

# A Decentralised Digital Identity Architecture

**Geoff Goodell**

Centre for Blockchain Technologies  
University College London  
g.goodell@ucl.ac.uk

**Tomaso Aste**

Centre for Blockchain Technologies  
University College London  
t.aste@ucl.ac.uk

## Abstract

Current architectures to validate, certify, and manage identity are based on centralised, top-down approaches that rely on trusted authorities and third-party operators. We approach the problem of digital identity starting from a human rights perspective, asserting that individual persons must be allowed to manage their personal information in a multitude of different ways in different contexts and that to do so, each individual must be able to create multiple unrelated identities. Therefore, we first define a set of fundamental constraints that digital identity systems must satisfy to preserve and promote human rights. With these constraints in mind, we then propose a decentralised, standards-based approach, using a combination of distributed ledger technology and thoughtful regulation, to facilitate many-to-many relationships among providers of key services. Our proposal for digital identity differs from others in its approach to trust in that we do not seek to bind credentials to each other or to a mutually trusted authority to achieve strong non-transferability. Because the system does not implicitly encourage its users to maintain a single aggregated identity that can potentially be constrained or reconstructed against their interests, individuals and organisations are free to embrace the system and share in its benefits.

## 1 Motivation

Establishing meaningful credentials for individuals and organisations in an environment in which the various authorities are not uniformly trustworthy presents a problem for currently deployed services, which are often based on hierarchical trust networks, all-purpose identity cards, and other artifacts of the surveillance economy. In the context of interactions between natural persons, identities are neither universal nor hierarchical, and a top-down approach to identity generally assumes that it is possible to impose a universal hierarchy. Consider “Zooko’s triangle,” which states that names can be global, secure, or human-meaningful, but not all three [1]. The stage names of artists may be globally unique and human-meaningful but are not really secure since they rely upon trusted authorities to resolve conflicts. The names that an individual assigns to friends or that a small community assigns to its members (“petnames” [2]) are secure and human-meaningful but not global. We extend the reasoning behind the paradox to the problem of identity itself and assert that the search for unitary identities for individual persons is problematic. It is technically problematic because there is no endogenous way to ensure that an individual has only one self-certifying name [3], there is no way to be sure about the trustworthiness or universality of an assigned name, and there is no way to ensure that an individual exists only within one specific community. More importantly, we assert that the ability to manage one’s identities in a multitude of different contexts, including the creation of multiple unrelated identities, is an essential human right.

This article is organised as follows. In the rest of this section, we offer some background on identity systems; we frame the problem space and provide examples of existing solutions. In Section 2, we introduce a set of constraints that serve as properties that a digital identity infrastructure must have to support human rights. In Section 3, we describe how a digital identity system with a fixed set of actors might operate and how it might be improved. In Section 4, we introduce distributed ledger technology to promote a competitive marketplace for issuers and verifiers of credentials and to constrain the interaction between participants in a way that protects the privacy of individual users. In Section 5, we consider how the system should be operated and maintained if it is to satisfy the human rights requirements. In Section 6, we suggest some potential use cases, and in Section 7 we conclude.

## 1.1 Manufacturing Trust

The current state-of-the-art identity systems, from technology platforms to bank cards, impose asymmetric trust relationships and contracts of adhesion on their users, including both the ultimate users as well as local authorities, businesses, cooperatives, and community groups. Such trust relationships, often take the form of a hierarchical trust infrastructure, requiring that users accept either a particular set of trusted certification authorities (“trust anchors”) or identity cards with private keys generated by a trusted third party. In such cases, the systems are susceptible to socially destructive business practices, corrupt or unscrupulous operators, poor security practices, or control points that risk coercion by politically or economically powerful actors.

Often, asymmetric trust relationships set the stage for security breaches. Rogue certification authorities constitute a well-known risk, even to sophisticated government actors [4, 5], and forged signatures have been responsible for a range of cyber-attacks including the Stuxnet worm, an alleged cyber-weapon believed to have caused damage to Iran’s nuclear programme [6], as well as a potential response to Stuxnet by the government of Iran [7]. Corporations that operate the largest trust anchors have proven to be vulnerable. Forged credentials were responsible for the Symantec data breach [8], and other popular trust anchors such as Equifax are not immune to security breaches [9]. Google has published a list of certification authorities that it thinks are untrustworthy [10], and IT administrators have at times undermined the trust model that relies upon root certification authorities [11]. Finally, even if their systems are secure and their operators are upstanding, trust anchors are only as secure as their ability to resist coercion, and they are sometimes misappropriated by governments [12].

Such problems are global, affecting the developed world and emerging economies alike. Identity systems that rely upon a single technology, a single implementation, or a single set of operators have proven unreliable [13, 14]. Widely-acclaimed national identity systems, including but not limited to X-Road in Estonia [15] and Aadhaar in India [16], are characterised by centralised control points, security risks, and surveillance.

Recent trends in technology and consumer services suggest that concerns about mobility and scalability will lead to the deployment of systems for identity management that identify consumers across a variety of different services, with a new marketplace for providers of identification services [17]. In general, the reuse of credentials has important privacy implications as a consumer’s activities may be tracked across multiple services or multiple uses of the same service. For this reason, the potential for a system to collect and aggregate transaction data must be evaluated whilst evaluating its impact on the privacy of its users.

While data analytics are becoming increasingly effective in identifying and linking the digital trails of individual persons, it has become correspondingly necessary to defend the privacy of individual users and implement instruments that allow and facilitate anonymous access to services. This reality was recognised by the government of the United Kingdom in the design of its GOV.UK Verify programme [19], a federated network of identity providers and services. However, the system as deployed has significant technical shortcomings with the potential to jeopardise the privacy of its users [20, 21], including a central hub and vulnerabilities that can be exploited to link individuals with the services they use [22].

Unfortunately, not only do many of the recently-designed systems furnish or reveal data about their users against their interests, but they have been explicitly designed to do so. For example, consider digital rights management systems that require users to identify themselves *ex ante* and then use digital watermarks to reveal their identities [23]. In some cases, demonstrable privacy has been considered an undesirable feature and designs that protect the user’s identity intrinsically are explicitly excluded, for example in the case of vehicular ad-hoc networks [18], with the implication that systems without exceptional access features are dangerous. Finally, of particular concern are systems that rely upon biometrics for identification. By binding identification to a characteristic that users (and in most cases even governments) cannot change, biometrics implicitly prevent a user from transacting within a system without connecting each transaction to each other and potentially to a permanent record. In recent years, a variety of US patents have been filed and granted for general-purpose identity systems that rely upon biometric data to create a “root” identity linking all transactions in this manner [24, 25].

## 1.2 Decentralised Approaches

The prevailing identity systems commonly require users to accept third parties as trustworthy. The alternative to imposing new trust relationships is to work with existing trust relationships by allowing users, businesses, and communities to deploy technology on their own terms, independently of external service providers. We propose a system-level approach to allow existing institutions and service providers to retain their relative authority and decision-making power without forcibly requiring them to cooperate with central authorities (such as governments and institutions), service providers (such as system operators), or the implementors of core technology. Our approach would encourage existing institutions and service providers to be self-sufficient, while concordantly allowing groups such as governments and consultants to act as advisors and regulators, not operators.

Modern identity systems are used to coordinate three activities: identification, authentication, and authorisation. The central problem to address is how to manage those functions in a decentralised context with no universally trusted authorities. Rather than trying to force all participants to use a specific new technology or platform, we suggest using a multi-stakeholder process to develop common standards that define a set of rules for interaction. Any organisation would be able to develop and use their own systems that would interoperate with those developed by any other organisation without seeking permission from any particular authority or agreeing to deploy any particular technology.

A variety of practitioners have recently proposed using a distributed ledger to decentralise the administration of an identity system [26], and we agree that the properties of distributed ledger technologies are appropriate for the task. In particular, distributed ledgers allow their participants to share control of the system. They also provide a common view of transactions ensuring that everyone sees the same transaction history.

Various groups have argued that distributed ledgers might be used to mitigate the risk that one powerful, central actor might seize control under the mantle of operational efficiency. However, it is less clear that this lofty goal is achieved in practice. Existing examples of DLT-enabled identity management systems include the following:

- *ShoCard* [27] is operated by a commercial entity that serves as a trusted intermediary [26].
- *Sovrin* [28] relies on a foundation to manage the set of approved certification authorities [29], and it is not clear that it could manage them with equanimity.
- *uPort* [30] does not rely upon a central authority, instead allowing for mechanisms such as social recovery. However, its design features an optional central registry that might introduce a means of linking together transactions that users would prefer to keep separate [26].

Researchers have proposed alternative designs to address some of the concerns.

A design suggested by Kaaniche and Laurent does not require a central authority for its blockchain infrastructure but does require a trusted central entity for its key infrastructure [31]. Coconut, a selective disclosure credential scheme, is designed to be robust against malicious authorities and may be deployed in a way that addresses such concerns [32]. We find that many systems such as these require users to bind together their credentials *ex ante* to achieve non-transferability, essentially following a design proposed by Camenisch and Lysyanskaya [33] that establishes a single “master key” that allows each user to prove that all of her credentials are related to each other. Even if users were to have the option to establish multiple independent master keys, service providers or others could undermine that option by requiring proof of the links among their credentials.

