TU Delft

Delft University of Technology

Distributed Attestation Revocation in Self-Sovereign Identity

Chotkan, Rowdy; Decouchant, Jérémie; Pouwelse, Johan

DOI
10.1109/LCN53696.2022.9843323

Publication date
2022

Document Version
Final published version

Published in
Proceedings of the 47th IEEE Conference on Local Computer Networks, LCN 2022

Citation (APA)
Chotkan, R., Decouchant, J., & Pouwelse, J. (2022). Distributed Attestation Revocation in Self-Sovereign Identity. In S. Oteafy, E. Bulut, & F. Tschorsch (Eds.), *Proceedings of the 47th IEEE Conference on Local Computer Networks, LCN 2022* (pp. 414-421). Article 9843323 (Proceedings - Conference on Local Computer Networks, LCN). IEEE. https://doi.org/10.1109/LCN53696.2022.9843323

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

https://www.openaccess.nl/en/you-share-we-take-care

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Distributed Attestation Revocation in Self-Sovereign Identity

Rowdy Chotkan
Distributed Systems
Delft University of Technology
Delft, The Netherlands
R.M.Chotkan-1@tudelft.nl

Jérémie Decouchant
Distributed Systems
Delft University of Technology
Delft, The Netherlands
J.Decouchant@tudelft.nl

Johan Pouwelse
Distributed Systems
Delft University of Technology
Delft, The Netherlands
J.A.Pouwelse@tudelft.nl

2022 IEEE 47th Conference on Local Computer Networks (LCN) | 978-1-6654-8001-7/22/$31.00 ©2022 IEEE | DOI: 10.1109/LCN53696.2022.9843323

***Abstract*—Self-Sovereign Identity (SSI) aspires to create a standardised identity layer for the Internet by placing citizens at the centre of their data, thereby weakening the grip of big tech on current digital identities. However, as millions of both physical and digital identities are lost annually, it is also necessary for SSIs to possibly be revoked to prevent misuse. Previous attempts at designing a revocation mechanism typically violate the principles of SSI by relying on central trusted components. This lack of a distributed revocation mechanism hampers the development of SSI. In this paper, we address this limitation and present the first fully distributed SSI revocation mechanism that does not rely on specialised trusted nodes. Our novel gossip-based propagation algorithm disseminates revocations throughout the network and provides nodes with a proof of revocation that enables offline verification of revocations. We demonstrate through simulations that our protocol adequately scales to national levels.**

***Index Terms*—Self-Sovereign Identity, revocation, offline verification**

I. Introduction

In our modern societies, citizens do not own their identities. The European Union recently announced that it would maintain a trusted and secure digital identity for each European citizen [1]. The majority of current digital identities are also maintained by Big Tech, which results in potential privacy issues as the digital presence of citizens can be monitored [2]. Furthermore, these digital identities can be revoked at the platform owner's discretion leading to loss of access to a plethora of other dependent connected services [3]. The *Self-Sovereign Identity* (SSI) concept overcomes these digital and societal issues by relying directly on the Internet, which currently does not embed any native method to determine who is communicating with whom [4]. As such, the SSI movement aims to create a standardised identity layer for the Internet, generating digital trust through verifiable identities and putting citizens at the centre of their data.

Previous works laid out the relevant principles and architectures of SSI [5]. However, in particular, SSI must be able to handle compromised identities, which might appear as a consequence of theft, loss, or a data breach. For the past five years in the USA, more than a million data breaches occurred annually [6], resulting in the loss of billions of credentials. Furthermore, 0.8% of UK passports [7] and 340,000 identity documents in The Netherlands [8] are lost annually. Revocation of these credentials is required to minimise further potential negative consequences.

Identity revocation remains a key technical challenge in Self-Sovereign Identity. As portrayed by Table I (further discussed in section III), previous SSI distributed revocation designs paved the way but they are still incomplete. Existing SSI and digital identity solutions such as Sovrin[1], Veramo[2] (formerly known as uPort) and IRMA[3] violate the principles of SSI itself by addressing revocation through centralisation and trusted third parties, whereas the cardinal requirement for SSI is an authoritarian-free ecosystem. Furthermore, recent natural disasters demonstrate that assuming the presence of always up-and-running digital infrastructure is not safe [9]. Digital identities are to be disaster-proof. Dependence on central parties for verification prevents offline usability and, moreover, introduces inherent inequalities in the network, leading to censorship or privacy issues [10].

In this paper, we address the identity revocation problem and alleviate an important issue that has been hampering the mass deployment of Self-Sovereign Identities. In a summary, we make the following contributions. We present the first revocation protocol for SSI that is fully distributed, supports offline verification of revocations and does not rely on additional trust assumptions. We evaluate this protocol using our pioneering serverless phone-to-phone infrastructure [11, 12], and a fully functioning SSI application that is backed by the Dutch government, which demonstrate the usability of distributed revocation on smartphones at a national level.

II. Problem Description

Because of theft or loss, digital identities may become compromised. To mitigate further damage, compromised credentials must be revoked. Revocation is also required when a credential becomes (prematurely) voided, e.g., an employee who is no longer employed by a company should no longer be given access to its infrastructure.

Figure 1 portrays the interactions between the three relevant parties of an SSI system following the definitions set out by the W3C [28]. An Issuer attests to a claim of a Subject by creating

This work was funded by NWO/TKI grant BLOCK.2019.004.

