworks enforced by the Financial Action Task Force (FATF) and the European Union's Anti-Money Laundering Directives (such as AMLD6 and the Transfer of Funds Regulation) present a fundamental counterweight to absolute cryptographic privacy. Financial institutions are legally barred from processing transactions for anonymous customers; they must conduct strict Customer Due Diligence (CDD) and enforce the "Travel Rule," transmitting originator and beneficiary identity data alongside funds.

This creates an apparent paradox: ZKP systems prioritize unobservability, while AML frameworks mandate full identification. To bridge this gap, modern decentralized architectures implement **Verifiable Threshold Escrow and**

In this architecture, during proof generation, the holder uses verifiable threshold public-key encryption (e.g., ElGamal or threshold Paillier over verifiable secret sharing) to encrypt identity attributes under the public keys of a distributed set of independent statutory trustees (e.g., judicial bodies, compliance ombudsmen, and banking supervisors). The ZKP circuit evaluates two simultaneous statements:

1. The user satisfies compliance predicates (e.g., "Customer has completed KYC, is not present on OFAC/sanctions lists, and resides in a compliant jurisdiction").
2. The ciphertext correctly encrypts the true identity witness under the regulatory escrow keys.

During standard operations, zero personal data is revealed, and verifiers observe only mathematical validity. In the event of confirmed financial crime, money laundering, or a court order, a threshold k-of-n quorum of trustees must co-sign to decrypt the identity attributes. This guarantees conditional privacy, reconciling self-sovereign cryptographic guarantees with international legal enforceability.

OPEN RESEARCH GAPS AND CHALLENGES

Despite foundational theoretical progress, several interdisciplinary challenges impede ubiquitous global deployment:

Client-Side Mobile Computational Constraints

While verification of zk-SNARKs and BBS+ proofs is computationally light (typically 2–15 ms), **proof generation** requires substantial processing power and memory. Generating a complex Groth16 or PLONK proof over large circuits involving nested hashing (such as SHA-256 or Keccak) requires hundreds of thousands of R1CS constraints, exhausting mobile memory and taking up to 10 seconds on consumer hardware.

While ZKP-friendly algebraic hash functions (e.g., Poseidon, Rescue, Reinforced Concrete) compress constraint counts by orders of magnitude, legacy enterprise databases and identity registries sign documents using standard SHA-256 or RSA. Consequently, identity circuits must evaluate non-native arithmetic, creating significant performance overhead on edge devices.

Post-Quantum Migration Paths

The overwhelming majority of production ZKP-SSI implementations rely on elliptic curve cryptography (ECC), specifically pairing-friendly curves such as BLS12-381, BN254, and secp256k1. Shor's algorithm on sufficiently scaled quantum architectures will reduce the discrete logarithm and elliptic curve Diffie-Hellman problems to polynomial time, invalidating the binding and zero-knowledge properties of Groth16, PLONK, and BBS+ signatures.

While zk-STARKs and lattice-based anonymous credentials offer post-quantum security, their proof sizes (40–100+ KB) and large public key footprints create substantial network transmission friction, particularly in offline identity verification environments. A unified post-quantum anonymous credential scheme that combines compact proof sizes with sub-second mobile prover latency remains an open research priority.

Semantic Interoperability and Universal Schema Registries

Even when systems share a proving back-end, differences in schema semantics prevent cross-border exchange. An attestation schema defined in the European Union (under eIDAS QEAA profiles) uses different semantic ontologies than an American driver's license issued under ISO/IEC 18013-5 (Mobile Driving License) or W3C VC specifications. Constructing universal zero-knowledge circuits capable of dynamically evaluating heterogeneous credential inputs without recompiling circuit parameters represents an unresolved engineering bottleneck.

The "Garbage In, Garbage Out" Paradox

Zero-knowledge proofs establish that an input satisfies specific mathematical assertions inside a circuit; they cannot guarantee the ontological truth of the underlying input in the physical world. If an issuer issues a valid credential to a fraudulent identity, ZKPs preserve the mathematical validity of that fraudulent credential. The decentralized identity ecosystem remains vulnerable to upstream root-of-trust compromises, synthetic identity creation, and coercion attacks (where a verifier forces a user to decrypt or disclose raw data via external administrative coercion).

FUTURE DIRECTIONS

To transition ZKP-enabled decentralized identity from experimental deployments to universal civil infrastructure, research must advance along four primary trajectories:

1. **Hardware-Accelerated Client Environments:** Transitioning mobile prover runtimes from single-threaded WebAssembly (WASM) to native, hardware-accelerated environments. Integrating specialized cryptographic coprocessors (NPUs and mobile GPUs) with mobile frameworks (e.g., Arkworks, WebGPU) can compress proving latency to under 300 ms for general-purpose circuits, providing instantaneous presentation interfaces.
2. **Recursive SNARKs and Incrementally Verifiable Computation (IVC):** Protocols utilizing recursive proof composition (such as Nova, SuperNova, and Halo2) allow multiple disparate identity attestations to be folded into a single succinct proof. A user could aggregate an academic degree, professional accreditation, tax compliance status, and criminal record clearance into one single-byte proof verification, without leaking the issuance sources or individual timestamps.
3. **Formal Verification of Identity Circuits:** As identity logic shifts into domain-specific languages (e.g., Circom, Noir, Leo), bugs in arithmetic circuits (such as under-constrained variables or field overflow vulnerabilities) can allow malicious provers to forge proofs of eligibility. Establishing automated formal verification and property-based fuzzing tools specifically tailored to decentralized identity circuits will be essential for civil-scale deployments.
4. **Standardized Identity Gateways for Cross-Border Recognition:** Multi-lateral identity federations should establish open cryptographic gateway protocols that bridge BBS+ W3C credentials, ISO 18013-5 mDLs, and eIDAS 2.0 QEAA standards. Developing universal zero-knowledge compiler pipelines will allow a single citizen wallet to resolve regulatory compliance dynamically, ensuring privacy protection regardless of cross-border verification jurisdictions.

