

Unlocking the Potential of Decentralised Digital Identification Systems for Smart Cities

Carla O. Castanho^{1,2}, Rafael Z. Frantz¹, Marcelo P. Chequin², Sandro Sawicki¹
Fabricia Roos-Frantz¹, Carlos Molina-Jimenez³, Jon Crowcroft³, Timothy Hobson⁴

¹Universidade Regional do Noroeste do Rio Grande do Sul
Ijuí – RS – Brazil

²Universidade Regional Integrada do Alto Uruguai e das Missões
Santiago – RS – Brazil

³University of Cambridge, UK

⁴The Alan Turing Institute, UK

{rzfrantz, sawicki, frfrantz}@unijui.edu.br

{carla.castanho, 102513}@urisantiago.br

{carlos.molina, jon.crowcroft}@cl.cam.ac.uk, thobson@turing.ac.uk

***Abstract.** We argue that the absence of central authorities and trust decentralisation in decentralised ID systems enable them to achieve openness, transparency, scalability, privacy and reliability. These are highly desirable properties in several application domains including smart city services. Decentralised ID systems are only emerging. Several candidates exist at an experimental stage of development, awaiting deployment and evaluation in realistic applications. In this paper, we discuss Trustchain. We explain the features that have motivated us to use it in the implementation of smart city services. Also, we share our experience gained from its local deployment.*

1. Introduction

We regard a digital ID (ID for short) as a unique string **created** and **associated (bound)** to a unique entity. Examples of entities are persons, devices and services. From here on, to simplify the discussion we will assume that the entity is a human, for example, Alice or Bob. Typically, the string is a data structure that includes several attributes such as name and address and even unique immutable data of the individual, for example, biometric information. Alice used her ID to identify herself: she presents it to another identity to claim that she is Alice. For example, she presents it to Bob (the administrator of a remote service) to gain access to Bob's service under the claim that she is entitled to because she is Alice. Bob grants access to Alice but only if she satisfies **authentication**: she is able to prove that she is Alice. Alice's string is not necessarily eternal, it might be subject to **revocation** (cease being valid) after certain time, after the occurrence of a number of events or after explicit revocation.

In the text above we have written in bold the fundamental operations associated to Alice's ID. Their implementation and execution is far from trivial, to the extent that in practice, fully fledged ID systems are needed to manage them. There are several of them with different architectures. The latter determines how these operations are implemented, for example, who (which entity) is responsible for creating and binding IDs to entities and

how authentication is performed. The architecture influences several properties of the ID system including trustworthiness, openness, transparency, scalability and privacy.

The corner stone of an ID system is the mechanism that creates and binds the IDs to the entities. It is called the root of trust. To a great extent, the trustworthiness of an ID system depends of how the root of trust is created. There are two approaches: centralised or decentralised. The centralised approach is widely used in the current Internet. The best example is the Web Public Key Infrastructure (Web PKI) which relies on certificates signed directly or through delegation of authority by a monopolising small group of certification authorities including DigiCert and Verisign that act as a centralised root of trust. Other examples are the Facebook ID and the Google ID. The salient feature of centralised ID systems is the inclusion of a centralisation entity (normally the owner of the ID system) that operates as a root of trust and, as such, is a central authority that centralises trust. The intrusive participation of the central authority is the source of several undesirable features (e.g. lack of openness and transparency; and data control for abusive monetisation) that make us hesitate about using them in our smart city applications. We believe that decentralised IDs match this requirement better than centralised alternatives.

In this paper we explain how decentralised ID systems offer a better fit for our requirements. Although they are only emerging, several systems have already been implemented including uPort [Panait et al. 2020], Veramo [Veramo 2025] and Trustchain [Hobson et al. 2023]. However, they are still at laboratory stage and awaiting realistic deployment, testing and evaluation. To help cover the gap, in this paper we contribute the experience gained from the installation and preliminary evaluation of Trustchain. Several Trustchain's features motivated our choice, in particular its dependence only on existing blockchain infrastructures. The remainder of this paper is structured as follows: Section 2 discusses digital ID systems similar to Trustchain; Section 3 introduces concepts that are central to our discussions; In Section 4 we share the experience gained from the installation of Trustchain; and Section 5 closes our discussion and points to future work.