[1] https://sovrin.org/ [2] https://veramo.io/ [3] https://irma.app/

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

TABLE I: Comparison of existing revocation solutions

	Domain	Type	Mature[1]	Description	No network operators	Offline availability	No authority interactivity	Offline verification	No SPOF	Full accuracy
This work (section V)	SSI	Attestation	✓	First fully distributed SSI revocation mechanism.	✓	✓	✓	✓	✓	✓
Abraham et al. [13]	SSI	Attestation	✓	Revocations stored on public permissioned blockchain.	✗	✓	✓	✓	✓	✓
Baars [14]	SSI	Credential	✗	Revocations stored on smart contracts.	✗	✗	✗	✗	✗	✓
Eschenauer and Gligor [15]	DSN	Node	✗	Single authority propagates revocations.	✗	✓	✗	✓	✗	✓
Haas et al. [16]	VANET	Certificate	✗	RSUs and v2v propagation.	✗	✓	✓	✓	✓	✗
IRMA [17]	SSI	Attestation	✓	Uses centralised database.	✗	✗	✗	✗	✗	✗
Laberteaux et al. [18]	VANET	Certificate	✗	RSUs and v2v propagation.	✗	✓	✓	✓	✓	✓
Lasla et al. [19]	C-ITS	Node	✗	Revocations stored on blockchain and RSUs.	✗	✓	✓	✗	✓	✓
Liau et al. [20]	P2P	Certificate	✗	Uses distribution points and P2P communication.	✗	✓	✗	✓	✓	✓
Popescu et al. [21]	DS	Certificate	✗	Revocations handled locally by authority.	✓	✗	✗	✓	✗	✗
Sovrin [22]	SSI	Attestation	✓	Uses public permissioned blockchain.	✗	✗	✓	✗	✓	✓
Speelman [23]	SSI	Credential	✗	Uses active verification with issuer.	✓	✗	✗	✗	✗	✓
Stokkink et al. [24, 25]	SSI	Attestation	✓	(Central) revocation registers, DLs, validity terms.	✓	✗	✗	✗	✗	✓
Veramo (uPort) [26]	SSI	Attestation	✗	Uses public permissionless blockchain.	✓	✗	✓	✗	✓	✓
Xu et al. [27]	SSI	Node	✗	List of accepted nodes stored on blockchain.	✗	✓	✓	✗	✓	✓

[1] Refers to maturity of the technology.

an attestation (step 1). A Verifier is then able to determine the validity of said claim by cryptographically verifying the attestation (step 2). In the instance that the attestation is to be revoked (step 3), the verification of the attestation by the Verifier must fail. The Subject can not be trusted to make the revocation available to the Verifier, as this sensibly goes against its own interest. Furthermore, revocations must be disseminated to any party that needs to verify the corresponding credential.

Bringing the Verifier in direct contact with the Issuer would go against the principles of Self-Sovereign Identity as this would defeat the purpose of attestations [5]. Furthermore, the cardinal requirement of SSI is that no third party is required or able to observe or otherwise interfere with the creation or verification of identity data [25].

Existing revocation mechanisms typically introduce centralised mechanics to handle revocations [17, 22] or require Proof-of-Work blockchains [26]. Both methods have limitations. First, relying on a centralised infrastructure may lead to censorship or privacy issues [10]. Second, blockchains suffer from privacy issues [29], low throughput and limited flexibility [30]. Furthermore, they are prone to legislation limiting their use [31].

The lack of a fully distributed revocation mechanism limits the mass deployment of Self-Sovereign Identities. It remains an open problem to design a revocation mechanism that does not depend on a central infrastructure or on third parties during verification and allows clients to independently verify credentials to prevent censorship.

We formulate the following problem description. First, an Issuer revokes their attestation for a credential. This revocation must be made apparent to all clients that may verify the credential and acknowledge the Issuer as an attestor. All clients are interconnected by a network overlay that they establish and maintain, and in which they have equal permissions. As direct communication with an Issuer or reliance on centralised infrastructure for verification goes against the principles of Self-Sovereign Identity, the propagation of revocations must be decentralised. Furthermore, neither an Issuer nor a receiving party can be expected to be online at all times, yet, all revocations are to be spread across the network in order to reach Verifiers. Furthermore, as all clients have equal permissions, they have to be able to individually decide whether or not to accept revocations.

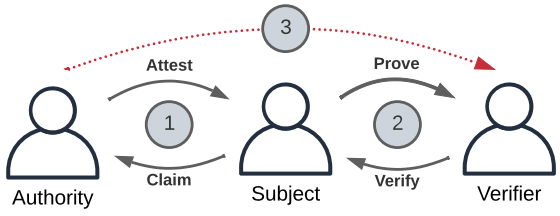

Fig. 1: Interactions in a Self-Sovereign Identity system.

III. Related Work

Table I summarises our analysis of the revocation state-of-the-art in identity systems, and precises their respective limitations. This table also compares our revocation mechanism (see section V) to the related work. We consider the following characteristics: maturity of the solution, network operator requirement, availability of offline revocations, reliance on interactions with authorities (i.e., Issuers in the case of SSI), possibility of offline verification of revocations, presence of single points of failure (SPOFs), and accuracy of the verification mechanism (e.g. false positives or false negatives). As one can see, our solution qualitatively outperforms existing solutions. We note that blockchains allow for the realisation of distributed revocation. However, they suffer from obstacles such as privacy and security issues and low throughput [30].

