



Co-Creating Circular
Resource Flows in Cities

constRuctive mEtabolic processes For materiaL fLOWs in
urban and peri-urban environments across Europe

Dyne.org foundation

REFLOW Portable Crypto Functions

Project deliverable 2.3

Authors: dr. Denis Roio, Alberto Ibrisevic, Andrea D'Intino

Due Date: 31/05/2021

Actual Submission Date: 28/05/2021

Start date of project: 01/06/2019

Project duration: 36 Months

Dissemination Level: Public ✓



*This project has received funding from the European Union's Horizon 2020
research and innovation programme under grant agreement number 820937.*



REFLOW is an EU Horizon 2020 research project running from 2019-2022, which aims to enable the transition of European cities towards circular and regenerative practices. More specifically, REFLOW uses Fab Labs and makerspaces as catalysers of a systemic change in urban and peri-urban environments, which enable, visualize and regulate “four freedoms”: free movement of materials, people, (technological) knowledge and commons, in order to reduce materials consumption, maximize multifunctional use of (public) spaces and envisage regenerative practices. The project will provide best practices aligning market and government needs in order to create favourable conditions for the public and private sector to adopt circular economy (CE) practices. REFLOW is creating new CE business models within six pilot cities: Amsterdam, Berlin, Cluj-Napoca, Milan, Paris and Vejle and assess their social, environmental and economic impact, by enabling active citizen involvement and systemic change to re-think the current approach to material flows in cities.

This document and the free and open source software referenced sits at the core of the innovative developments in REFLOW technical work-package and implements a novel signature scheme for the specific use-case of material passports whose integrity, provenance and portability is granted by means of provable cryptography.

The format of this document is that of an academic publication (Computer Science – Cryptography and Security, cs.CR) which being submitted to open publishing platforms and made freely available to the scientific community under the Creative Commons License (CC BY-NC-SA 4.0), see:

[Creative Commons — Attribution-NonCommercial-ShareAlike 4.0 International — CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/)

The latest version of this document is available on <https://github.com/dyne/reflow-crypto>



REFLOW

ZERO KNOWLEDGE MULTI PARTY SIGNATURES WITH APPLICATION TO DISTRIBUTED AUTHENTICATION

Denis Roio
Dyne.org

Alberto Ibrisevic
Dyne.org

Andrea D’Intino
Dyne.org

May 28, 2021

ABSTRACT

Reflow is a novel signature scheme supporting unlinkable signatures by multiple parties authenticated by means of a zero-knowledge credential scheme. Reflow integrates with blockchains to ensure confidentiality, authenticity and availability even when credential issuing authorities are offline. We implement and evaluate a Reflow smart contract for Zenroom and present an application related to multiple anonymous signatures by authenticated parties and their non-interactive verification. Reflow uses short and computationally efficient authentication credentials and signatures application scale linearly over multiple participants.

1 Introduction

Multi-party computation applied to the signing process allows the issuance of signatures without requiring any of the participating parties to disclose secret signing keys to each other, nor requires the presence of a trusted third-party to receive them and compose the signatures. However, established schemes have shortcomings. Existing protocols do not provide the necessary efficiency, re-randomization or blind issuance properties necessary for the application to trust-less distributed systems. Those managing to implement such privacy preserving features are prone to rogue-key attacks [Boneh et al., 2020] since they cannot grant that signatures are produced by legitimate key holders.

The lack of efficient, scalable and privacy-preserving signature schemes impacts distributed ledger technologies that support ‘smart contracts’ as decentralized or federated architectures where trust is not shared among all participants, but granted by one or more authorities through credential issuance for the generation of non-interactive and unlinkable proofs.

Reflow applies to the signature process a mechanism of credential issuance by one or more authorities for the generation of non-interactive and unlinkable proofs, resulting in short and computationally efficient signatures composed of exactly two group elements that are linked to each other. The size of the signature remains constant regardless of the number of parties that are signing, while the creden-

tial is verified and discarded after signature aggregation. While being signed, duplicates may be avoided by collecting unlinkable fingerprints of signing parties, as they would invalidate the final result. Before being able to sign, a one-time setup phase is required where the signing party collects and aggregates a signed credential from one or more authorities.

Our evaluation of the Reflow functions shows very promising results: basic session creation takes about *20ms*, while signing *73ms* and verification *40ms* on average consumer hardware.