2. Related work

Several decentralised ID systems have been suggested and implemented. An early system is uPort [Panait et al. 2020]. It was implemented on top of the Ethereum blockchain to enable users to securely and privately control their digital identities and personal data. Sovrin [Sovrin 2025] is another pioneer. We have ruled it out because its operation depends on the implementation and maintenance of a dedicated blockchain to support the decentralised ID systems. Without a clear economic incentive to motivate the participants that maintain the decentralised ledger, the Sovrin blockchain cannot be expected to survive in the long term. In fact, it is likely that it will soon be shut down. Hyperledger Indy [Indy 2025] is another decentralised ID system that suffers from a related, but distinct, problem. Every implementation requires the deployment of its own blockchain. The software is open source, but the infrastructure to run it and the institutional actors that will govern and maintain it must be provided by the community. Anyone wanting to deploy a system using Indy must first identify a consortium of trustworthy legal entities that are willing to spend significant resources to set up and maintain the network.

The General Data Protection Act (LGPD - Brazil) and the General Data Protection Regulation (GDPR - EU) impose stringent requirements for the processing of personal data. Decentralised identity systems — such as uPort, Sovrin, and Hyperledger Indy — adopt different approaches to ensure compliance: uPort prioritises user autonomy but con-

flicts with blockchain immutability. Sovrin leverages AnonCreds to enhance privacy, and is well suited to global applications. Hyperledger Indy enables bespoke private networks, but suffers from implementation complexity. To remain compliant, all systems store sensitive data off-chain and use zero-knowledge proofs (ZKPs) to guarantee privacy.

3. Decentralised Digital Identity

We define an **entity** as a person, organisation, computer application, or device connected to the Internet. Entities possess characteristics (also called **attributes**). In an ID system, each entity has a **Digital Identity (DI)** that includes information that can be used to identify and authenticate an entity programmatically. Figure 1 shows the key concepts used by decentralised ID systems to manage their DIDs¹.

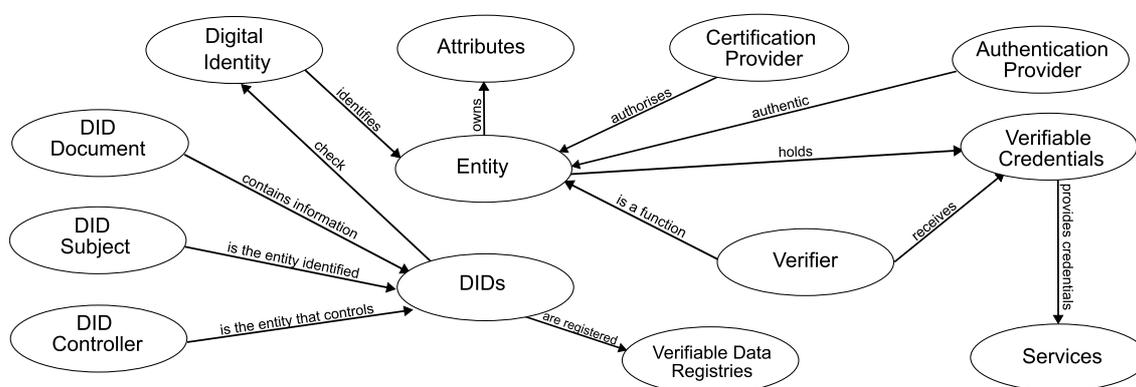


Figure 1. Key Concepts in Decentralised Digital Identities

The **Authentication Provider** is an entity or organisation that can be trusted to verify whether a credential is valid and has not been revoked. **Certification Providers** are entities or organisations responsible for issuing credentials. **Services** are provided by service providers to end users which are required to present credentials to be granted access. These credentials, are called **Verifiable Credentials** and are equivalent to physical credentials, such as passports or driving licences. A **Verifier** is an entity responsible for validating verifiable credentials.

Decentralised Identifiers (DIDs) are unique and used for identification. The **DID subject** is the entity identified by the DID. The **DID controller** is the entity (a person, organisation, or autonomous software) that has the authority to modify a DID document. **DID documents** are essentially decentralised public key certificates. They contain specific information associated to the DID subject: their cryptographic public keys and web URLs (or “service endpoints”). **Verifiable Data Registry** is a system that provides the facilities for the registration of DIDs for the creation and storage of DID documents.