As our key contribution addresses revocation, we focus on related work that discusses this topic. We found out that revocation in Self-Sovereign Identity is not widely discussed in academia, and as such, we selected works that address distributed revocation in a broader context. We organise the related works in groups that focus on the revocation of (SSI) credentials, certificates, and nodes.

Sovrin [22] and IRMA [17] propose the usage of cryptographic accumulators for *revocation of (SSI) credentials* based on the works of Camenish et al. [32]. A cryptographic accumulator is a probabilistic data structure that allows a large set of values to accumulate into a short witness value that can then be used to prove certain membership operations

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

(e.g. inclusion checks). In the aforementioned solutions, a subject provides a proof of non-revocability of their credentials through this witness value. A verifier can then check this proof using the witness value, which is published on the blockchain. Sovrin does not allow for offline verification of credentials as both the subject and the verifier are required to retrieve the latest witness value during the verification of a credential. Similarly, IRMA does not allow offline verification because communication with its infrastructure is required. Furthermore, cryptographic accumulators can be computationally expensive to the extent that it is discouraged to use them for each verification in IRMA [17] and their probabilistic nature is prone to false positives. Veramo (uPort) [26] uses a single Ethereum smart contract for marking attestations as revoked. The usage of the Ethereum blockchain requires synchronisation of blocks in order to guarantee certainty on stored revocations. Furthermore, a single smart contract introduces a security risk [33]. Xu et al. use a blockchain to store legitimate subjects, indirectly disallowing access for revoked subjects in the SSI system [27]. Updating this set is performed by the operators of the blockchain, which introduces centralised authorities. Abraham et al. propose the usage of a revocation list stored on a blockchain, on which consensus is reached through the nodes of the blockchain, maintained by operators [13]. Offline verification is achieved through the storage of this revocation list. As the revocation list is not stored per authority, clients require full storage of this list, leading to storage overhead. We note that all revocations in an SSI system can grow up to gigabytes of storage, which hinders the deployment on devices with low memory (e.g. smartphones). Furthermore, the usage of a blockchain introduces further overhead as clients have to synchronise blocks. Stokkink et al. propose a fully distributed SSI system using direct peer-to-peer communication [24, 25]. They allow for three revocation mechanisms: i) linkage to a central revocation register, belonging to the Issuer; ii) the usage of distributed ledger technology; iii) the usage of short validity terms. The centralised approach opposes the sovereignty of the protocol, something that is acknowledged by the authors, blockchains again suffer from the aforementioned issues and short validity terms place too much power in the hands of Issuers. Speelman implements the approaches of [24, 25], and additionally, proposes an active check as the main revocation method [23]. Baars proposes the usage of smart contracts for storing revocations of attestations [14], introducing security risks [33].

Mechanisms for the *revocation of PKI certificates* are present in traditional Public Key Infrastructures (PKIs) such as PKIX [34]. Broadly speaking, a PKI uses a Certificate Authority (CA) to publish a Certificate Revocation List (CRL), containing revoked certificates. In this structure, CAs are inherently central authorities, having relatively absolute power over revocations. These CAs, acting as trusted third parties, are central points of failure, suffer from MITM attacks, and are corruptable [35].

PGP's web of trust [36] attempted to overcome this by handling revocation in a decentralised fashion, in which the revocation of keys was handled by the owner through revocation certificates. These certificates indicate that the key was compromised and should therefore no longer be used. However, PGP and its web of trust have been shown to be impractical [37] and require central key servers. Another alternative to PKI is the *Decentralised Public Key Infrastructure* (DPKI) [35, 38]. DPKI proposes the usage of alternative storage solutions for storing revocations of public keys. The proposed solutions use blockchains and, thus, require synchronisation of blocks for verification, introducing overhead and possibly low throughput as discussed previously.

Laberteaux et al. discuss the revocation of PKI certificates in vehicular ad hoc networks (VANETS) through the distribution of CRLs [18]. Distribution is handled through Road Side Units (RSUs), serving as specialised nodes propagating the CRLs, and through epidemic spread between vehicles. The revocations are stored in Bloom filters. Haas et al. build upon this work by guaranteeing a certain degree of privacy by using group signatures [39] when requesting certificates from the CA [16]. However, the revocations are handled by a single CA and the reliance on Bloom filters introduces the possibility of false positives. Liau et al. propose the distribution of CRLs through direct peer updates, reducing the communication overhead caused by periodic CRL synchronisation [20]. Signatures over CRLs allow nodes to build trust in others. However, direct peer updates may prove to be suboptimal in the case of highly adaptive networks such as that of mobile devices. Propescu et al. discuss the revocation of certificates based on the clustering of clients and probabilistic auditing for honesty of revocation distributors [21]. This auditing is probabilistic in order to reduce performance requirements, however, this allows for malicious nodes to possibly exist for quite some time.

Eschenauer and Gligor discuss the *revocation of nodes* in distributed sensor networks [15]. Revocation is handled by a single node serving as an authority, delegating revocations to regular sensor clients. We note that the introduction of a single authority goes against the principles of SSI. Lasla et al. discuss the revocation of malicious vehicles in Cooperative Intelligent Transportation Systems (CITS) [19]. They use a blockchain for storing revocations through a distributed vehicle admission and revocation scheme. Again, we note that blockchains suffer from the aforementioned hurdles such as privacy and security issues.

