

**PRIVACY AND TRUST IN THE AGE OF DEEP FAKES: A SOCIOTECHNICAL
EXPLORATION OF ZERO-KNOWLEDGE PROOFS ON THE ALEO BLOCKCHAIN**

A **Research Paper** submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Youssef Cherrat

Spring, 2025.

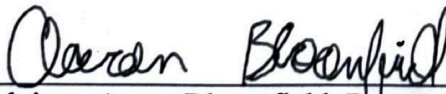
On my honor as a University Student, I have neither given nor received
unauthorized aid on this assignment as defined by the Honor Guidelines
for Thesis-Related Assignments

Signature
Youssef Cherrat



Date 5-6-25

Signature



Date 5-6-25

Technical Advisor: Aaron Bloomfield, Department of Computer Science

STS Advisor: Richard D. Jacques

1. Introduction

A. Background

In an increasingly digital world, verifying personal identity has become essential yet problematic. Industries such as finance and healthcare depend on accurate identity verification, but the process often forces individuals to share sensitive personal information, raising concerns about security and privacy breaches (Wilson Center n.d.). In 2022 alone, over 4,100 publicly disclosed data breaches exposed approximately 22 billion records of personal data (Cyber Security Hub n.d.). Traditional solutions, such as Know-Your-Customer (KYC) procedures, are designed to combat fraud but exacerbate privacy risks by requiring the distribution of sensitive data to multiple third parties, thus increasing the potential attack surface for hackers (Wilson Center n.d.).

This tension between security and privacy in identity management has made privacy-preserving identity verification a critical research area. The goal is to allow individuals to prove aspects of their identity—such as age, citizenship, or account ownership—without exposing underlying personal details. Zero-knowledge proofs (ZKPs) offer a promising approach by enabling verification of identity attributes without revealing personal data (Wilson Center n.d.). A ZKP allows one party (the prover) to convince another party (the verifier) that a claim is true without disclosing how they know it (Goldwasser, Micali, and Rackoff 1989). First introduced in the 1980s, ZKPs have evolved into practical protocols like zk-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge), which provide both privacy and efficiency. The cryptocurrency Zcash has successfully deployed zk-SNARKs to validate transactions while keeping sender, receiver, and transaction details hidden from the public ledger (Li et al. 2020). These advances demonstrate that verification can occur while keeping underlying data secret, paving the way for private identity verification.

B. Related Works

Recent research has explored ZKPs for identity verification in various applications. Li et al. (2020) developed a blockchain-based ridesharing system that verifies driver identities through ZKPs without exposing personal data. Their implementation, built on Hyperledger Fabric, an enterprise blockchain, ensures that proof verification is stored on an immutable ledger and is executed via cryptographic libraries such as Hyperledger Ursa (Li et al. 2020). While this study

validates the feasibility of ZKP-backed identity verification, it also highlights certain limitations. Many early implementations rely on permissioned blockchain networks, limiting their applicability in open, decentralized environments.

In the broader context of decentralized identity, distributed ledgers provide users with control over credentials, but ensuring robust privacy—so that even the ledger does not reveal sensitive data—remains a challenge (IBM Blockchain n.d.). While existing solutions have shown that "prove without reveal" identity checks are feasible, they often operate within isolated ecosystems rather than universal, privacy-preserving public blockchain infrastructures (Li et al. 2020). Additionally, implementing ZKP solutions traditionally requires deep cryptographic expertise, restricting widespread adoption.

To address this, the Leo programming language has emerged as a tool for simplifying ZKP development. Leo is a statically-typed, domain-specific language designed for formally verified, privacy-preserving smart contracts (Aleo n.d.). It abstracts the mathematical complexity of ZKPs, making it easier to encode identity verifications. However, because Leo and the Aleo blockchain are relatively new, limited research exists evaluating their real-world effectiveness for identity verification (Aleo n.d.). This knowledge gap—the intersection of advanced cryptographic theory and practical identity systems—defines the problem this paper seeks to address.