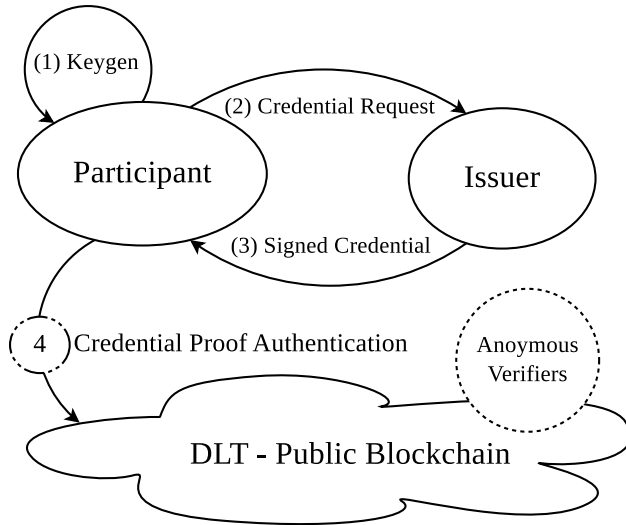
Overview

Reflow provides a production-ready implementation that is easy to embed in end-to-end encryption applications. By making it possible for multiple parties to anonymously authenticate and produce untraceable signatures, its goal is to leverage privacy-by-design scenarios that minimize the information exchange needed for document authentication.

The participation to a signature will be governed by one or more issuers holding keys for the one-time setup of signature credentials. The steps outlined below are represented in figure 1:

- ① Participant generates keys
- ② Participant sends a credential request to Issuer

Figure 1: Basic credential authentication



- ③ Issuer signs the credential request and sends it back
- ④ Participant can create anonymous credential proofs

Following this setup any participant will be able to produce a zero-knowledge proof of possession of the credential, which can be verified by anyone anonymously on the blockchain.

The base application of Reflow is the privacy preserving collective signature of digital documents for which the signed credential proof is a requirement to participate to any signature process.

A Reflow signature process is best described in 3 main steps: session creation, signature and verification.

1. Anyone creates a session A session may be created by anyone, no credentials are required, but only information that should be public: the public keys of participants holding a credential to sign, the public verifier (public signature key) of the issuer who has signed the credentials and at last a document to be signed. The steps below are represented in figure 2 and illustrate how a signature session is created:

- ① Participants will publish their public keys, available to anyone
- ② Anyone may indict a signature session by selecting a document, the public keys of signing participants and the issuer
- ③ The signature session is then published without disclosing the identity of any participant

A Reflow signature session can then be published for verification and its existence may be confidentially communicated to the participants elected to sign it. The possession of the session will allow to disclose the identity of the participants, but only that of the Issuer and the document (or the hash of it) signed.

Figure 2: Session Creation

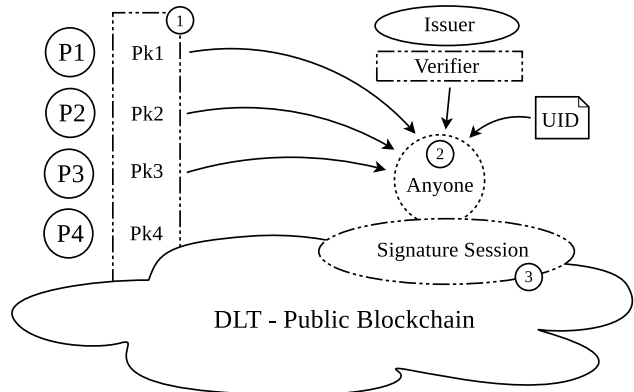



Figure 3: Sign and Verify



2. Participants sign the session Only elected participants that were initially chosen to sign the session may sign it, this is forced through credential authentication. Whenever they know about the existence of a session requesting their signature, they may chose to sign it. The steps below are illustrated in figure 3.

- ① Participants may be informed about the signature session and may create an anonymous signature to be added to the session
- ② Anyone may check that the anonymous signature is authentic and not a duplicate, then adds it to the session
- ③ Anyone may be informed about the signature session and be able to verify if the document is signed by all and only all participants

A delicate aspect of BLS signatures is avoiding double-signing: if a participant signs twice then the whole signature will never be valid. Relying on a stateless credential authentication alone does not avoid this case in Reflow, therefore we use a list of anonymous "fingerprints" of the signature related to the document being signed. Each signature will produce a participant fingerprint saved in a list that is checked against duplication before adding a new signature. This procedure adds significant computational overhead for sessions with a large number of participants, but it can be switched off in a system that avoids double-signing in its own architecture.

Anyone verifies signatures Until the session will have collected all the signatures of participants, its verification will not be valid. It is also impossible to know if all participants have signed the session or how many are missing. Anyone can verify the state of the signature session in any moment just by having the document and the session, as illustrated in figure 3 along with the signature process.