IV. System & Threat Model

We focus on the revocation of attestations and related SSI interactions in an *identity network*. An identity network is a peer-to-peer (P2P) network that implements the identity service that is maintained and used by its peers. The network can be openly joined by any peer (also called a client) at any time. We build upon the terminology of the W3C's DID to define the potential roles of peers [40]. A peer can assume one or several roles among those of Subject, Issuer and Verifier: a Subject is a client that holds credentials; an Issuer attests to a claim and is able to revoke its attestation; a Verifier verifies credentials. Any two clients are assumed to eventually be able

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

to directly communicate with each other. We assume that clients are always connected to a subset of other clients, which are called their neighbours, with which they exchange identity-related information. Moreover, clients are deemed not to be necessarily always online. However, when they are online they are reachable and can communicate with other peers. As such, nodes periodically exchange membership information with their neighbours, replace them, and exchange revocation information with them.

Adversaries or malicious actors may be present in the network. We assume that adversaries do not aid in maintaining the health of the network (via the spread of revocations discussed in section V). These actors are not able to drop arbitrary messages but are able to send fabricated messages (e.g., replayed) to other clients. These messages are discarded by honest nodes as they can be evaluated as invalid.

V. Architecture & Theoretical Analysis

Our revocation mechanism overcomes the hurdle of interactivity with authorities whilst enabling offline verification. Clients do not require to be online during verification of a credential, they merely require occasional synchronisation of revoked attestations by communicating with other nodes in the network. This is achieved through three concepts.

A. Trusted Issuers

In real life, a person has (relatively speaking) a choice of whether to acknowledge a certain authority. Following this fashion, this is also possible in our revocation architecture: each client manages, as we coin, a *Trusted Issuer Storage* (TIS). The TIS is a register containing the public keys of the Issuers that are trusted by a client. Each of these Issuers is referred to as a *Trusted Issuer* (TI). Hence, a distinction is made by each client, individually, on the trusted authorities in the network. A client exclusively accepts the revocations made by their TIs. The results of acceptance are the storage of the revocations by the client and further propagation of the revocations in the network. This significantly reduces storage requirements under the assumption that a client is not interested in revocations made by an Issuer that physically resides far away. TIs distinguish themselves from traditional TTPs as they are only able to influence the legitimacy of attestations that they created. Furthermore, they only aid in the verification process and, thus, do not provide a definite answer.

Issuers publish their revocations as sets containing the hashes of revoked credentials. These sets are uniquely identified using a label and subsequent sets only contain new revocations. Our implementation uses the SHA3-256 algorithm for hashing and an incremental integer for version labelling. The version label is unique for each Issuer, but not across revocation sets made by other Issuers. The Issuer signs its set of revocations and the label for authenticity. Hence, an issuer revokes its attestation by publishing the hash of the credential that the attestation belongs to. This counteracts the possibility for clients to hide a revoked attestation.

B. Attestation Revocation List

All received revocations are stored by a client for later reference in, what we coin, its *Attestation Revocation List* (ARL). The ARL is a register holding the revocations made by the TIs it trusts. It is, similarly to the TIS, stored and managed by each client individually. In the ARL, revocations are grouped by the TI that revoked the attestation and by the unique version label that is assigned by the TI. This makes it possible for a client to store duplicate revocations if multiple TIs revoked their attestation for the same credential. This is per design, as it allows Verifiers to build more trust in rejecting a revoked credential. Furthermore, storing revocations per TI counteracts malicious TIs from revoking attestations made by other Issuers.

As the number of revocations can grow to a large amount, we use Bloom filters [41] to speed up the verification process for Verifiers, whilst overcoming their probabilistic nature. All Verifiers, in addition to storing the attestations, add them to a Bloom filter. Using their filter, Verifiers first test whether an attestation belongs to the ARL, after which, upon success, the definitive search is performed (to handle possible false positives). Raya et al. [42] discuss the benefits of Bloom filters in Certificate Revocation Lists, which can provide similar speed improvements for the ARL, as both require validation of whether an item is part of a set of revoked items.

Furthermore, we note that the ARL can be replaced exclusively by a probabilistic data structure. A client may choose to accept the probabilistic nature of a Bloom filter over an exact membership check. Such clients are not able to aid in the propagation of the revocations, though the low memory requirements may sometimes prove to make the protocol suitable, e.g., for IoT devices. However, as a result, verification of credentials on the client may be affected by false positives. Whilst this does not explicitly impact security, it could lead to the false rejection of non-revoked credentials. As such, this is only suitable for Verifiers that expect to verify low numbers of credentials. For perspective: a Bloom filter of 907.24 KiB, using 10 hash functions, storing 100,000 items has a false positive probability of 1 in 1 billion.

C. Propagation

In order to disseminate revocations, the architecture requires a protocol that ensures that information is spread across the entire network, whilst also ensuring that unavailable clients receive the information at a later instance. For this, we use a gossip protocol with static re-transmissions. Gossip protocols are communication protocols that allow for a periodic exchange of data with (random) peers [43].