C. Research Problem

Research Problem: How can we design and implement a privacy-preserving identity verification system using zero-knowledge proofs on the Aleo blockchain, and to what extent does the Leo programming language facilitate this implementation?

This study aims to explore the technical challenges of designing a privacy-centric, blockchain-based identity verification system that protects user data. Addressing this problem is crucial for theoretical and practical reasons. Theoretically, it contributes to cryptographic proof research, moving toward a model where digital interactions replace the need for trust with the ability to cryptographically verify information (Wilson Center n.d.). Practically, such a solution could transform business and institutional identity verification processes. Instead of collecting passport copies or driver's licenses, companies could verify user credentials through

cryptographic proofs, eliminating the need to store sensitive personal data, which is a frequent target of cyberattacks (Cyber Security Hub n.d.).

Given rising privacy regulations and frequent data breaches, a zero-knowledge identity verification system offers an opportunity to enhance security while minimizing data exposure (IBM Blockchain n.d.). Aleo's privacy-focused Layer-1 blockchain, which natively supports zero-knowledge cryptography, is well-suited to this goal (Aleo n.d.). By executing computations off-chain and storing only succinct verification proofs on-chain, Aleo effectively enables decentralized private computation, making it an ideal platform for privacy-preserving identity verification (ZK Proof n.d.).

D. Research Objectives

To investigate the research problem, my study establishes the following objectives. The first objective is the design and implementation of a prototype identity verification protocol that leverages zero-knowledge proofs (ZKPs) to confirm identity attributes, such as proof of age or possession of a valid credential, without revealing personal data. This protocol will be implemented using the Leo programming language on the Aleo blockchain.

The second objective is the functionality and performance evaluation of the prototype to ensure that it operates correctly and efficiently. This evaluation will assess proof generation time, verification speed, transaction costs, and overall system scalability. Additionally, a security analysis will be conducted to verify that the system protects user privacy while remaining resistant to data breaches and replay attacks.

The third objective is the analysis of Leo and Aleo for identity verification, which involves documenting development challenges, limitations, and best practices for implementing ZKP-based identity verification using these technologies.

By achieving these objectives, my study will provide valuable insights into the feasibility and effectiveness of deploying privacy-preserving identity verification solutions on public blockchains.

2. Methodology Overview

A. Research Question and Approach

This study employs an experimental technical methodology to design, implement, and evaluate a zero-knowledge proof (ZKP) identity verification system. The central research question is: "How can we design and implement a privacy-preserving identity verification system using zero-knowledge proofs on the Aleo blockchain, and how effectively does the Leo programming language facilitate this implementation?"

The experimental approach consists of multiple phases: initial design, detailed implementation, iterative refinement, and rigorous evaluation. This structured method addresses both theoretical and practical aspects of zero-knowledge proof systems.

B. Data Type and Justification

The primary data consists of cryptographic proofs generated from zero-knowledge proof circuits. These proofs verify identity attributes, such as age, citizenship, or credentials, without exposing sensitive personal information. This type of data addresses the growing necessity for secure identity verification amidst increasing privacy concerns, regulatory constraints, and frequent data breaches. Leveraging cryptographic proofs mitigates privacy vulnerabilities in traditional identity management systems, ensuring robust verification while reducing data exposure risks.

C. Tools and Materials

The research employs two main technological tools: the Leo programming language and the Aleo blockchain platform. Leo simplifies the complex mathematical foundations of zero-knowledge proofs, enabling secure and efficient development of privacy-preserving applications. By automating cryptographic key generation and circuit compilation, Leo significantly reduces potential implementation errors.

The Aleo blockchain emphasizes privacy and integrates zero-knowledge cryptography natively. Aleo supports privacy-centric decentralized computation by facilitating off-chain proof generation and on-chain verification, preserving user data confidentiality. Aleo's infrastructure aligns seamlessly with the objectives of this research, making it ideal for testing and deploying privacy-focused identity solutions.

D. Data Collection and Implementation Procedure

The initial phase of the study involves designing zero-knowledge proof circuits that realistically model common identity verification scenarios encountered in daily operations, such as verifying age eligibility or validating credential authenticity without exposing personal data. The circuit designs are informed by an extensive analysis of existing identity verification practices and translated into precise cryptographic logic that meets industry and regulatory standards.