CONCLUSION

Decentralized Identity represents a paradigm shift away from fragile, surveillance-prone centralized directories toward user-centric data sovereignty. However, decentralization without advanced cryptography merely distributes data leakage vectors. Zero-Knowledge Proofs serve as the crucial structural bridge that resolves this challenge, enabling rigorous data minimization, forward secrecy, and mathematical unlinkability.

As documented in this review, while multi-message signature schemes like BBS+ provide immediate, lightweight solutions for selective disclosure, general-purpose zk-SNARKs and zk-STARKs offer the architectural flexibility necessary to evaluate complex predicates, range proofs, and verifiable compliance algorithms. Harmonizing these cryptographic constructions with cross-border regulatory frameworks—including GDPR, eIDAS 2.0, and international AML regimes—establishes that digital verification need not require sacrificing civil privacy or legal auditability. Addressing current performance bottlenecks, developing post-quantum proving constructions, and establishing universal semantic schemas will be key to establishing zero-knowledge decentralized identity as the cryptographic foundation for global digital interaction.

REFERENCES

1. Ben-Sasson, E., Bentov, I., Horesh, Y., & Riabzev, M. (2018). Scalable, transparent, and post-quantum secure computational integrity. *Cryptology ePrint Archive*, Report 2018/046.
2. Boneh, D., & Boyen, X. (2004). Short signatures without random oracles. In *Advances in Cryptology – EUROCRYPT 2004* (pp. 56–73). Springer, Berlin, Heidelberg.
3. Boneh, D., Boyen, X., & Shacham, H. (2004). Short group signatures. In *Advances in Cryptology – CRYPTO 2004* (pp. 41–55). Springer, Berlin, Heidelberg.
4. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018). Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy (SP)* (pp. 315–334). IEEE.

5. Camenisch, J., & Lysyanskaya, A. (2001). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in Cryptology – EUROCRYPT 2001* (pp. 93–118). Springer, Berlin, Heidelberg.
6. Camenisch, J., & Lysyanskaya, A. (2004). Signature schemes and anonymous credentials from bilinear maps. In *Advances in Cryptology – CRYPTO 2004* (pp. 56–72). Springer, Berlin, Heidelberg.
7. Dunphy, P., & Fabrega, F. A. (2018). A first look at identity management platforms on the blockchain. *IEEE Security & Privacy*, 16(4), 20–29.
8. European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*, L119, 1–88.
9. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1183 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework (eIDAS 2.0). *Official Journal of the European Union*, L 2024/1183.
10. Financial Action Task Force. (2021). *Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers*. FATF, Paris.
11. Glöckler, M., Busch, M., & Korsch, M. (2025). The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets. *Internet Policy Review*, 14(3), 1–28.
12. Goldwasser, S., Micali, S., & Rackoff, C. (1985). The knowledge complexity of interactive proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing (STOC '85)* (pp. 291–304). ACM.
13. Groth, J. (2016). On the size of pairing-based non-interactive arguments. In *Advances in Cryptology – EUROCRYPT 2016* (pp. 305–326). Springer, Berlin, Heidelberg.
14. Gupta, S. (2025). Zero-Knowledge Proofs For Privacy-Preserving Systems: A Survey Across Blockchain, Identity, And Beyond. *Engineering and Technology Journal*, 10(7), 412–428.
15. Končius, A. L., Košubienė, G., Atmanavičiūtė, L., & Masteika, S. (2025). Comparative Analysis and Implementation of Zero-Knowledge Proof Libraries for Digital Identity. *Vilnius University Research Series*, 1–14.
16. Looker, T., Kalos, M., & Lodderstedt, T. (2023). *BBS Signature Suite v1.0 for W3C Data Integrity*. W3C Candidate Recommendation.
17. Mühle, A., Grüner, A., Gayvoronskaya, T., & Meinel, C. (2018). A survey on essential components of a self-sovereign identity. *Computer Science Review*, 30, 80–86.
18. Parno, B., Howell, J., Gentry, C., & Raykova, M. (2013). Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy* (pp. 238–252). IEEE.
19. Sporny, M., Longley, D., & Chadwick, D. (2025). *Verifiable Credentials Data Model v2.0*. W3C Recommendation.
20. Tobin, A., & Reed, D. (2016). The Inevitable Rise of Self-Sovereign Identity. *Sovrin Foundation White Paper*, 1–24.
21. Windley, P., Reed, D., & Hardman, D. (2019). Hyperledger Indy: Directed identity on distributed ledgers. In *Blockchain and Web3 Strategies* (pp. 112–134). Linux Foundation.
22. World Wide Web Consortium (W3C). (2022). *Decentralized Identifiers (DIDs) v1.0: Core Architecture, Data Model, and Representations*. W3C Recommendation.
23. Yang, Q., Shen, Z., & Wu, L. (2025). SoK: Understanding zk-SNARKs: The gap between research and practice. *arXiv preprint arXiv:2502.02387*.
24. Zhou, L., Diro, A., Saini, A., Kaisar, S., & Hiep, P. C. (2024). Leveraging zero knowledge proofs for blockchain-based identity sharing: A survey of advancements, challenges and opportunities. *Journal of Information Security and Applications*, 80, 103678.