Configurable features may be introduced in the Reflow signature flow that may or may not disclose more information, for instance who has indicated the signature session, what documents are linked to signatures and how many participants were called to sign: this depends from the implementation and the metadata it may add to signature sessions or the communication protocols adopted. In any case the basic signature and verification flow of Reflow requires that only one identity is really made public and is that of the issuer.

Applications

Moving further in envisioning the possibilities opened by Reflow is important to state the possibility to aggregate (sum) signatures into compact multi-signatures, a core feature of our BLS based signature scheme [Boneh et al., 2018a].

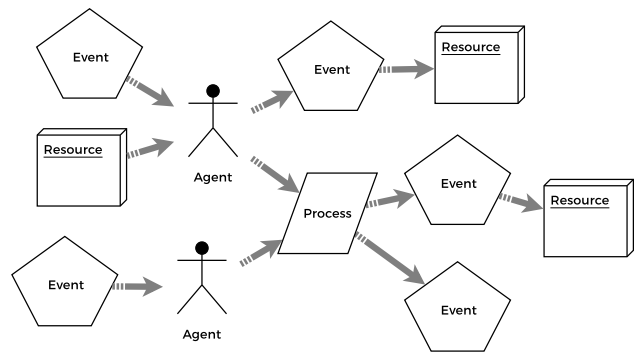
Need to Know. The base implementation exploiting this feature is that of a signature scheme for a single document split in separate sections to be signed by different participants: all the signatures can be later aggregated in a single one proving the whole document has been signed by all participants, without being disclosed to all of them in its entirety. This application helps to enforce the principles of need-to-know and least privilege to the access of information [Saltzer and Schroeder, 1975] and is useful for the realization of privacy-aware applications in various sectors, for instance for medical and risk mitigation analysis.

Disposable Identities A Disposable Identity is based on four properties common to other authentication and identification systems: verifiability, privacy, transparency, trustworthiness; then in addition introduces a fifth property: disposability. Disposability permits purpose-specific and context-driven authentication, to avoid linking the same identity across different authentication contexts [Goulden et al., 2021].

Reflow can be adopted by such an application to remove the need of context-free identifiers and implement authentication functions through a disposable identity whose traceability is bound to a context UID. Furthermore, being a signature scheme, Reflow can add the feature of context-free verifiable signatures (and multi-signatures) that are untraceable in public, but can be traced and even revoked in a known context by the means of context-specific signature fingerprints.

This scenario is relevant for the implementation of privacy-preserving public sector applications that allow authenti-

Figure 4: Valueflows



cation and signatures through disposable identity systems [Kranenburg and Gars, 2021].

Material Passport. Drawing on feature of multiple signature aggregation, Reflow can be used to implement a material passport for circular economy applications [Luscuere, 2017] to maintain the genealogy of a specific product, providing authenticated information about the whole set of actors, tools, collaborations, agreements, efforts and energy involved in its production, transportation and disposal [Dyne.org, 2020].

The provision of the information that forms the content of the material passport should be done by every actor in the supply chain and among the most important technical necessities for such an application are the confidentiality issues regarding access to information and the guarantees of the quality of information [Damen, 2012].

As an ideally simple and effective ontology we adopt a Resource-Event-Agent model [Laurier et al., 2018] and the ValueFlows vocabulary [Foster et al., 2017] to organize knowledge in a graph made of 3 main type of nodes:

- ① Events that are the means of creation and transformation of Resources
- ② Agents capable of creating Events and Processes
- ③ Processes where more than one Agent can consume and create Events

We then consider Resources as material passports made of the track and trace of all nodes - Events, Agents and Processes - they descend from. The material passport is an authenticated graph structure: in figure 4 the Resources on the right side have an UID which is the aggregation of all UIDs of elements leading to their existence.

The integrity of the material passport can be verified by recalculating the UID aggregation and see it matches the signed one attached to the Resource. In case one or more UIDs are wrong or missing, the Resource will not verify as valid. In brief is possible that:

- ① One or more Agents may interact to create material passports

- ② A material passport is the aggregation of all parent nodes
- ③ Anyone may verify the integrity and validity of any material passport
- ④ The verification of a material passport does not reveal the identity of Agents contained
- ⑤ One may export and import a material passport as a graph query

Reflow’s unlinkability of credentials and signatures satisfy the privacy requirement for the material passport, while the possibility to aggregate and link all the elements of its graph allow to group multiple signatures into a single compact one, without requiring any interaction with the previous signers. The material passport signature will then be the sum of all Agents, Processes and Events involved, created or consumed for it. The Reflow material passport is the authenticated, immutable and portable track record of all nodes connected in its graph: material passports can be signed on export and verified on import, which makes them reliably portable from one graph to another in a federated environment.