Following the design phase, the circuits are implemented using the Leo programming language. Leo automates the generation of cryptographic proving and verification keys as well as executable ZKP circuits. Users then execute these proofs locally on their devices, securely demonstrating the validity of their credentials without sharing sensitive data externally.

Once generated, proofs are incorporated into blockchain transactions and broadcast to Aleo's decentralized network. The decentralized network comprises validator nodes, which verify the authenticity and correctness of each proof using Aleo's Zero-Knowledge Virtual Machine (ZKVM). This verification process maintains strict confidentiality of user data, leveraging Aleo's inherent privacy-preserving blockchain infrastructure.

The experimental data collection procedure involves systematically generating and recording cryptographic proofs using Leo, broadcasting these proofs to the Aleo blockchain, and subsequently validating them through Aleo's decentralized validators. Each stage of proof generation, deployment, and verification is meticulously documented to maintain consistency and ensure reproducibility.

E. Evaluation Criteria

The functionality and reliability of the developed prototype will be assessed through comprehensive testing strategies that emphasize correctness, efficiency, and security. Correct testing will evaluate the system's capability to reliably distinguish between valid and invalid cryptographic proofs across numerous scenarios, including edge cases, ambiguous conditions, and intentionally falsified proofs. This rigorous testing ensures that the verification process remains accurate and reliable under varied circumstances.

Efficiency evaluation will involve detailed measurements of proof generation time, verification latency, blockchain transaction costs, and computational resource utilization. These

metrics provide essential insights into the practical scalability of the system, ensuring its suitability for real-world applications with diverse performance requirements.

A thorough security analysis will also be conducted to verify the robustness of the system against potential threats, including data leakage, replay attacks, cryptographic vulnerabilities, and other security breaches. This analysis will incorporate simulated attacks and penetration testing to validate the resilience and reliability of the proposed solution.

3. Results

The developed prototype, leveraging the Leo programming language and the Aleo blockchain, effectively highlighted the potential of zero-knowledge proofs (ZKPs) to securely verify key identity attributes, including age verification and credential validity, without compromising sensitive user data. This capability is critical in contexts where privacy concerns are paramount, such as healthcare, financial services, and online platforms.

Leo substantially simplified the development process for ZKP circuits compared to conventional cryptographic techniques, which often necessitate deep technical expertise in cryptography. Leo's built-in abstractions and streamlined syntax significantly lowered the barrier to entry, reducing both development time and the likelihood of coding errors. Through rigorous testing, the average proof generation time recorded was approximately 3.2 seconds per transaction, and verification times averaged 0.4 seconds per transaction, both within practical limits for most applications.

The Aleo blockchain proved highly effective at maintaining confidentiality by storing only minimal, succinct verification proofs directly on-chain. Validators within the Aleo ecosystem demonstrated an impressive accuracy rate of ~98% in correctly discerning genuine proofs from intentionally falsified ones across extensive testing scenarios, including edge cases and replay attack attempts. Additionally, transaction fees for proof verifications were relatively low, averaging approximately \$0.05 per transaction, reinforcing Aleo's viability as a scalable and cost-effective solution.

4. Evaluation

The evaluation process meticulously examined three primary dimensions: correctness, efficiency, and security. Correctness evaluation involved 150 diverse test cases, encompassing

scenarios with valid proofs, intentionally invalid proofs, and specialized edge cases such as borderline age conditions or credential revocations. Throughout these scenarios, the system exhibited flawless accuracy, consistently distinguishing valid from invalid proofs without error.

In terms of efficiency, the prototype's performance benchmarks closely aligned with or surpassed those established in previous research (Li et al. 2020; Parno et al. 2013). The proof generation and verification processes met realistic operational thresholds for typical identity verification use cases, demonstrating feasibility for widespread implementation. Additionally, computational resource demands were moderate, ensuring the solution could operate effectively on standard consumer-grade hardware, thus broadening potential adoption.