We imagine that the core technology underpinning these and similar approaches might be adapted to implement a protocol that is broadly compatible with what we describe in this article. However, we suspect that in practice they would need to be modified to encourage users to establish multiple, completely independent identities. In particular, service providers would not be able to assume that users have bound their credentials to each other *ex ante*, and if non-transferability is required, then the system would need to achieve it in a different way.

### 1.3 Participants in an Identity System

We shall use the following notation to represent the various parties that interact with a typical identity system:


- (1) A “*certification provider*” (**CP**). This would be an entity or organisation responsible for establishing a credential based upon foundational data. The credential can be used as a form of identity and generally represents that the organisation has checked the personal identity of the user in some way. In the context of digital payments, this might be a bank.
- (2) An “*authentication provider*” (**AP**). This would be any entity or organisation that might be trusted to verify that a credential is valid and has not been revoked. In current systems, this function is typically performed by a platform or network, for example a payment network such as those associated with credit cards.
- (3) An “*end-user service provider*” (**Service**). This would be a service that requires a user to provide credentials. It might be a merchant selling a product, a government service, or some other kind of gatekeeper, for example a club or online forum.
- (4) A *user* (**user**). This would be a human operator, in most cases aided by a device or a machine, whether acting independently or representing an organisation or business.

As an example of how this might work, suppose that a user wants to make an appointment with a local consular office. The consular office wants to know that a user is domiciled in a particular region. The user has a bank account with a bank that is willing to certify that the user is domiciled in that region. In addition, a well-known authentication provider is willing to accept certifications from the bank, and the consular office accepts signed statements from that authentication provider. Thus, the user can first ask the bank to sign a statement certifying that he is domiciled in the region in question. When the consular office asks for proof of domicile, the user can present the signed statement from the bank to the authentication provider and ask the authentication provider to sign a new statement vouching for the user’s region of domicile, using information from the bank as a basis for the statement, without providing any information related to the bank to the consular office.

## 2 Human Rights Constraints

We reflected on the various identity systems used today, including but not limited to residence permits, bank accounts, payment cards, transit passes, and online platform logins, and we found a plethora of features with weaknesses and vulnerabilities concerning privacy (and in some cases security) that could potentially infringe upon human rights. With this in mind, we identified the following eight fundamental constraints to frame the human rights requirements for technology infrastructure [35]:

1. *Minimise **control points** that can be used to co-opt the system.* A single point of trust is a single point of failure, and both state actors and technology firms have historically been proven to abuse such trust.
2. *Mitigate architectural characteristics that lead to **surveillance**.* Surveillance is about control as much as it is about discovery: people behave differently when they believe that their activities are being monitored or evaluated [36]. Powerful actors sometimes employ monitoring to create incentives for individual persons, for example to conduct marketing promotions or credit scoring operations. Such incentives may prevent individuals from acting autonomously, and the chance to discover irregularities, patterns, or even misbehaviour often does not justify such mechanisms of control.
3. *Do not impose **non-consensual** trust relationships upon beneficiaries.* It is an act of coercion for a service provider to require a client to maintain a direct trust relationship with a specific third-party platform provider or certification authority. Infrastructure providers must not explicitly or



**Figure 1:** A schematic representation of a generalised identity system. Users first establish a credential with the system, then use the system to verify the credential, and then use the verified identity to assert that they are authorised to receive a service.

implicitly engage in such coercion, which should be recognised for what it is and not tolerated in the name of convenience.

4. *Prevent solution providers from establishing a **monopoly position**.* Some business models are justified by the opportunity to achieve status as monopoly infrastructure. Monopoly infrastructure is problematic not only because it deprives its users of consumer surplus but also because it empowers the operator to dictate the terms by which the infrastructure can be used.
5. *Empower local businesses and cooperatives to establish their own **trust relationships**.* The opportunity to establish trust relationships on their own terms is important both for businesses to compete in a free marketplace and for businesses to act in a manner that reflects the interests of their communities.
6. *Empower service providers to establish their own **business practices and methods**.* Providers of key services must adopt practices that work within the values and context of their communities.
7. *Empower individual users to manage the **linkages** among their activities.* To be truly free and autonomous, individuals must be able to manage the cross sections of their activities, attributes, and transactions that are seen or might be discovered by various institutions, businesses, and state actors.
8. *Resist establishing potentially abusive **processes and practices**, including legal processes, that rely upon centralised infrastructure.* Infrastructure that can be used to abuse and control individual persons is problematic even if those who oversee its establishment are genuinely benign. Once infrastructure is created, it may in the future be used for other purposes that benefit its operators.

Experience shows that control points will eventually be co-opted by powerful parties, irrespective of the intentions of those who build, own, or operate the control points. Consider, for example, how Cambridge Analytica allegedly abused the data assets of Facebook Inc to manipulate voters in Britain and the US [37] and how the Russian government asserted its influence on global businesses that engaged in domain-fronting [38, 39]. The inherent risk that centrally aggregated datasets may be abused, not only by the parties doing the aggregating but also by third parties, implies value in system design that avoids control points and trusted infrastructure operators, particularly when personal data and livelihoods are involved.

### 3 A Digital Identity System

Various digital identity architectures and deployments exist today to perform the three distinct functions we mentioned earlier: identification, authentication, and authorisation [40]. We characterise the three functions:

- **IDENTIFICATION.** A user first *establishes* some kind of credential or identifier. The credential might be a simple registration, for example with an authority, institution, or other organisation. In other cases, it might imply a particular attribute. The implication might be implicit, as a passport might imply citizenship of a particular country or a credential issued by a bank might imply a banking

request	a request for a credential.
request $x$	a request for a credential with a parameter, $x$ .
$A(m)$	A message $m$ signed by party A.
identify	the foundational identifying elements that a user presents to a certification provider, encrypted so that other parties (including AP) cannot read them.
revoke-one	a message invalidating an earlier signature on a specific user credential.
revoke-all	a message invalidating all signatures by a certain key.
$[m]$	blinded version of message $m$ .
$A([m])$	blind signature of $m$ by A.
prove-owner $x^*$	proof of ownership of some private key $x^*$ , for example via a challenge-response (which would imply two extra messages not shown) or by using the key to sign a pre-existing secret created by the recipient and shared with the sender.
request-certs A	a request for all of the certificates on the ledger signed by A, followed by a response containing all of the matching certificates.
receipt	response from the distributed ledger system indicating that a transaction completed successfully.
object	physical, tamper-resistant object containing a unique receipt for a transaction.

**Table 1:** Notation used in the subsequent figures.

relationship, or it might be explicit, as in the case of attribute-backed credentials such as those issued by Identity Mixer [41, 42].<sup>1</sup>

- AUTHENTICATION. Next, when the provider of a service seeks to authenticate a user, the user must be able to *verify* that a credential in question is valid.
- AUTHORISATION. Finally, the user can use the authenticated credential to *assert* to the service provider that she is entitled to a particular service.


Figure 1 gives a pictorial representation of the functions. Table 1 defines the notation that we shall use in our figures. With the constraints enumerated in Section 2 taken as design requirements, we propose a generalised architecture that achieves our objectives for an identity system. The candidate systems identified in Section 1 can be evaluated by comparing their features to our architecture. Since we intend to argue for a practical solution, we start with a system currently enjoying widespread deployment.

### 3.1 SecureKey Concierge

As a baseline example of an identity framework, we consider a system that uses banks as certification providers whilst circumventing the global payment networks. *SecureKey Concierge* (SKC) [43] is a solution used by the government of Canada to provide users with access to its various systems in a standard way. The SKC architecture seeks the following benefits:

1. Leverage existing “certification providers” such as banks and other financial institutions with well-established, institutional procedures for ascertaining the identities of their customers. Often such procedures are buttressed by legal frameworks such as Anti-Money Laundering (AML) and “Know Your Customer” (KYC) regulations that broadly deputise banks and substantially all financial institutions [46] to collect identifying information on the various parties that make use of their services, establish the expected pattern for the transactions that will take place over time, and monitor the transactions for anomalous activity inconsistent with the expectations [47].
2. Isolate service providers from personally identifying bank details and eliminate the need to share specific service-related details with the certification provider, whilst avoiding traditional authentication service providers such as payment networks.

<sup>1</sup>We do not describe how to use attribute-backed credentials here.



**Figure 2: A stylised schematic representation of the SecureKey Concierge (SKC) system.** The parties are represented by the symbols “CP”, “AP”, “User”, and “Service”, and the arrows represent messages between the parties. The numbers associated with each arrow show the sequence, and the symbol following the number represents the contents of a message. First, the service provider (Service) requests authorisation from the user, who in turn sends identifying information to the authentication provider (AP) to share with the certification provider (CP). If the CP accepts the identifying information, it sends a signed credential  $u$  to the AP, which in turn issues a new credential  $u'$  for consumption by the Service, which can now authorise the user.