This paper makes four key contributions

- We describe the signature scheme underlying Reflow, including how key generation, signing and verification operate (Section 2). The scheme is an application of the BLS signature scheme [Boneh et al., 2018b] fitted with features to grant the unlinkability of signatures and to secure it against rogue-key attacks.
- We describe the credential scheme underlying Reflow, including how key generation, issuance, aggregation and verification of credentials operate (Section 3). The scheme is an application of the Coconut credential scheme [Sonnino et al., 2018] that is general purpose and can be scaled to a fully distributed issuance that is re-randomizable.
- We implement a Zencode scenario of Reflow to be executed on and off-chain by the Zenroom VM, complete with functions for public credential issuance, signature session creation and multi-party non-interactive signing (Section 4). We evaluate the performance and cost of this implementation on on-site and on-line platforms leveraging end-to-end encryption (Section 5).
- We implement an efficient end-to-end encryption scheme for the authentication of graph data structures as observed in the material passport use-case. The scheme allows to simplify the complexity of track-and-trace implementations by making it sufficient to proceed by one level of depth and verify the integrity of aggregated signatures at each step.

Notations and assumptions

We will adopt the following notations:

- \mathbb{F}_p is the prime finite field with p elements (i.e. of prime order p);
- E denotes the (additive) group of points of the curve BLS-383 [Scott, 2017] which can be described with the Weierstrass form $y^2 = x^3 + 16$;
- E_T represents instead the group of points of the twisted curve of BLS-383, with embedding degree $k = 12$. The order of this group is the same of that of E ;

We also require defining the notion of a cryptographic pairing. Basically it is a function $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$, where $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T are all groups of same order n , such that satisfies the following properties:

- i. *Bilinearity*, i.e. given $P_1, Q_1 \in \mathbb{G}_1$ and $P_2, Q_2 \in \mathbb{G}_2$, we have

$$e(P_1 + Q_1, P_2) = e(P_1, P_2) \cdot e(Q_1, P_2)$$

$$e(P_1, P_2 + Q_2) = e(P_1, P_2) \cdot e(P_1, Q_2)$$
- ii. *Non-degeneracy*, meaning that for all $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$, $e(g_1, g_2) \neq 1_{\mathbb{G}_T}$, the identity element of the group \mathbb{G}_T ;
- iii. *Efficiency*, so that the map e is easy to compute;
- iv. $\mathbb{G}_1 \neq \mathbb{G}_2$, and moreover, that there exist no efficient homomorphism between \mathbb{G}_1 and \mathbb{G}_2 .

For the purpose of our protocol we will have $\mathbb{G}_1 = E_T$ and $\mathbb{G}_2 = E$, and $\mathbb{G}_T \subset \mathbb{F}_{p^{12}}$ is the subgroup containing the n -th roots of unity, where n is the order of the groups E and E_T . Instead $e : E_T \times E \rightarrow \mathbb{G}_T$ is the *Miller pairing*, which in our work is encoded as the method `miller(ECP2 P, ECP Q)`.

To conclude, the credential scheme in section 3 uses non-interactive zero-knowledge proofs (for short NIZK proofs) to assert knowledge and relations over discrete logarithm values. They will be represented using the notation introduced by Camenisch and Stadler [1997] as

$$\text{NIZK}\{(x, y, \dots) : \text{text statements about } x, y, \dots\}$$

2 Signature

A *BLS signature* is a signature scheme whose design exploits a cryptographic pairing. As for other well known algorithm such as ECDSA, it will work following these three main steps:

- **Key Generation phase.** For a user who wants to sign a message m , a secret key sk is randomly chosen uniformly in \mathbb{F}_n , where n is the order of the groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$. The corresponding public key pk is the element $sk \cdot G_2 \in E_T$;
- **Signing phase.** The message m is first hashed into the point $U \in E$, which in our scheme is done by the method `hashtopoint`; the related signature is then given by $\sigma = sk \cdot U$;
- **Verification phase.** For an other user that wants to verify the authenticity and the integrity of the message m , it needs to

1. parse m, pk and σ
2. hash the message m into the point U and then check if the following identity holds,

$$e(pk, U) = e(G_2, \sigma)$$

If verification passes it means that σ is a valid signature for m and the protocol ends without errors.