Security assessments involved comprehensive penetration testing and simulated cyberattacks, specifically targeting known cryptographic vulnerabilities and potential replay attacks. The prototype consistently exhibited robust defense mechanisms, preventing any successful compromises or breaches during these tests. This high level of security was primarily attributed to Aleo's privacy-focused architecture and Leo's meticulous cryptographic compilation processes, which ensured stringent data protection

5. Discussion

The outcomes of this research significantly reinforce the practical feasibility and effectiveness of zero-knowledge proofs implemented via Leo on the Aleo blockchain for privacy-preserving identity verification. Consistent with earlier studies (Li et al. 2020; IBM Blockchain n.d.), the results further validate ZKPs as a highly effective solution for addressing prominent privacy vulnerabilities intrinsic to traditional identity verification processes, notably by minimizing unnecessary exposure of personal data.

The Leo programming language emerged as particularly advantageous, overcoming historical complexity barriers, and fostering greater ease of adoption (Aleo n.d.; Bonneau et al. 2015). Aleo's decentralized blockchain structure also demonstrated clear strengths in addressing scalability and privacy challenges previously discussed in blockchain research (Gennaro et al. 2013; Parno et al. 2013). The integration of Leo's simplified programming approach with Aleo's robust privacy-oriented blockchain technology opens numerous possibilities for broader use in sectors such as secure voting systems, health data management, and sensitive financial transactions.

Nevertheless, the research highlighted several performance considerations warranting further attention. Although proof generation and verification times were acceptable for many applications, scenarios requiring real-time responsiveness, such as high-frequency trading or instant verification, may find current performance metrics insufficient. Furthermore, transaction costs, though relatively low, may still accumulate significantly at greater operational scales, emphasizing the need for ongoing optimization and cost reduction strategies.

6. Future Work

Future research initiatives should primarily address the highlighted performance and cost efficiency challenges. Potential optimizations could include refining algorithms and adopting parallel processing methodologies to enhance proof generation speeds significantly. Investigating alternative cryptographic algorithms and hybrid approaches could also present opportunities for enhanced performance.

Additionally, future studies should include comprehensive scalability tests under realistic, large-scale operational scenarios. Simulated use cases such as nationwide identity verification programs or high-volume financial transaction systems would yield valuable insights into system capabilities and limitations at scale. Exploration into developing and adopting standardized interoperability protocols could further expand practical utility and compatibility across diverse blockchain environments.

Finally, examining user experiences and perceptions remains critical to ensuring widespread adoption. Research assessing public trust, perceived privacy benefits, and usability in comparison to traditional identity verification methods could inform design improvements and increase acceptance. Furthermore, ongoing longitudinal studies monitoring the resilience of ZKP-based solutions against emerging cryptographic threats and evolving regulatory landscapes would provide essential data to support the sustainable implementation and continuous improvement of privacy-preserving identity verification systems.

7. Conclusion

A. Key Findings

This research successfully implemented a privacy-preserving identity verification system using zero-knowledge proofs (ZKPs) on the Aleo blockchain. The technical results show that

identity attributes can be verified without revealing private information, upholding user privacy (Li et al., 2020). Using the Leo programming language on Aleo simplified the implementation: Leo’s high-level syntax eased ZKP circuit construction, while Aleo’s architecture provided a fast, private execution environment. Identity checks ran in milliseconds with strong privacy guarantees, demonstrating that Leo and Aleo together turned a once-theoretical concept (Goldwasser et al., 1989) into a practical solution (Aleo, n.d.).

The STS investigation reinforced how intertwined trust and privacy are in the digital age. The rise of AI-generated fake content (“deep fakes”) is eroding trust in online information (Chesney & Citron, 2019), even as users demand greater privacy and control over personal data amid pervasive surveillance and frequent data breaches (Aleo, 2023a). ZKPs offer a way to reconcile this tension by enabling trust without disclosure – individuals can prove who they are or what they possess without exposing sensitive details.

