BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF INFORMATION TECHNOLOGY

Excel@FIT 2026
http://excel.fit.vutbr.cz

zk-Census: Privacy-Preserving Decentralized Polling via Zero-Knowledge Eligibility and Verifiable Identity

Tomáš Hanák*

Abstract
Traditional censuses and surveys rely on centralized agencies that operate as "black boxes," requiring absolute trust from participants. This model lacks transparency and is vulnerable to data manipulation, identity fraud, and privacy breaches. This paper proposes a decentralized system utilizing cryptographic proofs alongside a trusted execution environment (TEE) to ensure verifiable and private data collection. The architecture incorporates an independent identity management bridge to verify participant attributes while maintaining their anonymity. The implemented solution allows respondents to prove their eligibility based on demographic criteria while strictly enforcing a "one person, one vote" policy to prevent manipulation through Sybil accounts. The resulting platform provides a mathematically grounded alternative for conducting a census, replacing institutional trust with transparent and verifiable processes that protect both data integrity and user privacy.

*xhanakt00@stud.fit.vut.cz, *Faculty of Information Technology, Brno University of Technology*

1. Introduction

[Motivation] Public opinion polling and demographic data collection are critical components of modern strategic planning. However, the current landscape forces a compromise between data integrity and respondent privacy. Centralized survey agencies cannot be independently audited, meaning there is nothing to prevent them from inviting biased participants or deliberately tampering with the results. This model is inherently based on strong institutional trust, which we aim to eliminate. In an era where digital identities are easily forged, conducting a reliable census without exposing sensitive personal data remains a significant challenge.

[Problem definition] Because centralized agencies lack public transparency, respondents cannot independently verify if data was manipulated or omitted. Furthermore, preventing Sybil attacks (users submitting multiple fake responses) typically requires strict identity verification, which discourages participation due to privacy concerns. An ideal solution must guarantee mathematical verifiability, Sybil resistance, and absolute privacy.

[Existing solutions] Current solutions fail to balance privacy and integrity. Traditional platforms require absolute institutional trust, while standard Web3 voting completely lacks privacy. Various blockchain-based e-voting protocols, such as BBB-Voting [1], SBvote [2], Provotum [3], and the Open Vote Network (OVN) [4], achieve strong end-to-end verifiability and self-tallying properties. However, much like advanced zero-knowledge frameworks such as Semaphore [5], these systems generally lack a seamless bridge to real-world identity authorities required for efficiently filtering demographic attributes.

[Our solution] We propose a decentralized census utilizing zk-SNARKs [6] for primary privacy and a TEE for confidential tallying. While the TEE ensures that ongoing results remain hidden, the core guarantees of user anonymity and eligibility are mathematically enforced by zk-SNARKs. Participants receive signed credentials from a trusted Identity Authority, which are used locally to generate cryptographic proofs. This allows respondents to prove eligibility and cast votes on-chain without exposing their real identity.

[Contributions] This work replaces institutional trust with verifiable cryptography, securely binds real-world identity to anonymous voting, ensures strict Sybil resistance via zero-knowledge nullifiers, and guarantees the confidentiality of ongoing survey results.

2. Oasis Sapphire

The system is built on the Oasis Sapphire blockchain [7], which provides confidentiality by processing transactions within Trusted Execution Environments (TEE) (Figure 1). Unlike public ledgers where all states are visible, Sapphire executes encrypted transactions inside secure enclaves. This ensures that raw data remains hidden even from node operators, satisfying our requirement for maintaining tally confidentiality during an active census phase.

3. System Architecture

The decentralized census system operates in distinct phases, ensuring both data integrity through blockchain and user anonymity through cryptography. These phases are visualized on the poster.

3.1 Survey Creation

The lifecycle of a poll begins with the creator defining specific demographic criteria, such as a minimum age or a required geographical region. As shown in Figure 2, these parameters are embedded into a smart contract and deployed to the blockchain. This establishes an immutable and transparent set of rules for the given census.

3.2 Identity Provider

A functional decentralized census requires a trusted digital identity to verify respondent attributes without compromising privacy. Since a universal infrastructure is not yet available, the project is designed with the upcoming European Digital Identity (EUDI) framework [8] in mind for future adoption.

As illustrated in Figure 3, our current implementation utilizes MojeID [9] to authenticate that participants are real individuals. However, because MojeID does not natively provide ZK-friendly hashes, we employ an intermediate Identity Authority. This server retrieves the identity data and signs it along with the user's address using the Poseidon hash function [10]. This optimized signature allows the client to locally generate a zero-knowledge proof of eligibility while keeping all sensitive data private.