4. Trustchain for DIDs

Trustchain is free and open-source software designed for the development of decentralised digital identity systems. The primary innovation of Trustchain lies in its ability to establish a Public Key Infrastructure (PKI) that supports digital identification without requiring trusted intermediaries to manage or maintain the infrastructure. As a result, the only trusted entities within the network are institutions recognised and legitimised

¹<https://www.w3.org>

by the user community itself — namely, organisations that already possess the credibility to issue credentials or provide services. This decentralised approach eliminates reliance on third parties, thereby enhancing both the security and autonomy of the system [Hobson et al. 2023].

4.1. Main Features

In addition to eliminating the need for trusted third parties to maintain the digital infrastructure, Trustchain offers additional advantages that integrate security, efficiency, and respect for user privacy [Hobson et al. 2023]:

- **Open access:** Anyone can deploy Trustchain without requiring prior authorisation.
- **Low cost:** The software is free and does not require the creation of a new blockchain, thereby reducing operational costs and complexity.
- **High security:** The proof-of-work mechanism makes the root of trust independently verifiable by every user, and counterfeiting practically impossible.
- **Transparency:** Verifiers fully see the legal entities in each process.
- **Privacy:** Credentials are stored locally on the holder’s device, ensuring greater control over personal information.

Trustchain distinguishes itself from other decentralised identification systems through two key advantages. Firstly, the peer-to-peer infrastructure it employs already exists and operates autonomously, regardless of the specific use case for digital identification. This eliminates the need to develop or maintain additional infrastructure, thereby simplifying implementation and reducing costs. Secondly, users can verify the authenticity and integrity of received information directly on their own devices, without relying on intermediaries or trusted third parties. These combined features enhance efficiency, security, and autonomy, thereby strengthening trust in the system as a whole.

The primary technical distinction of Trustchain is that its security relies on proof-of-work (PoW) generated through mining on the Bitcoin network, which produces approximately 800 EH/s. At this scale, it becomes computationally and economically infeasible, even for large corporations, to retroactively alter the Bitcoin blockchain, as the cost of an attack increases linearly over time. Other decentralised identification systems do not use PoW, instead opting for mechanisms such as proof-of-stake (PoS) or proof-of-authority (PoA), which prevent users from independently verifying the authenticity of public keys. In these systems, trust in the digital infrastructure and its operators is required, as verification depends on digital signatures, which require a pre-existing public key for validation. In contrast, Trustchain enables users to directly verify the authenticity and publication date of cryptographic keys, eliminating the need to rely on third parties or the underlying infrastructure [Hobson et al. 2023].

4.2. How to deploy Trustchain

Trustchain runs in standard operating systems. We have successfully deployed it on a virtual machine running Linux Debian 12, instantiated in a server located in our lab of Unijui. We will describe now the installation procedure and the tests conducted to verify correctness of the installation and to get a feeling of Trustchain’s capabilities.

We installed ION (Identity Overlay Network), the Rust programming language, and Trustchain itself. Trustchain uses ION as its DID method implementation, meaning it is used to perform operations for creating and publishing the DIDs. No trust is placed

in the ION software because the result of each DID operation is subsequently verified cryptographically. ION executes read and write operations on the Bitcoin ledger and the distributed Inter Planetary File System (IPFS). Trustchain also needs Node.js with its package manager and MongoDB which we installed successfully. To demonstrate Trustchain's capabilities we executed the following steps:

We set up a Wallet and collected tBTC tokens: We collected some tBTC (fake cryptocurrency used for software development) after setting up a Wallet, to pay for transactions: DID documents are stored in IPFS (at zero cost) and are simultaneously recorded in the Bitcoin ledger by embedding a hash of the DID content in a Bitcoin transaction, which incurs a transaction fee. For testing purposes we used the Bitcoin Testnet whose native token tBTC serves as payment for transactions.

We published our Root DID: A legally authorised entity is designated as the root DID subject. This involves publishing a DID containing a set of cryptographic public keys belonging to that entity. The root authority discloses the calendar date on which its DID was published, a detail that can be independently verified through the proof-of-work mechanism. In a smart cities context, various communication channels may be used to publicise the date, such as newspapers, billboards, advertisements, TV and radio.

We issued Downstream DIDs signed by the Root DID: The root entity now has the capability to sign the DIDs of other legal entities, creating verifiable public key certificate chains known as "downstream DIDs" (dDIDs). Each credential issuing authority is assigned a dDID. To do this they run their own Trustchain node and make a request to the upstream entity, invoking a challenge-response protocol to verify their identity and the dDID content. The only critical consideration is that all nodes on the network must be configured with the same root DID timestamp.