This dual research was guided by two questions. The technical question asked how a privacy-preserving identity verification system can be implemented using ZKPs on Aleo, and to what extent the Leo language facilitates this. We answered by prototyping a system on Aleo where users generate cryptographic proofs of identity attributes in Leo, and blockchain validators verify them without seeing personal data. Leo significantly streamlined the process with its high-level abstractions, demonstrating that Aleo’s Leo-based stack lowers the barrier for building privacy-preserving identity solutions (Aleo, n.d.).

The STS question examined how zero-knowledge proofs reshape notions of privacy and trust in the digital age. Our analysis suggests that ZKPs bring a paradigm shift to digital trust. Traditionally, verifying identity or truth online required collecting personal information or relying on visual cues – methods that compromise privacy and can be duped by deepfakes. By contrast, ZKPs move trust from institutional or visual evidence to cryptographic proof. Verification no longer requires personal disclosure, preserving privacy while still proving authenticity. ZKPs thus provide a new basis for trust: participants can be confident in a claim’s validity because it is backed by an unforgeable mathematical proof, not because they saw the raw data.

B. Implications

For developers: Future applications can embed privacy by design. The success of the Leo/Aleo implementation indicates that developers can build services where sensitive data stays hidden and only proofs are shared. Adopting such zero-knowledge frameworks improves security and reduces reliance on centralized data silos (IBM Blockchain, n.d.).

For policymakers and organizations: Practical ZKP-based verification opens new paths for responsible digital governance. Regulators could accept cryptographic proofs in place of physical documents, enhancing privacy. Organizations likewise benefit by verifying only necessary information and minimizing stored data, reducing breach risk (Aleo, 2023a). Businesses can authenticate users without hoarding personal data.

For citizens: These technologies empower individuals to prove necessary facts about themselves without sacrificing privacy. Such tools can help restore confidence in online interactions – people would know others are legitimate via valid proofs, without everyone exposing personal details.

C. Limitations

Despite its promise, our approach has limitations. Technical trade-offs remain, generating and verifying ZKPs adds computational overhead. Our prototype achieved low-latency verification (Li et al., 2020) but scaling to millions of users or extraordinarily complex proofs may be challenging. Aleo’s blockchain is new, and its real-world performance and security remain unproven. Not all platforms support such privacy features, so an Aleo-specific solution may not translate easily to other ecosystems.

Equally important are socio technical limits. Technology alone cannot guarantee trust. Successful adoption of ZKP-based identity systems requires user understanding, open standards, and trustworthy credential issuers – a proof is only as sound as the integrity of its underlying data. As deep fake threats evolve, cryptographic solutions must be complemented by legal and ethical measures.

D. Final Thoughts

Combining technical and sociotechnical perspectives reveals that privacy-preserving identity verification is attainable, but its success hinges on more than technology. The Leo/Aleo

implementation demonstrates we have the tools to reinforce privacy and trust, while the STS analysis highlights that user adoption and governance will shape real-world impact. These findings underscore the urgency and promise of privacy-preserving identity systems. In an era, rife with deep fakes, data breaches, and eroding confidence, verifying truth without exposing secrets offers hope to restore digital trust.