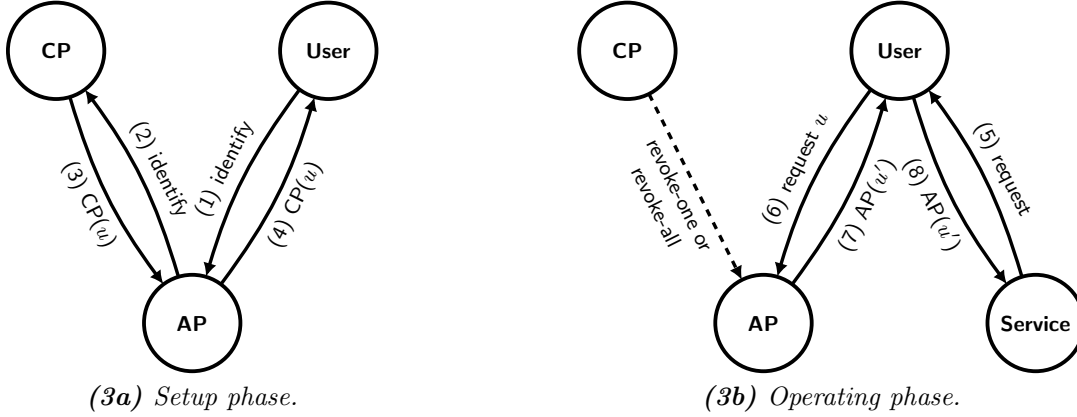
Figure 2 offers a stylised representation of the SKC architecture, as interpreted from its online documentation [43]. When a user wants to access a service, the service provider sends a request to the user (1)<sup>2</sup> for credentials. The user then sends encrypted identifying information (for example, bank account login details) to the authentication provider (2), which in this case is SKC, which then forwards it to the certification provider (3). Next, the certification provider responds affirmatively with a “meaningless but unique” identifier  $u$  representing the user, and sends it to the authentication provider (4). The authentication provider then responds by signing its own identifier  $u'$  representing the user and sending the message to the user (5), which in turn passes it along to the service provider (6). At this point the service provider can accept the user’s credentials as valid. The SKC documentation indicates that SKC uses different, unlinked values of  $u'$  for each service provider.

### 3.2 A Two-Phase Approach

We might consider modifying the SKC architecture so that the user does not need to log in to the CP each time it requests a service. To achieve this, we divide the protocol into two phases, as shown in Figure 3: a *setup phase* (Figure 3a) in which a user establishes credentials with an “certification provider” (CP) for use with the service, and an *operating phase* (Figure 3b) in which a user uses the credentials in an authentication process with a service provider. So, the setup phase is done once, and the operating phase is done once per service request. In the setup phase, the user first sends authentication credentials, such as those used to withdraw money from a bank account, to an authentication provider (1). The authentication provider then uses the credentials to authenticate to the certification provider (2), which generates a unique identifier  $u$  that can be used for subsequent interactions with service providers and sends it to the authentication provider (3), which forwards it to the user (4). Then, in the operating phase, a service provider requests credentials from the user (5), which in turn uses the previously established unique identifier  $u$  to request credentials from the authentication provider (6). This means that the user would implicitly maintain a relationship with the authentication provider, including a way to log in. The authentication provider then verifies that the credentials have not been revoked by the certification provider. The process for verifying that the CP credential is still valid may be offline, via a periodic check, or online, either requiring the AP to reach out to the AP when it intends to revoke a credential or requiring the AP to send a request to the CP in real-time. In the latter case, the AP is looking only for updates on the set of users who have been through the setup phase, and it does not need to identify which user has made a request. Once the AP is satisfied, it sends a signed certification of its identifier  $u'$  to the user (7), which forwards it to the service provider as before (8).

Unfortunately, even if we can avoid the need for users to log in to the CP every time they want to use a service, the authentication provider itself serves as a trusted third party. Although the SKC architecture may eliminate the need to trust the existing payment networks, the authentication provider maintains the mapping between service providers and the certification providers used by individuals to request services. It is also implicitly trusted to manage all of the certification tokens, and there is no way to ensure that it does not choose them in a way that discloses information to service providers or

<sup>2</sup>The numbers in italics correspond to messages in the figure indicated, in this case Figure 2.



**Figure 3:** A schematic representation of a modified version of the SKC system with a stateful authentication provider and one-time identifiers for services. The user first establishes credentials in the setup phase (3a). Then, when a service provider requests credentials from the user in the operating phase (3b), the user reaches out to the authentication provider for verification, which assigns a different identifier  $u'$  each time.

certification providers. In particular, users need to trust the authentication provider to use identifiers that do not allow service providers to correlate their activities, and users may also want to use different identifiers from time to time to communicate with the *same* service provider. As a monopoly platform, it also has the ability to tax or deny service to certification providers, users, or service providers according to its own interests, and it serves as a single point of control vulnerable to exploitation. For all of these reasons, we maintain that the SKC architecture remains problematic from a public interest perspective.

### 3.3 A User-Oriented Identity Architecture

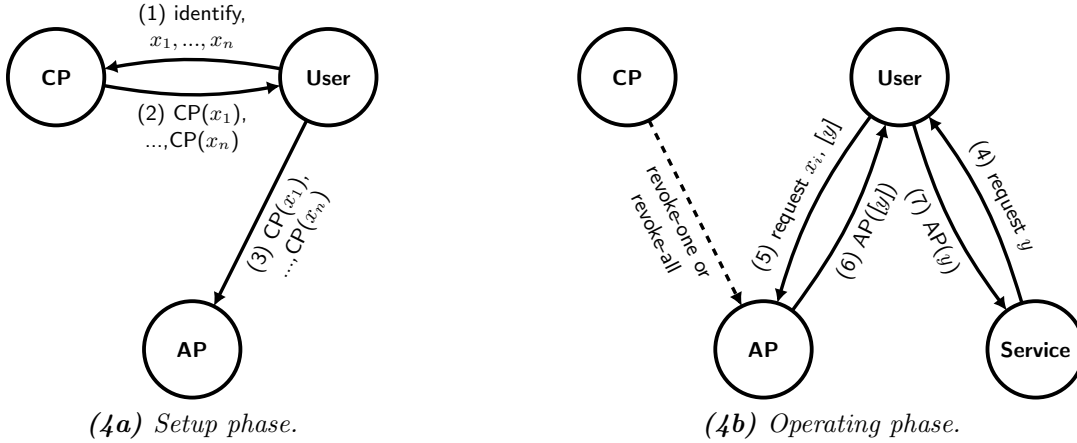
In the architecture presented in Section 3.2, the authentication provider occupies a position of control. In networked systems, control points confer economic advantages on those who occupy them [44], and the business incentives associated with the opportunity to build platform businesses around control points have been used to justify their continued proliferation [45].

However, control points also expose consumers to risk, not only because the occupier of the control point may abuse its position but also because the control point itself creates a vector for attack by third parties. For both of these reasons, we seek to prevent an authentication provider from holding too much information about users. In particular, we do not want an authentication provider to maintain a mapping between a user and the particular services that a user requests, and we do not want a single authentication provider to establish a monopoly position in which it can dictate the terms by which users and service providers interact. For this reason, we put the user, and not the authentication provider, in the centre of the architecture.

For a user to be certain that she is not providing a channel by which authentication providers can leak her identity or by which service providers can trace her activity, then she must *isolate* the different participants in the system. The constraints allow us to define three *isolation objectives* as follows:

1. *Have users generate unlinked identifiers on hardware that they own and trust.* Unless they generate the identifiers themselves, users have no way of knowing for sure whether identifiers assigned to them do not contain personally identifying information. For users to verify that the identifiers will not disclose information that might identify them later, they would need to generate random identifiers using devices and software that they control and trust.
2. *Ensure that authentication providers do not learn about the user's identity or use of services.* Authentication providers that require foundational information about a user, or are able to associate





**Figure 4: A schematic representation of a digital identity system with a user-oriented approach.** The new protocol uses user-generated identifiers and blind signatures to isolate the authentication provider. The authentication provider cannot inject identifying information into the identifiers, nor can it associate the user with the services that she requests.

different requests over time with the same user, are in a position to collect information beyond what is strictly needed for the purpose of the operation. The role of the authentication provider is to act as a neutral channel that confers authority on certification providers, time-shifts requests for credentials, and separates the certification providers from providers of services. Performing this function does not require it to collect information about individual users at any point.

3. *Ensure that information given to service providers is not shared with authentication providers..* The user must be able to credibly trust that his or her interaction with service providers must remain private.

The communication among the four parties that we propose can be done via simple, synchronous (e.g. HTTP) protocols that are easily performed by smartphones and other mobile devices. The cryptography handling public keys can be done using standard public-key-based technologies.

### 3.4 Repositioning the User to be in the Centre

Figure 4 shows how we can modify the system shown in Figure 3 to achieve the three isolation objectives defined above.

Figure 4a depicts the new setup phase. First, on her own trusted hardware, the user generates her own set of identifiers  $x_1, \dots, x_n$  that she intends to use, at most once each, in future correspondence with the authentication provider. Generating the identifiers is not computationally demanding and can be done with an ordinary smartphone. By generating her own identifiers, the user has better control that nothing encoded in the identifiers that might reduce her anonymity. The user then sends both its identifying information and the identifiers  $x_1, \dots, x_n$  to the certification provider (1). The certification provider then responds with a set of signatures corresponding to each of the identifiers (2). The user then sends the set of signatures to the authentication provider for future use (3).