In order to maintain a low overhead and allow for selective revocation updates, revocations are propagated using advertisements. Gossiping nodes advertise their known revocations, after which a receiving node is able to selectively request revocations. Advertisements are structured as key-value pairs of the digest of a TI's public key and the latest version known by the gossiping client. A receiving client requests updates by sending back key-value pairs of the digests of each TI's public

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

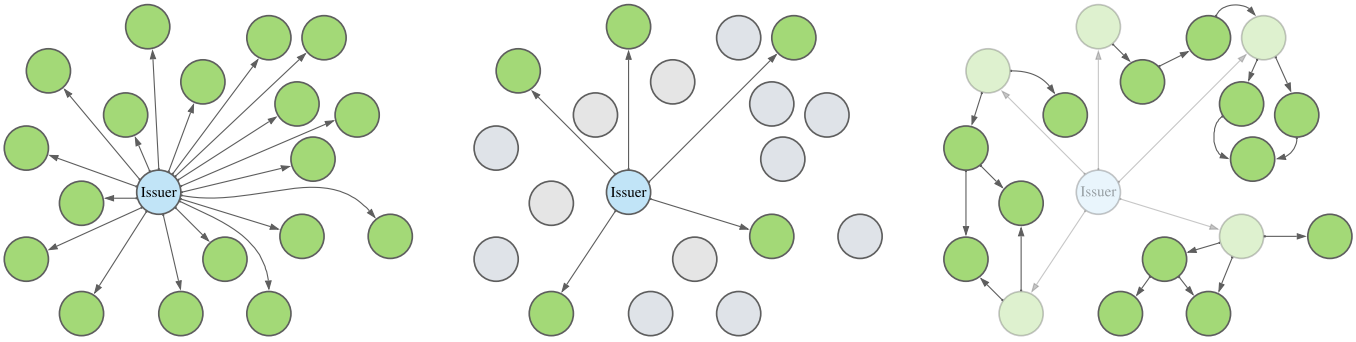

(a) Centralised revocation propagation (b) Direct propagation from Issuer (c) Indirect propagation through gossip

Fig. 2: Fully distributed revocation

key it is interested in and the lowest version it is missing. The gossiping client then sends back all the requested revocations. Authenticity is guaranteed using the signatures. Based on the publication dates attached to revocations, clients are able to ignore revocations, optimising storage usage as they may no longer be relevant in the system due to expired validity terms of the corresponding attestations.

Figure 2 illustrates the communication of revocations from an initial Issuer to a set of clients. Figure 2a portrays the traditional client-server model in which all revocations are received by the clients through a single central party, indicated by the green-coloured clients. This system is limited by the requirement of a direct link between all nodes and the central Issuer. A central Issuer may introduce security and availability issues and possibly leads to censorship.

Our gossip-based approach utilises the structure visible in Figure 2b and Figure 2c. Figure 2b portrays the initial gossip from an Issuer to a set of neighbours. In this instance, all clients are honest, acknowledge the Issuer, and are online during the propagation window (the time at which the Issuer has its gossip iteration). After which, the clients continue to propagate the revocations, in the same fashion, eventually spanning across the network. This is depicted in Figure 2c. Thus, after the initial gossip by the Issuer, no further interaction with said Issuer is required for the propagation of the revocations. This figure also shows relays to clients that are already gossiped to, indicating that clients have multiple opportunities to receive the latest revocations.

Our algorithm has a worst-case runtime of $\mathcal{O}(n)$, with n being the network size, as in the worst case a single node updates the entire network. However, it is expected to be logarithmic with respect to the number of nodes, as each node that has received the latest revocations from a TI can gossip to the remaining uninformed nodes, speeding up the propagation time for the remaining nodes. Furthermore, not the entire network is interested in all revocations.

As becomes apparent from this description, dishonest nodes pose no large threat to the propagation of revocations. They could introduce a slight delay due to fewer clients gossiping information or due to the spread of fabricated revocations, which will be discovered by receiving clients. We do note that depending on the network topology, Eclipse attacks [44] are a possibility that can be circumvented through direct connections with Issuers wherever possible.

VI. Algorithms & Simulation

In order to realise the proposed revocation mechanism discussed in section V, each client in the network runs three algorithms. A gossiping client runs algorithm 1, which enables the periodic advertisement of revocations. First, a random subset of peers is generated using the node selection function (line 2). This function generates a set of peers from the neighbours known by the client. Next, an advertisement is gossiped to each of these peers (lines 3-4). Finally, the gossiping client awaits the start of the next gossip interval (line 5). An advertisement consists of pairs of Issuer public keys and their latest known version of revocations.

A node receiving an advertisement runs algorithm 2. In this procedure, the node verifies whether any TI is present in the advertisement (lines 2-3). Then it verifies, using the `FindMissingVersion` routine, whether it misses or is uninformed about revocation versions belonging to the Trusted Issuer (lines 4-5). More specifically, `FindMissingVersion` determines on an advertisement containing the revocation version v_i part of revocations made by TI a_i with public key pk_i whether $\exists(v_j, pk_i) \in \mathcal{ARL}$ such that $\forall(v_k, pk_i) \in \mathcal{ARL}$ it holds that $(v_j \geq v_k \wedge v_j < v_i) \vee (v_{j+1} \notin \mathcal{ARL} \wedge v_{j+1} < v_i)$. If this is the case, an update is requested from the gossiping client for the respective Issuer and the lowest missing version (line 6). The advertising node verifies whether it advertised to the node recently and sends the revocations.

Following the reception of requested revocations, a node executes algorithm 3. This procedure verifies the relevance of the revocations (line 1) and their validity (line 2). This validity check is performed by verifying the attached signature over the revocations and their version using the public key of the TI. Finally, the revocations are stored in the ARL (line 3).