We set up a HTTP server for issuance and verification: Trustchain incorporates a built-in HTTP server for the issuance and verification of credentials via an HTTP API. The credential issuer's dDID is used as a parameter in the configuration. Furthermore, it is necessary to define the Fully Qualified Domain Name (FQDN) of the HTTP server, which serves as the endpoint for the mobile application. To ensure secure communication, Transport Layer Security (TLS) was implemented; this requires access to port 443 that our network administration agreed to open for us. We used CertBot to generate a Let's Encrypt certificate for the issuing server to ensure encrypted and reliable connections.

We installed the Trustchain Mobile app on an Android device: Credential holders and verifiers only need to install the Trustchain Mobile app. Upon launching the app, a personal DID is created from a newly generated public-private key pair. The user also inputs the root DID date, thereby enabling accurate identification of the legitimate root DID. With this configuration, the app verifies the signature and service endpoints (URLs) of any entity whose DID forms part of the user's trust network.

We run a practical experiment with credential issuance: We developed a practical example that simulates the verification of a digital driving licence. In this scenario, a new credential conforming to the ISO-18013 standard was generated by the dDID subject playing the role of vehicle licensing agency. The HTTP server generates a QR code which the user scans using the mobile wallet app, directing them to the credential issuance API endpoint. Because the server endpoint is included in the issuer's dDID document, the user can verify it before agreeing to any interaction. Upon successful verification, the user can trust the identity of the issuing authority and, consequently, accept their credential.

We verified the credential: Subsequently, a second user can perform device-to-device verification. The first user selects the attributes they wish to share, and the application generates a verifiable presentation where the non-disclosed attributes are redacted while preserving the valid signature of the issuer. This presentation is displayed as a QR code. Any other device equipped with the Trustchain Mobile app may scan the code to confirm the authenticity of the credential, thereby ensuring both privacy and security throughout the information-sharing process.

5. Conclusions and Future Work

We have presented Trustchain (a decentralised ID system) and highlighted its main features (e.g., openness, low cost and high security). We installed it successfully after about 60 hours of work conducted by an undergrad and a PhD student without previous experience in blockchain but with remote guidance from one of the Trustchain’s developers. Most of the installation hurdles were due to documentation gaps—we have covered them now in the Git documentation. We also faced and solved technical problems that emerged mainly from erroneous configurations (e.g. of databases and network connections) generated by default configuration files. We estimate that with this experience, a re-installation would take about 20 hours without remote assistance. We have shown that Trustchain is ready for practical work and all on top of open source existing technologies—which saved long hours of programming and coordination effort, for example, to implement a brand new decentralised ledger. The tests reported here are only a proof of concept. The next steps is to conduct more demanding and realistic tests to validate Trustchain scalability. For example, can it scale up to hundreds or thousands of DIDs? How many links can a chain of downstream DIDs contain and how long does it take to verify a credential? We will find the answers to these and similar questions in several applications that we have in our to-do list. We will use it to issue DIDs to students and staff of a smart campus that we are creating at Unijui and for issuing credentials to citizens of Santa Rosa city to access smart city services.

Acknowledgements

Research partially funded by the Co-ordination for the Brazilian Improvement of Higher Education Personnel (CAPES) and the Brazilian National Council for Scientific and Technological Development (CNPq) under project grants 311011/2022-5, 309425/2023-9, 402915/2023-2. EPSRC/EP/X015785/1 (G115169) grant funded Carlos and Jon.

References

- Hobson, T., France, L., Greenbury, S., Hare, L., and Wochner, P. (2023). Trustchain – trustworthy decentralised public key infrastructure for digital credentials. *IET Conference Proceedings*, 2023(14):31–40.
- Indy, H. (2025). Hyperledger indy. <https://www.lfdecentralizedtrust.org/projects/hyperledger-indy>.
- Panait, A.-E., Olimid, R. F., and Stefanescu, A. (2020). Analysis of uport open, an identity management blockchain-based solution. In Gritzalis, S., Weippl, E. R., Kotsis, G., Tjoa, A. M., and Khalil, I., editors, *Trust, Privacy and Security in Digital Business*, pages 3–13, Cham. Springer International Publishing.
- Sovrin (2025). Sovrin foundation. <https://sovrin.org>.
- Veramo (2025). Performant and modular apis for verifiable data and ssi. <https://veramo.io>.