Figure 4b depicts the new operating phase. First, the service sends a request to the user along with a new *nonce* (one-time identifier)  $y$  corresponding to the request (4). The user then applies a blinding function to the nonce  $y$ , creating a blinded nonce  $[y]$ . The user chooses one of the identifiers  $x_i$  that she had generated during the setup phase and sends that identifier along with the blinded nonce  $[y]$  to the authentication provider (5). Provided that the signature on  $x_i$  has not been revoked, the authentication provider confirms that it is valid by signing  $[y]$  and sending the signature to the user (6). The user in turn “unblinds” the signature on  $y$  and sends the unblinded signature to the service provider (7). The use of

blind signatures ensures that the authentication provider cannot link what it sees to specific interactions between the user and the service provider.

### 3.5 Architectural Considerations


To satisfy the constraints listed in Section 2, all three process steps (identification, authentication, and authorisation) must be isolated from each other. Although our proposed architecture introduces additional interaction and computation, we assert that the complexity of the proposed architecture is both parsimonious and justified:

1. If the certification provider were the same as the service provider, then the user would be subject to direct control and surveillance by that organisation, violating Constraints 1, 2, and 7.
2. If the authentication provider were the same as the certification provider, then the user would have no choice but to return to the same organisation each time it requests a service, violating Constraints 1 and 3. That organisation would then be positioned to discern patterns in its activity, violating Constraints 2 and 7. There would be no separate authentication provider to face competition for its services as distinct from the certification services, violating Constraint 4.
3. If the authentication provider were the same as the service provider, then the service provider would be positioned to compel the user to use a particular certification provider, violating Constraints 1 and 3. The service provider could also impose constraints upon what a certification provider might reveal about an individual, violating Constraint 2, or how the certification provider establishes the identity of individuals, violating Constraint 6.
4. If the user could not generate her own identifiers, then the certification provider could generate identifiers that reveal information about the user, violating Constraint 2.
5. If the user were not to use blind signatures to protect the requests from service providers, then service providers and authentication providers could compare notes to discern patterns of a user's activity, violating Constraint 7.

The proposed architecture does not achieve its objectives if either the certification provider or the service provider colludes with the authentication provider; we assume that effective institutional policy will complement appropriate technology to ensure that sensitive data are not shared in a manner that would compromise the interests of the user.

## 4 A Decentralised Identity Architecture

A significant problem remains with the design described in Section 3.4 in that it requires  $O(n^2)$  relationships among authentication providers and certification providers (i.e., with each authentication provider connected directly to each certification provider that it considers valid) to be truly decentralised. Recall that the system relies critically upon the ability of a certification provider to *revoke* credentials issued to users, and authentication providers need a way to learn from the certification provider whether a credential in question has been revoked. Online registries such as OCSP [48], which are operated by a certification provider or trusted authority, are a common way to address this problem, although the need for third-party trust violates Constraint 1. The issue associated with requiring each authentication provider to establish its own judgment of each candidate certification provider is a *business* problem rather than a technical one. Hierarchical trust relationships emerge because relationships are expensive to maintain and introduce risks; all else being equal, business owners prefer to have fewer of them. Considered in this context, concentration and lack of competition among authentication providers makes sense. If one or a small number of authentication providers have already established relationships with a broad set of certification providers, just as payment networks such as Visa and Mastercard have done with a broad set of banks, then the cost to a certification provider of a relationship with a new authentication provider would become a barrier of entry to new authentication providers. The market for authentication could fall under the control of a monopoly or cartel.



**Figure 5:** *A schematic representation of a decentralised identity system with a distributed ledger. The user is not required to interact directly with the distributed ledger (represented by a dashed circle) and can rely upon the services offered by certification providers and authentication providers.*

#### 4.1 Introducing Distributed Ledger Technology

We propose using *distributed ledger technology* (DLT) to allow both certification providers and authentication providers to proliferate whilst avoiding industry concentration. The distributed ledger would serve as a standard way for certification providers to establish relationships with any or all of the authentication providers at once, or vice-versa. The ledger itself would be a mechanism for distributing signatures and revocations; it would be shared by participants and not controlled by any single party. Figure 5 shows that users would not interact with the distributed ledger directly but via their chosen certification providers and authentication providers. Additionally, users would not be bound to use any particular authentication provider when verifying a particular credential and could even use a different authentication provider each time. Provided that the community of participants in the distributed ledger remains sufficiently diverse, the locus of control would not be concentrated within any particular group or context, and the market for authentication can remain competitive.

Because the distributed ledger architecture inherently does not require each new certification provider to establish relationships with all relevant authentication providers, or vice-versa, it facilitates the entry of new authentication providers and certification providers, thus allowing the possibility of decentralisation.

We argue that a distributed ledger is an appropriate technology to maintain the authoritative record of which credentials have been issued (or revoked) and which transactions have taken place. We do not trust any specific third party to manage the list of official records, and we would need the system to be robust in the event that a substantial fraction of its constituent parts are compromised. The distributed ledger can potentially take many forms, including but not limited to blockchain, and, although a variety of fault-tolerant consensus algorithms may be appropriate, we assume that the set of node operators is well-known, a characteristic that we believe may be needed to ensure appropriate governance.

If implemented correctly at the system level, the use of a distributed ledger can ensure that the communication between the certification provider and the authentication provider is limited to that which is written on the ledger. If all blind signatures are done without including any accompanying metadata, and as long as the individual user does not reveal which blind signature on the ledger corresponds to the unblinded signature that he or she is presenting to the authentication provider for approval, then nothing on the ledger will reveal any information about the individual persons who are the subjects of the certificates. We assume that the certification authorities would have a limited and well-known set of public keys that they would use to sign credentials, with each key corresponding to the category of individual persons who have a specific attribute. The size of the anonymity set for a credential, and therefore the effectiveness of the system in protecting the privacy of an individual user with that

credential, depends upon the generality of the category. We would encourage certification authorities to assign categories that are as large as possible. We would also assume that the official set of signing keys used by certification providers and authentication providers is also maintained on the ledger, as doing so would ensure that all users of the system have the same view of which keys the various certification providers and authentication providers are using.

## 4.2 Achieving Decentralisation with a Distributed Ledger

Figure 6 shows how the modified architecture with the distributed ledger technology would work. Figure 6a shows the setup phase. The first two messages from the user to the certification provider are similar to their counterparts in the protocol shown in Figure 4a. However, now the user also generates  $n$  asymmetric key pairs  $(x_i, x_i^*)$ , where  $x_i$  is the public key and  $x_i^*$  is the private key of pair  $i$ , and it sends each public key  $x_1, \dots, x_n$  to the certification provider (1). Then, rather than sending the signed messages to the authentication provider via the user, the certification provider then instead writes the signed certificates directly to the distributed ledger (2). Importantly, the certificates would not contain any metadata but only the public key  $x_i$  and its bare signature; eliminating metadata is necessary to ensure that there is no channel by which a certification provider might inject information that might be later used to identify a user. Figure 6b shows the operating phase, which begins when a service provider asks a user to authenticate and provides some nonce  $y$  as part of the request.

The certification provider can revoke certificates simply by transacting on the distributed ledger and without interacting with the authentication provider at all. Because the user and the authentication provider are no longer assumed to mutually trust one another, the user must now prove to the authentication provider that the user holds the private key  $x_i^*$  when the user asks the authentication provider to sign the blinded nonce  $[y]$  (4). At this point we assume that the authentication provider maintains its own copy of the distributed ledger and has been receiving updates. The authentication provider then refers to its copy of the distributed ledger to determine whether a credential has been revoked, either because the certification provider revoked a single credential or because the certification provider revoked its own signing key. Provided that the credential has not been revoked, the authentication provider signs the blinded nonce  $[y]$  (5), which the user then unblinds and sends to the service provider (6). The following messages are carried out as they are done in Figure 4b.

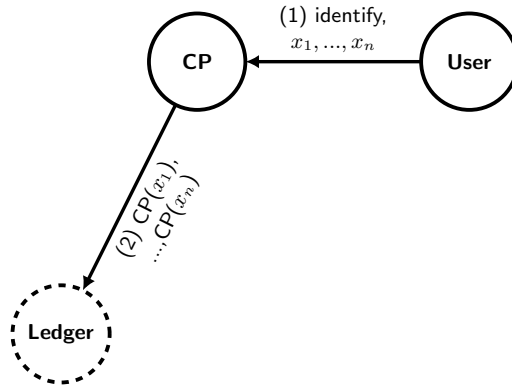
We assume that each certification provider and authentication provider has a distinct signing key for credentials representing each possible policy attribute, and we further assume that each possible policy attribute admits for a sufficiently large anonymity set to not identify the user, as described in Section 4.1. A policy might consist of the union of a set of attributes, and because users could prevent arbitrary subsets of the attributes to authentication providers and service providers, we believe that in most cases it would not be practical to structure policy attributes in such a manner that one attribute represents a qualification or restriction of another. Additionally, at a system level, authentication providers and service providers must not require a set of attributes, either from the same issuer or from different issuers, whose combination would restrict the anonymity set to an extent that would potentially reveal the identity of the user.

