BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF INFORMATION TECHNOLOGY

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

Sybil-Resistant Identity Systems in Decentralized Environments

Michal Ľaš*

Abstract
Almost all identity management systems, whether centralized or decentralized, use a centralized approach to achieve Sybil resistance – that is, to prevent the creation of a large number of fake identities. Typically, these systems require verification via a phone number or government-issued ID. **This project aims to design a decentralized identity system that ensures Sybil resistance while preserving user anonymity and sensitive data.** The proposed solution utilizes facial biometric verification and social graph analysis. To ensure decentralization and protect the privacy of sensitive data, the solution employs blockchain technology, Zero-Knowledge Proofs (ZKPs), and Trusted Execution Environments (TEEs). Although the resulting system serves its purpose, the price of decentralization is a slow bootstrapping process for users and greater overall system complexity. The main advantage is decentralization. This solution is one of the few viable alternatives to the currently dominant centralized solutions.

*xlasmi00@stud.fit.vut.cz, *Faculty of Information Technology, Brno University of Technology*

1. Introduction

In today's digital age, almost everyone has some form of digital identity. Identity management is typically centralized, with a single entity maintaining a database of all users, managing registration, and controlling the entire system, including user data. An alternative is Decentralized Identifiers (DIDs) [1], standardized by the World Wide Web Consortium (W3C). DIDs enable individuals to manage their sensitive data in accordance with the principle of self-sovereignty. However, DIDs alone cannot ensure Sybil resistance – that is, prevent the creation of a large number of fake identities. Yet many services require this property, such as electronic voting, online auctions, or social networks. In practice, centralized solutions are almost exclusively used, where such systems require a phone number or government ID for initial verification. In a decentralized world, there are only a few solutions to this problem, and they usually fall short of protecting privacy or being fully decentralized.

The goal is not to replace DIDs, but to design an identity system that, when combined with DIDs, would ensure Sybil resistance. There are several parameters that such a system should meet. The most important ones are anonymity and unlinkability to real-world identity, decentralization, transparency, inclusivity, and scalability. Ideally, the system should not allow anyone to own more than one virtual identity or to link all virtual identities of any real person, so that during verification, it is possible to check whether a single real person is deceiving the system.

Few similar systems exist, and research summarizing them is limited. The most notable examples include World Network [2] and BrightID [3]. World Network uses a special device called the Orb, which scans people's irises and stores the data in a database. After successful biometric verification, a person receives a special token that proves their uniqueness. The main drawback is the centralization of the hardware (the Orb) and the server, and the fact that biometric data is highly sensitive personal information. BrightID uses social graph analysis to detect Sybil entities. Users who trust each other form connections that make up a social graph. The disadvantage is that the graph is public, and even though entities appear under pseudonyms, it can reveal social ties.

The solution proposed in this project utilizes facial biometrics and private social graph analysis. The primary focus is on preserving the privacy of biometrics and social connection data through ZKPs and TEEs.

2. Architecture

The proposed system uses the Oasis protocol [4, 5], a blockchain that executes smart contracts in TEEs and ZKPs to provide recursive proofs of random paths in a social graph without requiring knowledge of the entire graph. Key components include:

- **Oasis Sapphire:** Acts as confidential smart contract layer. It holds the encrypted state.
- **Oasis ROFL (The Off-chain Compute Engine):** Acts as stateless, verifiable compute layer.
- **Database:** Storage of encrypted facial biometrics embeddings.
- **Blokchain or L2:** Cheaper blockchain or L2 for storing commitments in a Merkle Tree.

The system works by requiring users to register, connect with other users, and build a social graph, so they can subsequently demonstrate random paths within it to prove they are well connected to other honest users.

3. Registration

The registration process is shown in Figure 1. The user captures a 3D face scan with their phone and C2PA manifest [6]. This 3D scan is encrypted using ROFL's public key and sent to ROFL for verification. ROFL decrypts the 3D scan within the TEE, verifies the signature and the validity of the C2PA manifest, performs liveness checks, checks that the timestamp in the metadata is not too old, calculates the nullifier, and verifies that this scan has not yet been used. It then creates an embedding from the 3D scan, queries the DB to verify that the user is not yet registered, and inserts the new embedding into the DB. ROFL returns the signed result and the embedding to the user, who saves the embedding locally and uses the signed result to create a profile on the blockchain using Sapphire.

Since facial recognition can fail, the user will still be registered, albeit in a limited capacity, and will need to obtain endorsements from other users to complete the registration process.

4. Connection Establishment

The connection establishment process is illustrated in the Figure 2. To establish a connection, two people must meet in person, take a photo of each other with the C2PA manifest, a mutually signed challenge, and exchange their saved embeddings, along with any other data, such as identifiers. Both will use ROFL for mutual verification. ROFL verifies that the hashes of the provided embeddings for both entities are recorded on Sapphire in their profiles, verifies the C2PA manifest, re-verifies the timestamp and nullifier as during registration, and matches the faces in the photos against the provided embeddings. If everything passes, ROFL generates a signed connection commitment.

At least one user submits the ROFL-signed commitment to Sapphire. Sapphire verifies the mutual user challenge and the ROFL signature and checks whether the connection already exists. Since the connection commitment is just a hash of identifiers, Sapphire blinds the commitment with a secret. Finally, Sapphire emits an event that the relayer listens for, and the relayer adds the blinded commitment to the Merkle tree. The blinded commitment is also returned to the user if the relayer fails.

5. Private Random Walks

Using connection commitments, it is possible to generate random walks through a social graph even though the graph itself is not stored anywhere, and each user knows only their own connections. Each user can prove a connection to another person without revealing that person's identity. These proofs are shared with one another and can be used recursively in subsequent proofs, thereby extending the path. Essentially, this is an Incrementally Verifiable Computation (IVC).

For example, entity A is connected to entity B, and B is connected to C. A can create a proof that they know a path to entity B. B takes this proof and confirms that they are connected to A, and that they also know a path to entity C. This process is illustrated in Figure 3 for four users.

Users must thus demonstrate the existence of multiple paths that do not pass through the same individuals. The principle of this method is based on graph algorithms for Sybil analysis. It is assumed that Sybil entities will be weakly connected to the honest part of the graph and will not be able to obtain enough distinct paths.

6. Conclusion

This project has successfully demonstrated a design for a robust, decentralized identity management system that exhibits a high level of resistance to Sybil attacks. The system safeguards sensitive user data while maintaining anonymity and unlinkability to real-world identities. While being accessible to anyone with a mobile phone.

References

[1] Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Orie Steele, and Christopher

Allen. Decentralized identifiers (dids) v1.0. Technical report, World Wide Web Consortium, 7 2022.

[2] World. World whitepaper. Technical report.

[3] BrightID. Brightid. whitepaper, 1 2022.

[4] Oasis Protocol Foundation. The oasis blockchain platform. Technical report, Oasis Labs, 6 2020. White paper.

[5] Oasis Protocol Foundation. Oasis network documentation. online, 2026.

[6] Coalition for Content Provenance and Authenticity. C2pa specification. online, 2026. Version 1.3.