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

A. Simulation

The analysis of the mechanism is two-fold. Firstly, we discuss a simulation showcasing scalability amongst a relatively high number of clients using the aforementioned algorithms. Secondly, we showcase analysis through the deployment on smartphones in section VII. The simulation was performed on a system with an Intel i7-6700HQ CPU clocked at 2.60 GHz and 16 GB of RAM.

Algorithm 1: Revocation advertisement gossip

input : $\mathcal{C}$ Set of neighbours
$\mathcal{A}$ Set of known Issuer-version pairs
t_g Gossip interval
n_g Gossip amount
output: Revocation advertisements

```
while True do
    C_g ← SelectPeers (C, n_g);
    foreach c_i ∈ C_g do
        GossipRevocations(c_i, A);
    Wait(t_g);
```

Algorithm 2: Revocation update request procedure

input : $\mathcal{A}$ Set of known Issuer-version pairs
output: Revocation update request

```
On reception of A by Client c_i;
for Issuer a_i, Version v_j in A do
    if a_i ∈ TIS then
        v_local ← FindMissingVersion(a_i);
        if v_local < v_j then
            RequestUpdate(c_i,a_i,v_local);
```

Algorithm 3: Revocation reception

input : R Set of revocations
v_i Revocations version
s_i Signature
pk_i Issuer public key
output: $R \subseteq \mathcal{ARL}$

```
if Issuer pk_i in TIS then
    if Verify(pk_i, s_i, v_i|R) then
        ARL ← ARL ∪ R;
    else
        ⊥
else
    ⊥
```

The simulation has been performed through the mimicking of gossip of 340,000 revocations between clients released by a single client serving as a revoking Issuer. Each client runs the three algorithms and acknowledges the Issuer as a TI. The measurements of the simulation were gathered by simulating

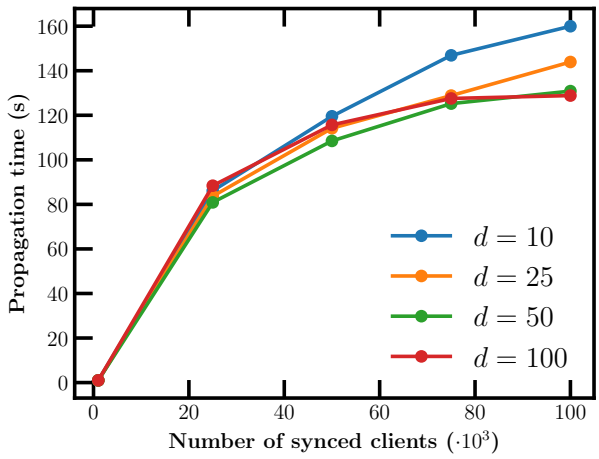

Fig. 3: Revocation scaling

the execution time of the algorithm until full propagation of the revocations across a simulated network comprised of 100,000 nodes. The network is simulated using a weighted regular graph with differing degrees. A uniformly distributed weight between 0 and 20ms is introduced per link to simulate the impact of network latency, based on the global average reported by [45]. Furthermore, the cost associated with the transfer of 340,000 SHA3-256 hashes of 32 bytes is simulated using the global average upload speed of around 65 Mbps [45]. The revocations are released on $t = 0$ by the Issuer. Each simulation is repeated 5 times. Due to multiple paths existing to each node, a client has multiple opportunities to receive revocations. Revocations are gossiped to all neighbours on reception.

B. Simulation Results

The averages of the timings are visible in Figure 3. As expected, increasing the number of neighbours d per client leads to lower propagation times. The simulation using $d = 100$, however, shows a small decrease in performance when operating under fewer clients. Nonetheless, at 100,000 clients, this simulation leads to the lowest propagation time as expected. The observed decrease in performance, under fewer clients, can be explained by the overhead introduced by gossiping to 100 clients upon receiving new revocations. The results seem to indicate that the propagation increases logarithmically with respect to the number of clients and that the number of neighbours should be bound, not only due to performance limitations but also due to overhead in communication introduced by advertising to a larger number of clients. Overall, the simulation portrays that the algorithm is able to achieve realistic timings, taking up to 160 seconds to achieve propagation. In blockchain solutions, the propagation time depends on when the next block containing the revocations is mined, where the average block time on e.g. Bitcoin is 10 minutes plus transaction fees.

VII. Implementation & Performance Analysis

Sections V & VI presented a novel fully distributed revocation algorithm with offline verification capabilities for Self-

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

Sovereign Identity systems. Based on this architecture, two implementations have been made using a pioneering phone-to-phone infrastructure we developed previously[4] [11, 12]. This serverless Web3 fabric allows for direct client-to-client communication, enabling a fully distributed infrastructure at the core of the solution. Using this implementation, we wrote an application for Android, backed by the Dutch National Office for Identity Data (RvIG) [46], showcasing the feasibility on smartphones. The resulting implementations can be found in our public repositories[5,6].

The analysis of the implementation has been performed in a test setup measuring the time required to gossip revocations between an Issuer and three Verifier clients running on smartphones. For revocations, we generated a dataset of 1 million revoked 32 bytes SHA3-256 hashes, a format used by the implementation. Revocations were split up into sets of 1000 in order to minimise the impact of a single packet loss. For the default parameters, the gossip-interval t_g was set to 100 ms in order to maximise the throughput of gossip. The number of selected peers n_g was set to 5 as our phone-to-phone infrastructure uses 20 simultaneous connections per default. However, due to the network size of the test setup, the parameters are of minor impact.