## 4.3 Operating the System Offline

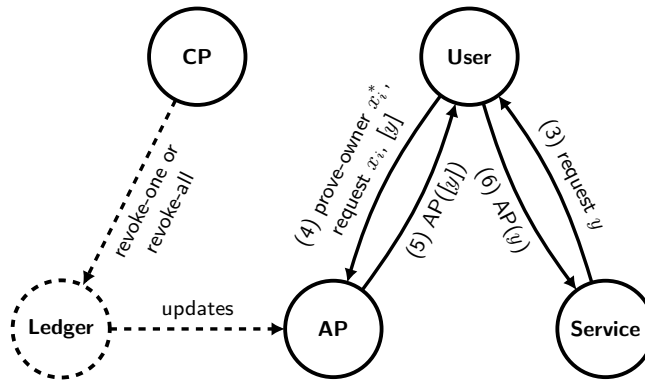
The proposed approach can also be adapted to work *offline*, specifically when a user does not have access to an Internet-connected authentication provider at the time that it requests a service from a service provider. This situation applies to two cases: first, in which the authentication provider has only intermittent access to its distributed ledger peers (perhaps because the authentication provider has only intermittent access to the Internet), and second, in which the user does not have access to the authentication provider (perhaps because the user does not have access to the Internet) at the time that it requests a service.

In the first case, note that the use of a distributed ledger allows the authentication provider to avoid the need to send a query in real-time.<sup>3</sup> If the authentication provider is disconnected from the

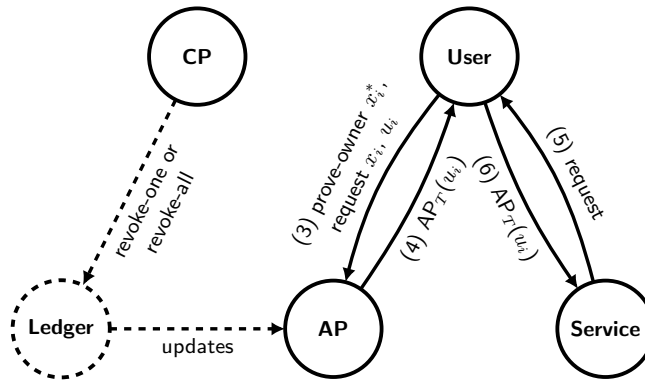
<sup>3</sup>Not sending the query over the network may also improve the privacy of the transaction.



(6a) Setup phase.



(6b) "Online mode" operating phase.



(6c) "Offline mode" operating phase.

**Figure 6: A schematic representation of one possible decentralised digital identity system using a distributed ledger.** Diagrams (6a) and (6b) show the setup and operating phases for an initial sketch of our design, which uses a distributed ledger to promote a scalable marketplace that allows users to choose certification providers and authentication providers that suit their needs. Diagram (6c) shows a variation of the operating phase that can be used in an offline context, in which the user might not be able to communicate with an up-to-date authentication provider and the service provider at the same time.

network, then it can use its most recent version of the distributed ledger to check for revocation. If the authentication provider is satisfied that the record is sufficiently recent, then it can sign the record with a key that is frequently rotated to indicate its timeliness, which we shall denote by  $AP_T$ . We presume that  $AP_T$  is irrevocable but valid for a limited time only. If the authentication provider is disconnected from its distributed ledger peers but still connected to the network with the service provider, then it can still sign a nonce from the service provider as usual.

In the second case, however, although the user is disconnected from the network, the service provider still requires an indication of the timeliness of the authentication provider’s signature. The generalised solution is to adapt the operating phase of the protocol as illustrated by Figure 6c. Here, we assume that the user knows in advance that she intends to request a service at some point in the near future, so she sends the request to the authentication provider pre-emptively, along with a one-time identifier  $u_i$  (3). Then, the authentication provider verifies the identifier via the ledger and signs the one-time identifier  $u_i$  with the time-specific key  $AP_T$  (4). Later, when the service provider requests authorisation (5), the user responds with the signed one-time identifier that it had obtained from the authentication provider (6). In this protocol, the service provider also has a new responsibility, which is to keep track of one-time identifiers to ensure that there is no duplication.


#### 4.4 Achieving Unlinkability with Blinded Credentials

Unfortunately, the architecture described in Sections 4.2 and 4.3 has an important weakness as a result of its reliance on the revocation of user credentials. Because the credential that an certification provider posts to the ledger is specifically identified by the user at the time that the user asks the authentication provider to verify the credential, the certification provider may collude with individual authentication providers to determine when a user makes such requests. Even within the context of the protocol, an unscrupulous (or compromised, or coerced) certification provider may post revocation messages for all of the credentials associated with a particular user, hence linking them to each other.

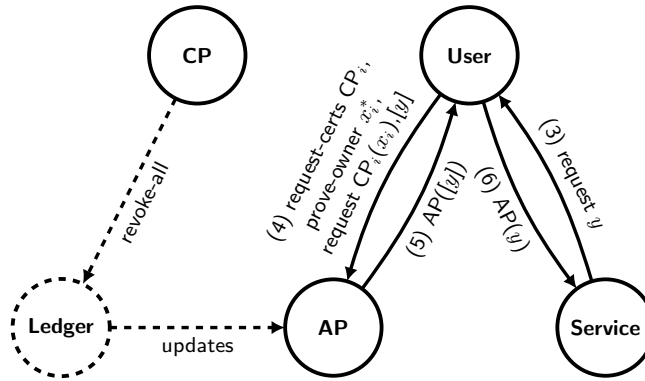
For this reason, we recommend modifying to the protocol to prevent this attack by using blinded credentials to improve its metadata-resistance. Figure 7 shows how this would work. Rather than sending the public keys  $x_i$  directly to the certification provider, the user sends blinded public keys  $[x_i]$ , one for each of a series of specific, agreed-upon time intervals (1). which in turn would be signed by the certification provider using a blind signature scheme that does not allow revocation (1). The certification provider would not sign all of the public keys and publish the certificates to the ledger immediately; instead, it would sign them and post the certificates to the ledger at the start of each time interval  $i$ , in each instance signing the user keys with a key of its own specific to that time interval,  $CP_i$  (2). If a user expects to make multiple transactions per time interval and desires those transactions to remain unlinked from each other, the user may send multiple keys for each interval.

When the time comes for the user to request a service, the user must demonstrate that it is the owner of the private key corresponding to a (blinded) public key that had been signed by the certification provider. So the user must first obtain the set of all certificates signed by  $CP_i$ , which it can obtain from the authentication provider via a specific request, `request-certs`. Then it can find the blind signature on  $[x_i]$  from the list and unblind the signature to reveal  $CP_i(x_i)$ . It can then send this signature to the authentication provider along with its proof of ownership of  $x_i^*$  as before.

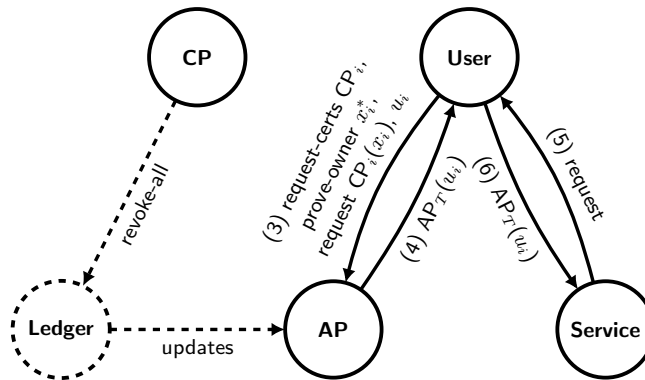
This version of the protocol is the one that we recommend for most purposes. Although the `request-certs` exchange might require the user to download a potentially large number of certificates, such a requirement would hopefully indicate a large anonymity set. In addition, there may be ways to mitigate the burden associated by the volume of certificates loaded by the client. For example, we might assume that the service provider offers a high-bandwidth internet connection that allows the user to request the certificates anonymously from an authentication provider. Alternatively, we might consider having the certification provider subdivide the anonymity set into smaller sets using multiple well-known public keys rather than a single  $CP_i$ , or we might consider allowing an interactive protocol between the user and the authentication provider in which the user voluntarily opts to reduce her anonymity set, for example by specifying a small fraction of the bits in  $[x_i]$  as a way to request only a subset of the certificates.



(7a) Setup phase.



(7b) "Online mode" operating phase.



(7c) "Offline mode" operating phase.

**Figure 7: A schematic representation of a metadata-resistant decentralised identity architecture.** This version of the design represents our recommendation for a generalised identity architecture. By writing only blinded credentials to the ledger, this version extends the design shown in Figure 6 to resist an attack in which the certification provider can expose linkages between different credentials associated with the same user. Diagrams (7a) and (7b) show the setup and operating phases, analogous to the online example shown in Figure 6; Diagram (7c) shows the corresponding offline variant.