3.3 zk-SNARK Generation

Following the acquisition of a valid identity credential, the respondent generates a zero-knowledge proof locally on their device (Figure 4). This logic is implemented using Circom [11], a language for designing arithmetic circuits.

Inside the circuit, the application verifies the user's private identity attributes against the Authority's signature and public survey requirements (e.g., minimum age). If satisfied, it outputs a valid zk-SNARK proof, allowing the respondent to prove eligibility without exposing raw personal data to the network.

3.4 Voting & On-Chain Verification

Upon generating the proof, the respondent submits an encrypted vote alongside the zk-SNARK proof to the network (Figure 5). The smart contract takes over to verify the cryptographic proof and process the vote. To prevent Sybil attacks and double voting, the contract checks a uniquely generated nullifier. If the proof is valid and the nullifier is unused, the vote is accepted, while the actual choice remains hidden within the TEE.

3.5 Confidential Tallying & Results

The final phase, shown in Figure 6, utilizes the TEE to guarantee the confidentiality of the ongoing census. The smart contract is programmed to reject any requests for current standings while the voting is active. Only after the predefined voting period expires can the creator or respondents request the final outcome. At that point, the contract decrypts the state and returns the aggregated tally, ensuring that early results cannot manipulate the final outcome.

3.6 Performance and Cost Evaluation

On-chain measurements demonstrate that the proposed architecture is highly cost-effective and scalable (Table 1). Deploying the smart contract requires approximately 2.3 million gas (~$0.026), while the execution of a complex zk-SNARK verification and vote casting consumes approximately 415,000 gas. This keeps the transaction cost for voters well below $0.01 at current network rates.

4. Conclusions

This work presents a decentralized approach to public opinion polling and censuses that reduces the reliance on centralized agencies. By establishing zk-SNARKs as the primary layer for verifying eligibility and ensuring anonymity, and utilizing the Oasis Sapphire TEE strictly for confidential tallying, the implemented system provides a fully verifiable and privacy-preserving platform. A key component of this solution is the identity bridge, which securely connects real-world authentication with anonymous, Sybil-resistant on-chain voting.

References

[1] Ivan Homoliak, Zengpeng Li, and Pawel Szalachowski. Bbb-voting: Self-tallying end-to-end verifiable 1-out-of-k blockchain-based boardroom voting. In *Blockchain*, pages 297–306, 2023.

[2] Ivana Stančíková and Ivan Homoliak. Sbvote: Scalable self-tallying blockchain-based voting, 2022.

[3] Christian Killer, Bruno Rodrigues, Eder John Scheid, Muriel Franco, Moritz Eck, Nik Zaugg, Alex Scheitlin, and Burkhard Stiller. Provotum: A blockchain-based and end-to-end verifiable remote electronic voting system. In *2020 IEEE 45th Conference on Local Computer Networks (LCN)*, pages 172–183, 2020.

[4] Mohamed Seifelnasr, Hisham S. Galal, and Amr M. Youssef. Scalable open-vote network on ethereum. Cryptology ePrint Archive, Paper 2020/033, 2020.

[5] Privacy and Scaling Explorations (PSE). Semaphore: A zero-knowledge protocol for anonymous signaling on ethereum. Official Documentation, 2021. https://semaphore.pse.dev/.

[6] Junkai Liang, Daqi Hu, Pengfei Wu, Yunbo Yang, Qingni Shen, and Zhonghai Wu. SoK: Understanding zk-SNARKs: The gap between research and practice. Cryptology ePrint Archive, Paper 2025/172, 2025.

[7] Oasis Protocol Project. The oasis blockchain platform. Technical Report, 2020. https://docs.oasis.io/general/oasis-network/papers/.

[8] European Commission. European digital identity wallet (eudi). Official Documentation, 2025. https://eudi.dev/.

[9] CZ.NIC. Mojeid: Bezpečná digitální identita. Online, 2026. https://www.mojeid.cz/.

[10] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. Cryptology ePrint Archive, Paper 2019/458, 2019.

[11] Marta Belles-Munoz, Miguel Isabel, Jose Munoz-Tapia, Albert Rubio, and Jordi Baylina. Circom: A circuit description language for building zero-knowledge applications. *IEEE Transactions on Dependable and Secure Computing*, PP:1–18, 01 2022.