TABLE II

Phone	Propagation time (s)
Galaxy s10	1066
Pixel 2 XL (emul.)	903
Pixel 4 (emul.)	801

A. Revocation Amount

Table II showcases the revocation scaling in a system of 1 client gossiping revocations and 3 clients receiving revocations. Our results indicate that the propagation time scales linearly with respect to the number of revocations. One million revocations take up to 1066 seconds or just under 18 minutes. As this can be deemed more than 4 years' worth of revocations [8] in the Netherlands, we deem this scalability usable as the propagation is expected to grow logarithmic with respect to the number of clients in larger networks.

Compared to the simulations discussed in section VI, the performance is worse. We note that this can be explained mostly due to communication overhead caused by UDP packet splitting. The tremendous amount of packets led to many packet drops, in turn leading to the loss of specific revocation versions. As the reference implementation naively provides the gossiping client with a lower bound of missing versions, the additional network traffic of already gossiped versions causes additional packet losses. This snowballing effect worsens the performance of the algorithm. As such, the investigation of other network protocols or more sophisticated handling of packet loss can prove to significantly improve performance. However, the achieved performance can be deemed usable.

[4] Official (Python) documentation: https://py-ipv8.readthedocs.io/en/latest/
[5] Infrastructure: https://github.com/Tribler/kotlin-ipv8
[6] Android application: https://github.com/Tribler/trustchain-superapp

VIII. Conclusion

This paper addresses revocation in Self-Sovereign Identity systems. We deem revocation to be the last remaining open issue for SSI to become a feasible contender for the next generation of identity systems. We proposed the first fully distributed revocation mechanism requiring no interactivity with any central parties, whilst adhering to the principles of the SSI paradigm. Revocations are propagated through the network using a gossip-based protocol, in which the acknowledgement of revocations is up to the discretion of Verifiers. Offline verification is enabled through local storage and no dependency on revoking Issuers. The feasibility of this revocation mechanism has been validated using a fully distributed SSI implementation. Our results show that fully distributed SSI is feasible on modern smartphones and that this is a promising direction to further explore. We conclude that our proposed architecture is a sizeable step towards placing identity back into the hands of the citizens.

References

[1] The European Commission, "Commission Recommendation (EU) 2021/946 of 3 June 2021 on a common Union Toolbox for a coordinated approach towards a European Digital Identity Framework," Jun. 2021. [Online]. Available: http://data.europa.eu/eli/reco/2021/946/oj/eng

[2] O. Tene and J. Polonetsky, "Big data for all: Privacy and user control in the age of analytics," *Nw. J. Tech. & Intell. Prop.*, vol. 11, p. xxvii, 2012.

[3] R. Rogers, "Deplatforming: Following extreme internet celebrities to telegram and alternative social media," *European Journal of Communication*, vol. 35, no. 3, pp. 213–229, 2020.

[4] K. Cameron, "The laws of identity," *Microsoft Corp*, vol. 5, pp. 8–11, 2005.

[5] A. Mühle, A. Grüner *et al.*, "A survey on essential components of a self-sovereign identity," *Computer Science Review*, vol. 30, pp. 80–86, Nov. 2018.

[6] J. Johnson, "Annual number of data breaches and exposed records in the United States from 2005 to 2020," Mar. 2021. [Online]. Available: https://www.statista.com/statistics/273550/data-breaches-recorded-in-the-united-states-by-number-of-breaches-and-records-exposed/

[7] HM Passport Office, Border Force *et al.*, "Report your lost or stolen passport," June 2018. [Online]. Available: https://www.gov.uk/government/news/report-your-lost-or-stolen-passport

[8] G. Brummelkamp, J. Wils, and J. Boog, "Evaluatie proeftuin vermissing reisdocumenten," Panteia, Zoetermeer, Tech. Rep., 12 2015.

[9] P. Zhang and G. Rui, "China floods: 'digital dark age' after disaster wreaks havoc on internet and electricity," July 2021. [Online]. Available: https://www.scmp.com/news/people-

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.