## 4.5 Adapting the Design for Spending Tokens

The architecture defined in Section 4.4 can also be adapted to allow users to spend tokens on a one-time basis. This option may be of particular interest for social and humanitarian services, in which tokens to be used to purchase food, medicine, or essential services may be issued by a government authority or aid organisation to a community at large. In such cases, the human rights constraints are particularly important. Figure 8 shows how a certification provider might work with an issuer of tokens to issue spendable tokens to users, who may in turn represent themselves or organisations.

Figure 8a shows the setup phase. We assume that the service provider first tells the user that it will accept tokens issued by the certification provider (1). The user then sends a set of  $n$  newly-generated public keys, along with any required identity or credential information, to the certification provider (2). We assume that the tokens are intended to be fungible, so the certification provider issues  $n$  new, fungible tokens on the ledger (3). We need not specify the details of how the second message works in practice; depending upon the trust model for the ledger it may be as simple as signing a statement incrementing the value associated with the certification provider’s account, or it may be a request to move tokens explicitly from one account to another. Then, the certification provider signs a set of  $n$  messages, each containing one of the blinded public keys, and sends them to the user (4). The messages will function as promissory notes, redeemable by the user that generated the keys, for control over the tokens.

Figure 8b shows the operating phase. When a service provider requests a token (5), the user sends a message to an authentication provider demonstrating both that it has the right to control the token issued by the certification provider and that it wishes to sign the token over to the service provider (6). The authentication provider, who never learns identifying information about the user, lodges a transaction on the ledger that assigns the rights associated with the token to the service provider (7), generating a receipt (8). Once the transaction is complete, the authentication provider shares a receipt with the user (9), which the user may then share with the service provider (10), who may now accept that the payment as complete.


Like the “main” architecture described in Section 4.4, the “token” architecture can also be configured to work in an offline context by modifying the operating phase. Figure 8c shows how this would work. The user requests from the authentication provider one or more physical “objects”, which may take the form of non-transferrable electronic receipts or physical tokens, that can be redeemed for services from the service provider (5). The authentication provider sends the objects to the user (8), who then redeems them with the service provider in a future interaction (9,10).

## 5 Governance Considerations

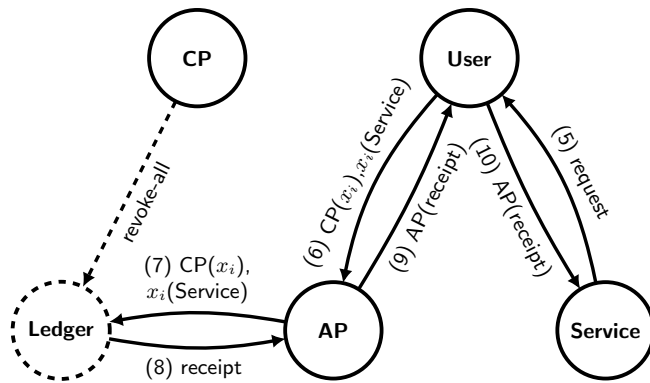
An important challenge that remains with the distributed ledger system described in Section 4 is the management of the organisations that participate in the consensus mechanism of the distributed ledger. We believe that this will require the careful coordination of local businesses and cooperatives to ensure that the system itself does not impose any non-consensual trust relationships (Constraint 3), that no single market participant would gain dominance (Constraint 4), and that participating businesses and cooperatives will be able to continue to establish their own business practices and trust relationships on their own terms (Constraint 5), even while consenting to the decisions of the community of organisations participating in the shared ledger. We believe that our approach will be enhanced by the establishment of a multi-stakeholder process to develop the protocols for participating in the distributed ledger and ultimately facilitate a multiplicity of different implementations of the technology needed to participate. Industry groups and regulators will still need to address the important questions of determining the rules and participants. We surmise that various organisations, ranging from consulting firms to aid organisations, would be positioned to offer guidance to community participants in the network, without imposing new constraints.

The case for a common, industry-wide platform for exchanging critical data via a distributed ledger is strong. Analogous mechanisms have been successfully deployed in the financial industry. Established mechanisms take the form of centrally-managed cooperative platforms such as SWIFT [49], which securely

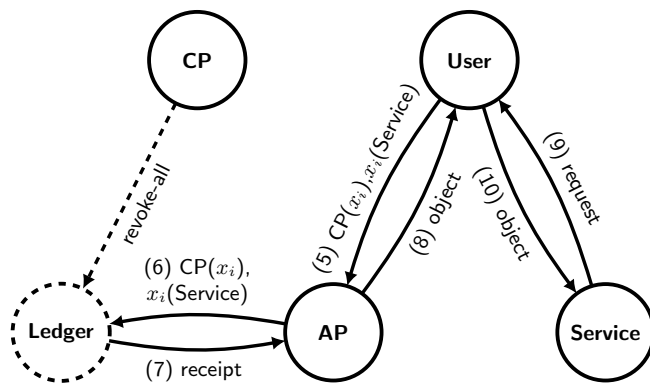




(8a) Setup phase.



(8b) "Online mode" operating phase.



(8c) "Offline mode" operating phase.

**Figure 8: A schematic representation of a decentralised identity architecture for exchanging tokens.** The protocol represented in this figure can be used to allow users to pay for services privately. Diagrams (8a) and (8b) show the setup and operating phases, analogous to the online example shown in Figures 6 and 7; Diagram (8c) shows the corresponding offline variant.

carries messages on behalf of financial markets participants, while others take the form of consensus-based industry standards, such as the Electronic Data Interchange (EDI) standards promulgated by X12 [50] and EDIFACT [51]. Distributed ledgers such as Ripple [52] and Hyperledger [53] have been proposed to complement or replace the existing mechanisms.

## 5.1 Open Standards

For the digital identity infrastructure, we suggest that the most appropriate application for the distributed ledger system described in Section 4 would be a technical standard for business transactions promulgated by a self-regulatory organisation working concordantly with government regulators. A prime example from the financial industry is Regulation NMS [54], which led to the dismantling of a structural monopoly in electronic equities trading in the United States. Although the US Securities and Exchange Commission had the authority to compel exchanges to participate in a national market system since 1975, it was not until thirty years later that the SEC moved to explicitly address the monopoly enjoyed by the New York Stock Exchange (NYSE). The Order Protection Rule imposed by the 2005 regulation (Rule 611) “requir[ed] market participants to honor the best prices displayed in the national market system by automated trading centers, thus establishing a critical linkage framework” [55]. The monopoly was broken, NYSE member firms became less profitable, and NYSE was ultimately bought by Intercontinental Exchange in 2013 [56].

We believe that distributed ledgers offer a useful mechanism by which self-regulatory organisations satisfy regulations precisely intended to prevent the emergence of control points, market concentration, or systems whose design reflects a conflict of interest between their operators and their users.

## 5.2 No Master Keys, No Early Binding

An important expectation implicit to the design of our system is that users can establish and use as many identities as they want, without restriction. This means not only that a user can choose which credentials to show to relying parties, but also that a user would not be expected to bind the credentials to each other in any way prior to their use. Such a binding would violate Constraint 7 as given in Section 2. In particular, given two credentials, there should be no way to know or prove that they were issued to the same individual or device. This property is not shared by some schemes for non-transferrable anonymous credentials that encourage users to bind together credentials to each other via a master key or similar mechanism [33] as described in Section 1.2.

If it were possible to prove that two or more credentials were associated with the same identity, then an individual could be forced to associate a set of credentials with each other inextricably, and even if an individual might be given an option to reveal only a subset of his or her credentials to a service provider at any given time, the possibility remains that an individual might be compelled to reveal the linkages between two or more credentials. For example, the device that an individual uses might be compromised, revealing the master key directly, which would be problematic if the same master key were used for many or all of the individual’s credentials. Alternatively, the individual might be coerced to prove that the same master key had been associated with two or more credentials.

The system we describe explicitly does not seek to rely upon the *ex ante* binding together of credentials to achieve non-transferability or for any other purpose. We suggest that the specific desiderata and requirements for non-transferability might vary across use cases and can be addressed accordingly. Exogenous approaches to achieve non-transferability might have authentication providers require users to store credentials using trusted escrow services or physical hardware with strong counterfeiting resistance such as two-factor authentication devices. Endogenous approaches might have authentication providers record the unblinded signatures on the ledger once they are presented for inspection, such that multiple uses of the same credential become explicitly bound to each other *ex post* or are disallowed entirely. Recall that the system assumes that credentials are used once only and that certification providers would generate new credentials for individuals in limited batches, for example at a certain rate over time.

	assert entitlements	spend tokens
private sector	membership programmes	mobile communication services
public sector	library access programme eligibility	public transportation

**Table 2:** *A categorisation matrix for use cases.* We divide the universe of use cases into four categories based upon whether the purpose is to assert entitlements or to spend tokens and upon whether the services in question are operated by the public sector or the private sector, and we include some examples.

## 6 Use Cases

We anticipate that there might be many potential use cases for a decentralised digital identity infrastructure that affords users the ability to manage the linkages among their credentials. Table 2 presents one view of how the use cases might be divided into four categories on the basis of whether the purpose is to assert entitlements or spend tokens and upon whether the services in question are operated by the public sector or the private sector.