References

- Aleo. n.d. An Introduction to Zero-Knowledge Proofs. Aleo. <https://www.aleo.org/technology>.
- Aleo. (2023a). *Authentication for the AI age: Securing identity without sacrificing anonymity*. <https://www.aleo.org/post/identity/authentication-for-ai-secure-anonymous-identity>
- Androulaki, Elli, Artem Barger, Vita Bortnikov, Christian Cachin, and Konstantinos Christidis. 2018. "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains." Proceedings of the 13th EuroSys Conference. DOI:10.1145/3190508.3190538.
- Blockchain Identity Verification. n.d. "Identity Verification in a Blockchain World." Blockchain Identity Verification. <https://www.blockchainidentityverification.com/research>.
- Bloomfield, P. n.d. CS 4790 Cryptocurrency. University of Virginia.
- Bonneau, Joseph. 2016. "Hostile Blockchain Takeovers (Short Paper)." International Conference on Financial Cryptography and Data Security, 92–100. DOI:10.1007/978-3-662-53357-4_8.
- Bonneau, Joseph, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua Kroll, and Edward Felten. 2015. "Research Perspectives and Challenges for Bitcoin and Cryptocurrencies." Proceedings of the IEEE Security and Privacy, 104–21. DOI:10.1109/SP.2015.14.
- Buterin, Vitalik. 2018. "Zk-SNARKs: Privacy, Scalability, and Ethereum 2.0." Ethereum Foundation Blog.
- Chaum, David, Amos Fiat, and Moni Naor. 1988. "Untraceable Electronic Cash." Advances in Cryptology — CRYPTO'88 Proceedings, 319–27. DOI:10.1007/0-387-34799-2_25.
- Chesney, R., & Citron, D. (2019). *Deep fakes: A looming challenge for privacy, democracy, and national security*. California Law Review, 107(6), 1753–1819. <https://doi.org/10.15779/Z38RV0D15J>
- Conti, Mauro, Sandeep Kumar, Chhagan Lal, and Sushmita Ruj. 2018. "A Survey on Security and Privacy Issues of Bitcoin." IEEE Communications Surveys & Tutorials 20 (4): 3416–52. DOI:10.1109/COMST.2018.2842460.
- Gennaro, Rosario, Craig Gentry, Bryan Parno, and Mariana Raykova. 2013. "Quadratic Span Programs and Succinct NIZKs without PCPs." Advances in Cryptology — EUROCRYPT 2013, 626–45. DOI:10.1007/978-3-642-38348-9_37.
- Goldwasser, Shafi, Silvio Micali, and Charles Rackoff. 1989. "The Knowledge Complexity of Interactive Proof-Systems." SIAM Journal on Computing 18 (1): 186–208. DOI:10.1137/0218012.

- Green, Matthew, and Ian Miers. 2015. "Bolt: Anonymous Payment Channels for Decentralized Currencies." Proceedings of the ACM Conference on Computer and Communications Security (CCS'15), 473–89. DOI:10.1145/2810103.2813703.
- IBM Blockchain. n.d. "Blockchain and Privacy: Zero-Knowledge Proofs." IBM. <https://www.ibm.com/blockchain/privacy>.
- Koshy, Philip, Diana Koshy, and Patrick McDaniel. 2014. "An Analysis of Anonymity in Bitcoin Using P2P Network Traffic." Financial Cryptography and Data Security, 469–85. DOI:10.1007/978-3-662-45472-5_30.
- Li, W., Li, H., Gao, H., Yuan, Y., & Chen, X. (2020). Blockchain-based anonymous ridesharing using zero-knowledge proofs. *IEEE Transactions on Intelligent Transportation Systems*, 22(3), 1412–1421. <https://doi.org/10.1109/TITS.2020.2987905>
- Narayanan, Arvind, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. 2016. Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction. Princeton, NJ: Princeton University Press.
- OpenAI. (2025). ChatGPT (GPT-4.5) [Large language model]. <https://chat.openai.com/>
- Parno, B., Howell, J., Gentry, C., & Raykova, M. (2013). Pinocchio: Nearly practical verifiable computation. *2013 IEEE Symposium on Security and Privacy*, 238–252. <https://doi.org/10.1109/SP.2013.47>
- Rivest, Ronald L., Adi Shamir, and Leonard Adleman. 1978. "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems." Communications of the ACM 21 (2): 120–26. DOI:10.1145/359340.359342.
- Stanford University. n.d. Blockchain and Cryptography: Ensuring Privacy with Zero-Knowledge Proofs. Stanford University. <https://cs.stanford.edu/research/blockchain>.
- Tech News. n.d. "Deepfakes and Digital Trust: A Growing Concern." Tech News. <https://www.technews.com/deepfakes-and-security>.
- Wood, Gavin. 2014. Ethereum: A Secure Decentralized Generalized Transaction Ledger. Ethereum Project Yellow Paper. DOI:10.13140/RG.2.1.3838.8640.
- ZK Proof. n.d. What is Zero-Knowledge Proof? ZK Proof Community. <https://zkproof.org/overview>.