culture/environment/article/3142544/china-floods-digital-dark-age-after-disaster-wreaks
[10] D. Khovratovich and J. Law, "Sovrin: digital identities in the blockchain era," The Sovrin Foundation, Tech. Rep., 2017. [Online]. Available: https://sovrin.org/library/sovrin-digital-identities-in-the-blockchain-era/
[11] G. Halkes and J. Pouwelse, "Udp nat and firewall puncturing in the wild," in *NETWORKING*, 2011.
[12] N. Zeilemaker, B. Schoon, and J. Pouwelse, "Dispersy bundle synchronization," *TU Delft*, 2013.
[13] A. Abraham, S. More *et al.*, "Revocable and offline-verifiable self-sovereign identities," *TrustCom*, 2020.
[14] D. Baars, "Towards self-sovereign identity using blockchain technology," Master's thesis, 2016.
[15] L. Eschenauer and V. D. Gligor, "A key-management scheme for distributed sensor networks," in *CCS*, 2002.
[16] J. J. Haas, Y. C. Hu, and K. P. Laberteaux, "Efficient certificate revocation list organization and distribution," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 3, pp. 595–604, Mar. 2011.
[17] G. Alpár, F. van den Broek *et al.*, "Irma: practical, decentralized and privacy-friendly identity management using smartphones," *HotPETs 2017*, 2017.
[18] K. P. Laberteaux, J. J. Haas, and Y.-C. Hu, "Security certificate revocation list distribution for vanet," in *ACM International workshop on VANET*, 2008, pp. 88–89.
[19] N. Lasla, M. Younis *et al.*, "Efficient Distributed Admission and Revocation Using Blockchain for Cooperative ITS," in *NTMS*, Mar. 2018, pp. 1–5.
[20] C. Y. Liau, S. Bressan, and K.-L. Tan, "Efficient Certificate Revocation : A P2P Approach," *HICSS'05*, 2005.
[21] B. C. Popescu, B. Crispo, and A. S. Tanenbaum, "A certificate revocation scheme for a large-scale highly replicated distributed system," in *ISCC*, 2003.
[22] A. Tobin and D. Reed, "The inevitable rise of self-sovereign identity," *The Sovrin Foundation*, 2016.
[23] T. Speelman, "Self-Sovereign Identity: Proving Power over Legal Entities," Master's thesis, 2020.
[24] Q. Stokkink and J. Pouwelse, "Deployment of a blockchain-based self-sovereign identity," in *IEEE iThings, GreenCom, CPSCom and SmartData*, 2018.
[25] Q. Stokkink, D. Epema, and J. Pouwelse, "A Truly Self-Sovereign Identity System," *arXiv:2007.00415*, 2020.
[26] C. Lundkvist, R. Heck *et al.*, "Uport: A platform for self-sovereign identity," Oct. 2016. [Online]. Available: https://blockchainlab.com/pdf/uPort_whitepaper_DRAFT20161020.pdf
[27] J. Xu, K. Xue *et al.*, "An identity management and authentication scheme based on redactable blockchain for mobile networks," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 6, 2020.
[28] M. Sporny, D. Longley, and D. Chadwick, "Basic concepts," in *Verifiable Credentials Data Model 1.0*. W3C, Nov. 2019. [Online]. Available: https://www.w3.org/TR/vc-data-model/
[29] J. Yli-Huumo, D. Ko *et al.*, "Where is current research on blockchain technology?—a systematic review," *PloS one*, vol. 11, no. 10, p. e0163477, 2016.
[30] L. Hughes, Y. K. Dwivedi *et al.*, "Blockchain research, practice and policy: Applications, benefits, limitations, emerging research themes and research agenda," *International Journal of Information Management*, vol. 49, pp. 114–129, 2019.
[31] R. Xie, "Why china had to ban cryptocurrency but the us did not: A comparative analysis of regulations on crypto-markets between the us and china," *Wash. U. Global Stud. L. Rev.*, vol. 18, p. 457, 2019.
[32] J. Camenisch and A. Lysyanskaya, "Dynamic accumulators and application to efficient revocation of anonymous credentials," in *CRYPTO*, 2002.
[33] P. Praitheeshan, L. Pan *et al.*, "Security analysis methods on ethereum smart contract vulnerabilities: a survey," *arXiv:1908.08605*, 2019.
[34] IETF, "Public-key infrastructure (x.509) (pkix)." [Online]. Available: https://datatracker.ietf.org/wg/pkix/
[35] C. Allen, A. Brock *et al.*, "Decentralized public key infrastructure. a white paper from rebooting the web of trust," 2015.
[36] P. Zimmermann, "Why I Wrote PGP," 1999. [Online]. Available: https://www.philzimmermann.com/EN/essays/WhyIWrotePGP.html
[37] A. Whitten and J. D. Tygar, "Why johnny can't encrypt: A usability evaluation of pgp 5.0." in *USENIX Sec.*, 1999.
[38] C. Fromknecht, D. Velicanu, and S. Yakoubov, "A decentralized public key infrastructure with identity retention." *IACR Cryptol. ePrint Arch.*, vol. 2014, p. 803, 2014.
[39] D. Chaum and E. Van Heyst, "Group signatures," in *Workshop on the Theory and Application of of Cryptographic Techniques*. Springer, 1991, pp. 257–265.
[40] M. Sporny, D. Longley, and D. Chadwick, "Verifiable credentials data model 1.0," Nov. 2019. [Online]. Available: https://www.w3.org/TR/vc-data-model/
[41] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.
[42] M. Raya, P. Papadimitratos *et al.*, "Eviction of Misbehaving and Faulty Nodes in Vehicular Networks," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, 2007.
[43] M. Kwiatkowska, G. Norman, and D. Parker, "Analysis of a gossip protocol in prism," *ACM SIGMETRICS*, vol. 36, no. 3, pp. 17–22, 2008.
[44] M. Castro, P. Druschel *et al.*, "Secure routing for structured peer-to-peer overlay networks," *ACM SIGOPS Operating Systems Review*, vol. 36, no. SI, pp. 299–314, 2002.
[45] Ookla, "Speedtest global index," July 2021. [Online]. Available: https://www.speedtest.net/global-index
[46] RvIG, "De mens centraal bij Self Sovereign Identity," 2020. [Online]. Available: https://magazines.rvig.nl/idee/2020/13/de-mens-centraal-bij-self-sovereign-identity

Authorized licensed use limited to: TU Delft Library. Downloaded on February 06,2023 at 09:43:03 UTC from IEEE Xplore. Restrictions apply.