Use cases that involve asserting entitlements might include asserting membership of a club for the purpose of gaining access to facilities, accessing the resources of a public library, or demonstrating eligibility for a discount, perhaps on the basis of age, disability, or financial hardship, at the point of sale. Use cases that involve spending tokens can potentially be disruptive, particularly in areas that generate personally identifiable information. Tokens for public transportation (for example, pay-as-you-go or monthly bus tickets) can be purchased with cash in one instance and then spent over time with each trip unlinkable to each of the others. Mobile phone service subscribers can have their devices present tokens to cellular towers without revealing their specific identities, allowing them to avoid tracking.

## 7 Conclusions and Future Work

We argue that the ability of individuals to create and maintain multiple unrelated identities is a fundamental, inalienable human right. For the digital economy to function while supporting this human right, individuals must be able to control and limit what others can learn about them over the course of their many interactions with services, including government and institutional services.

We have introduced an open digital identity architecture that promotes the implementation of identity architectures that satisfy constraints that we consider essential to the protection of human rights, and we believe that a combination of strong technology and thoughtful policy will be necessary to promote and ensure the implementation, deployment, and use of technology that satisfies them. We have elaborated eight requirements for technology infrastructure and demonstrated that they can be achieved by means

of a decentralised architecture. We discussed issues associated with scalability and governance, also demonstrating how tokens can be spent via such a system.

The specific mechanism for fostering a community of participating organisations will depend upon the relationship between those organisations and the group that ultimately assumes the role of ensuring that the system does not impose non-consensual trust relationships on its users. It must be noted that any system that puts control in the hands of end-users carries the burden of education, both for the well-functioning of the system as well as for safeguarding its role in protecting the public interest. Future research, therefore, must include case studies of how similar systems have been developed, deployed, and maintained over time, in a variety of different social and business contexts.

## 8 Acknowledgements

We thank Valerie Khan, Edgar Whitley, and Paul Makin for their thoughtful insights. Geoff Goodell is also an associate of the Centre for Technology and Global Affairs at the University of Oxford. We acknowledge the Engineering and Physical Sciences Research Council (EPSRC) for the BARAC project (EP/P031730/1) and the European Commission for the FinTech project (H2020-ICT-2018-2 825215). Tomaso Aste acknowledges the Economic and Social Research Council (ESRC) for funding the Systemic Risk Centre (ES/K0 02309/1).

## References

- [1] Wilcox-O’Hearn, Z. “Names: Decentralized, Secure, Human-Meaningful: Choose Two.” [online] <https://web.archive.org/web/20011020191610/http://zooko.com/distnames.html> [retrieved 2018-04-21]
- [2] Stiegler, M. “An Introduction to Petname Systems.” February 2005. [online] <http://www.skyhunter.com/marcs/petnames/IntroPetNames.html> [retrieved 2018-05-11]
- [3] Douceur, J. “The Sybil Attack.” *IPTPS ’01 Revised Papers from the First International Workshop on Peer-to-Peer Systems*, 2002-03-07, pp. 251–260. [online] <https://www.freehaven.net/anonbib/cache/sybil.pdf> [retrieved 2018-05-11]
- [4] Vanderburg, E. “A Certified Lack of Confidence: The Threat of Rogue Certificate Authorities.” TCDI Blog. [online] <https://www.tcdi.com/the-threat-of-rogue-certificate-authorities/> [retrieved 2018-05-11]
- [5] Charette, R. “DigiNotar Certificate Authority Breach Crashes e-Government in the Netherlands.” *IEEE Spectrum*, 2011-09-09. [online] <https://spectrum.ieee.org/riskfactor/telecom/security/diginotar-certificate-authority-breach-crashes-egovernment-in-the-netherlands> [retrieved 2018-05-11]
- [6] Kushner, D. “The Real Story of Stuxnet.” *IEEE Spectrum*, 2013-02-26. [online] <https://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet> [retrieved 2018-05-11]
- [7] Eckersley, P. Electronic Frontier Foundation Technical Analysis, 2011-03-23. [online] <https://www.eff.org/deeplinks/2011/03/iranian-hackers-obtain-fraudulent-https> [retrieved 2018-05-11]
- [8] Goodin, D. “Already on probation, Symantec issues more illegit HTTPS certificates.” *Ars Technica*, 2017-01-20. [online] <https://arstechnica.com/information-technology/2017/01/already-on-probation-symantec-issues-more-illegit-https-certificates/> [retrieved 2018-05-11]
- [9] Equifax Inc. “2017 Cybersecurity Incident & Important Consumer Information.” 2018-03-01. [online] <https://www.equifaxsecurity2017.com/> [retrieved 2018-05-11]

- [10] Chirgwin, R. “Google publishes list of Certificate Authorities it doesn’t trust.” *The Register*, 2016-03-23. [online] [https://www.theregister.co.uk/2016/03/23/google\\_now\\_publishing\\_a\\_list\\_of\\_cas\\_it\\_doesnt\\_trust/](https://www.theregister.co.uk/2016/03/23/google_now_publishing_a_list_of_cas_it_doesnt_trust/) [retrieved 2018-05-11]
- [11] Slashdot. “School Tricks Pupils Into Installing a Root CA.” 2014-03-09. [online] <https://news.slashdot.org/story/14/03/09/0225224/school-tricks-pupils-into-installing-a-root-ca> [retrieved 2018-05-11]
- [12] Bright, P. “Govt, certificate authorities conspire to spy on SSL users?” *Ars Technica*, 2010-03-29. [online] <http://files.cloudprivacy.net/ssl-mitm.pdf> [retrieved 2018-05-11]
- [13] Goodin, D. “Flaw crippling millions of crypto keys is worse than first disclosed.” *Ars Technica*. 2017-11-06. [online] <https://arstechnica.com/information-technology/2017/11/flaw-crippling-millions-of-crypto-keys-is-worse-than-first-disclosed/> [retrieved 2018-04-21]
- [14] Moon, M. “Estonia freezes resident ID cards due to security flaw.” *Engadget*. 2017-11-04. [online] <https://www.engadget.com/2017/11/04/estonia-freezes-resident-id-cards-security-flaw/> [retrieved 2018-04-21]
- [15] Thevoz, P. “Diving into a ‘Digital Country’: e-Estonia.” *Medium*. 2016-09-16. [online] <https://medium.com/@PhilippeThevoz/diving-into-a-digital-country-e-estonia-af561925c95e> [retrieved 2018-04-21]
- [16] Tully, M. “The problem with Aadhaar cards is the way they are being pushed by the State.” *Hindustan Times*, 2017-12-09. [online] <https://www.hindustantimes.com/analysis/the-problem-with-aadhaar-cards-is-the-way-they-are-being-pushed-by-the-state/story-RT1WUXgF3ck4rsoN1zKXUI.html> [retrieved 2018-04-21]
- [17] Wagner, R. “Identity and Access Management 2020.” *ISSA Journal*, June 2014. [online] <http://www.issa.org/resource/resmgr/JournalPDFs/feature0614.pdf> [retrieved 2019-01-11]
- [18] Shuhaimi, N and Juhana, T. “Security in vehicular ad-hoc network with Identity-Based Cryptography approach: A survey.” 7th International Conference on Telecommunication Systems, Services, and Applications (TSSA), 2012-10-30. [online] <http://ieeexplore.ieee.org/abstract/document/6366067/> [retrieved 2019-01-11]
- [19] Government Digital Service (UK). “GOV.UK Verify: Guidance.” 2018-12-14. [online] <https://www.gov.uk/government/publications/introducing-govuk-verify/introducing-govuk-verify> [retrieved 2019-02-15]
- [20] Brandao, L, Christin, N, Danezis, G, and Anonymous. “Toward Mending Two Nation-Scale Brokered Identification Systems.” *Proceedings on Privacy Enhancing Technologies* **2015**(2), pp. 135-155. [online] <http://www0.cs.ucl.ac.uk/staff/G.Danezis/papers/popets15-brokid.pdf> [retrieved 2019-01-11]
- [21] Whitley, E. “Trusted Digital Identity Provision: GOV.UK Verify’s Federated Approach.” Center for Global Development Policy Paper 131, 2018-11-07. [online] <https://www.cgdev.org/sites/default/files/Trusted-Digital-ID-Provision-govuk.pdf> [retrieved 2019-02-15]
- [22] O’Hara, K, Whitley, E, and Whittall, P. “Avoiding the Jigsaw Effect: Experiences with Ministry of Justice Reoffending Data.” Monograph, 2011-12-19, modified 2018-06-06. [online] <https://eprints.soton.ac.uk/273072/1/AVOIDING%2520THE%2520JIGSAW%2520EFFECT.pdf> [retrieved 2018-07-26]
- [23] Thomas, T. “Joint Watermarking Scheme for Multiparty Multilevel DRM Architecture.” *IEEE Transactions on Information Forensics and Security* **4**(4), 2009-09-29, pp. 758–767. [online] <http://ieeexplore.ieee.org/abstract/document/5272486/> [retrieved 2019-01-11]
- [24] Liu, S et al. “Method and System for Implementing Authentication on Information Security.” United States Patent Application US20080065895A1, 2008-03-13. [online] <https://patents.google.com/patent/US20080065895A1/en> [retrieved 2019-01-11]

- [25] Thackston, J. “System and method for verifying user identity in a virtual environment.” US Patent Grant US10153901B2, 2018-12-11. [online] <https://patents.google.com/patent/US10153901B2/en> [retrieved 2019-01-11]
- [26] Dunphy, P and Petitcolas, F. “A First Look at Identity Management Schemes on the Blockchain.” arXiv:1801.03294v1 [cs.CR], 2018-01-10. [online] <https://arxiv.org/abs/1801.03294> [retrieved 2019-01-11]
- [27] SITA. “Travel Identity of the Future – White Paper.” ShoCard, May 2016. [online] <https://shocard.com/wp-content/uploads/2016/11/travel-identity-of-the-future.pdf> [retrieved 2019-01-11]
- [28] Tobin, A and Reed, D. “The Inevitable Rise of Self-Sovereign Identity.” The Sovrin Foundation, 2016-09-29. [online] <https://sovrin.org/wp-content/uploads/2017/06/The-Inevitable-Rise-of-Self-Sovereign-Identity.pdf> [retrieved 2019-01-11]
- [29] Aitken, R. “IBM Blockchain Joins Sovrin’s ‘Decentralized’ Digital Identity Network To Stem Fraud.” Forbes, 2018-04-05. [online] <https://www.forbes.com/sites/rogeraitken/2018/04/05/ibm-blockchain-joins-sovrins-decentralized-digital-identity-network-to-stem-fraud/> [retrieved 2019-01-11]
- [30] Lundkvist, C et al. “uPort: A Platform for Self-Sovereign Identity.” Draft Version, 2016-10-20. [online] [http://blockchainlab.com/pdf/uPort\\_whitepaper\\_DRAFT20161020.pdf](http://blockchainlab.com/pdf/uPort_whitepaper_DRAFT20161020.pdf) [retrieved 2019-01-11]
- [31] Kaaniche, N and Laurent, M. “A blockchain-based data usage auditing architecture with enhanced privacy and availability.” IEEE 16th International Symposium on Network Computing and Applications (NCA), 2017-10-30. [online] <https://ieeexplore.ieee.org/abstract/document/8171384/> [retrieved 2019-01-11]
- [32] Sonnino, A et al. “Coconut: Threshold Issuance Selective Disclosure Credentials with Applications to Distributed Ledgers.” arXiv:1802.07344v3 [cs.CR], 2018-08-08. [online] <https://arxiv.org/abs/1802.07344> [retrieved 2019-01-11]
- [33] Camenisch, J and Lysyanskaya, A. “An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation.” Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2001: Advances in Cryptology), 2001-04-15, pp. 93–118.
- [34] Chaum, D. “Blind signatures for untraceable payments.” In *Advances of Cryptology Proceedings of Crypto*. **82**(3), pp. 199–203, 1983. <http://www.hit.bme.hu/~buttyan/courses/BMEVIHIM219/2009/Chaum.BlindSigForPayment.1982.PDF> [retrieved 2018-02-23]
- [35] Goodell, G and Aste, T. “Blockchain technology for the public good: Design constraints in a human rights context.” To appear.
- [36] Mayo, E. “Hawthorne and the Western Electric Company.” In *The Social Problems of an Industrial Civilization*. Boston: Division of Research, Harvard Business School, 1945. [online] [www.practicessurvival.com/wa\\_files/Hawthorne\\_20Studies\\_201924\\_20Elton\\_20Mayo.pdf](http://www.practicessurvival.com/wa_files/Hawthorne_20Studies_201924_20Elton_20Mayo.pdf) [retrieved 2019-01-06]
- [37] Kosłowska, H et al. “The Cambridge Analytica scandal is wildly confusing. This timeline will help.” *Quartz*, 2018-03-29. [online] <https://qz.com/1240039/the-cambridge-analytica-scandal-is-confusing-this-timeline-will-help/> [retrieved 2018-04-20]
- [38] Lunden, I. “Russias Telegram ban that knocked out 15M Google, Amazon IP addresses had a precedent in Zello.” *TechCrunch*, 2018-04-19. [online] <https://techcrunch.com/2018/04/17/russias-telegram-ban-that-knocked-out-15m-google-amazon-ip-addresses-had-a-precedent-in-zello/> [retrieved 2018-04-20]
- [39] Savov, V. “Russias Telegram ban is a big, convoluted mess.” *The Verge*, 2018-04-17. [online] <https://www.theverge.com/2018/4/17/17246150/telegram-russia-ban> [retrieved 2018-04-20]

- [40] Riley, S. “Its Me, and Heres My Proof: Why Identity and Authentication Must Remain Distinct.” Microsoft TechNet Security Viewpoint, 2006-02-14. [online] <https://technet.microsoft.com/en-us/library/cc512578.aspx> [retrieved 2018-05-11]
- [41] IBM Research Zurich. “IBM Identity Mixer.” [online] [https://www.zurich.ibm.com/identity\\_mixer/](https://www.zurich.ibm.com/identity_mixer/) [retrieved 2018-04-21]
- [42] Camenisch, J and Lysyanskaya, A. “A Signature Scheme with Efficient Protocols.” *Lecture Notes in Computer Science*, vol. 2576. Springer, 2003. [online] [http://rd.springer.com/chapter/10.1007%2F3-540-36413-7\\_20](http://rd.springer.com/chapter/10.1007%2F3-540-36413-7_20) [retrieved 2017-05-10]
- [43] SecureKey Technologies, Inc. “Trust Framework – SecureKey Concierge in Canada, SKUN-117.” 2015-09-09. [online] <http://securekey.com/wp-content/uploads/2015/09/SK-UN117-Trust-Framework-SecureKey-Concierge-Canada.pdf> [retrieved 2018-05-11]
- [44] Trossen, D et al. “Value Chain Dynamics in the Communication Industry.” Value Chain Dynamics Working Group (VCDWG), MIT Communications Futures Program and Cambridge University Communications Research Network, December 2005. [online] <https://cfp.mit.edu/docs/core-edge-dec2005.pdf> [retrieved 2019-01-11]
- [45] Ramakrishnan, K and Selvarajan, R. “Transforming the Telecom Value Chain with Platformization.” TATA Consultancy Services White Paper, 2017. [online] <https://www.tcs.com/content/dam/tcs/pdf/Industries/communication-media-and-technology/Abstract/Transforming%20the%20Telecom%20Value%20Chain%20with%20Platformization.pdf> [retrieved 2019-01-11]
- [46] GOV.UK. “Money Laundering Regulations: who needs to register.” 2014-10-23. [online] <https://www.gov.uk/guidance/money-laundering-regulations-who-needs-to-register> [retrieved 2017-05-28]
- [47] Better Business Finance. “What are the AML and KYC obligations of a Bank in the UK?” [online] <https://www.betterbusinessfinance.co.uk/aml-and-kyc/what-are-the-aml-and-kyc-obligations-of-a-bank-in-the-uk> [retrieved 2017-05-28]
- [48] Juniper Networks. “Understanding Online Certificate Status Protocol.” 2018-02-27. [online] [https://www.juniper.net/documentation/en\\_US/junos/topics/concept/certificate-ocsp-understanding.html](https://www.juniper.net/documentation/en_US/junos/topics/concept/certificate-ocsp-understanding.html) [retrieved 2018-05-11]
- [49] Society for Worldwide Interbank Financial Telecommunication. “Discover SWIFT.” [online] <https://www.swift.com/about-us/discover-swift> [retrieved 2018-04-21]
- [50] The Accredited Standards Committee. “ASC X12.” [online] <https://web.archive.org/web/20140927153741/http://www.x12.org/x12org/about/index.cfm> [retrieved 2018-04-21]
- [51] United Nations Economic Commission for Europe. “Trade Programme – Trade – UNECE.” [online] <http://www.unece.org/tradewelcome/trade-programme.html> [retrieved 2018-04-21]
- [52] Ripple – One Frictionless Experience To Send Money Globally. [online] <https://ripple.com/> [retrieved 2019-01-11]
- [53] IBM. “Hyperledger: blockchain collaboration changing the business world.” [online] <https://www.ibm.com/blockchain/hyperledger.html> [retrieved 2019-01-11]
- [54] Securities and Exchange Commission (US). *REGULATION NMS*. 17 CFR PARTS 200, 201, 230, 240, 242, 249, and 270; Release No. 34-51808; File No. S7-10-04, RIN 3235-AJ18, 2005-06-09. [online] <https://www.sec.gov/rules/final/34-51808.pdf> [retrieved 2019-01-30]
- [55] Securities and Exchange Commission Historical Society. “2000s: Timeline.” [online] <http://www.sechistorical.org/museum/timeline/2000-timeline.php> [retrieved 2018-04-19]
- [56] Reuters. “ICE completes takeover of NYSE.” [online] <https://www.reuters.com/article/us-ice-nyse-sprecher/ice-completes-takeover-of-nyse-idUSBRE9AB16V20131112> [retrieved 2018-04-19]