Proof of the verification algorithm: By using the definitions of the elements involved and exploiting the property of the pairing e we have

$$\begin{aligned} e(pk, U) &= e(sk \cdot G_2, U) \\ &= e(G_2, U)^{sk} \\ &= e(G_2, sk \cdot U) \\ &= e(G_2, \sigma) \end{aligned}$$

□

BLS signatures present some interesting features. For instance, the length of the output σ competes to those obtained by ECDSA and similar algorithms; in our specific case, by using BLS-383 [Scott, 2017], it will be 32 Bytes long, which is typically a standard nowadays. Then, since this curve is also pairing-friendly, meaning that (with the assumption made on e) signature and verification are obtained in very short time. Moreover, BLS supports also aggregation, that is the ability to aggregate a collection of multiple signatures σ_i (each one related to a different message m_i) into a singular new object σ , that can be validated using the respective public keys pk_i in a suitable way. This is possible thanks to the fact that $\sigma_i \in G_1 \forall i$, giving to the algorithm an homomorphic property. We will show now how this last feature can be attained in the context of a multi-party computation using the same message m but different participants.

Session Generation

After the key generation step we introduce a new phase called **session generation**, where the signature is initialized; anyone willing to start a signing session on a message m will create:

1. a random r and its corresponding point $R = r \cdot G_2$
2. the sum of R and all pk supposed to participate to the signature such as $P = R + \sum_i pk_i$
3. the unique identifier UID of the session calculated as hash to point of the message m , such as $U = H(m) \in E$, where H is a combination of a cryptographic hash function (treated as a random oracle) together with an encoding into elliptic curve points procedure
4. the first layer of the signature $\sigma \leftarrow r \cdot U$, later to be summed with all other signatures in a multi-party computation setup resulting in the final signature as $\sigma \leftarrow r \cdot U + \sum_i sk_i \cdot U$

5. the array of unique fingerprints ζ_i of each signature resulting from the credential authentication (see section 3)

After this phase is terminated, every participants involved in the session start their own signing phase during the session, producing (from the same message m) their respective σ_i 's. The final signature σ is then computed in this way: first of all let us call $\sigma_0 = r \cdot U$, then supposing k participants have already aggregated their σ_i , obtaining a partial signature S_k , the $(k + 1)$ -th one will compute

$$S_{k+1} = S_k + \sigma_k = \sum_{i=0}^{k+1} \sigma_i$$

Finally, the resulting output will be $\sigma = S_N$, where N is the total number of the signers of the session. In order to verify that σ is valid we compute $P = R + \sum_{i=0}^N pk_i$, where $R = r \cdot G_2$, working as a public key with respect to the nonce r , which instead is kept secret. Verification is then performed by checking if the following identity holds

$$e(P, U) = e(G_2, \sigma)$$

If verification passes without errors it means that σ is a valid aggregated signature of m .

Proof. By recalling that $\sigma = r \cdot U + \sum_{i=1}^N \sigma_i$, $P = R + \sum_{i=1}^N pk_i$, by using the property of the pairing e we have

$$\begin{aligned} e(P, U) &= e\left(R + \sum_{i=1}^N pk_i, U\right) \\ &= e\left(\left(r + \sum_{i=1}^N sk_i\right) \cdot G_2, U\right) \\ &= e(G_2, U)^{r + \sum_{i=1}^N sk_i} \\ &= e\left(G_2, \left(r + \sum_{i=1}^N sk_i\right) \cdot U\right) \\ &= e\left(G_2, r \cdot U + \sum_{i=1}^N \sigma_i\right) \\ &= e(G_2, \sigma) \end{aligned}$$

□

We conclude this section with a final consideration on this feature. We recall that in the generation of the aggregated signature σ we used as a starting point the variable $\sigma_0 = r \cdot U$, but in literature it is also common to find instead simply the base point G_2 . The choice of randomizing it (providing that the random number generator acts as an oracle) helps in preventing replay attacks, since the signature generated by the process is linked to the session in which is produced, for if an attacker managed to get some information from σ , it would be difficult to use it in order to forge new signatures.

3 Credential

Following the guidelines of Coconut, the credentials issuing scheme works as follows:

1. the issuer generates its own keypair (s_k, v_k) , where $s_k = (x, y) \in \mathbb{Z}^2$ is the pair of secret scalars (the signing key) and $v_k = (\alpha, \beta) = (x \cdot G_2, y \cdot G_2)$ is the verifying key, made by the related pair of public points over E_T ;
2. the user i , with its respective keys (sk_i, PK_i) make a credential request on its secret attribute $ck_i \in \mathbb{Z}$ to the issuer, represented by λ which contains a zero-knowledge proof π_s of the authenticity of user i ;
3. the issuer, after having received λ , verifies the proof π_s at its inside, and if it passes, then releases to user i a credential $\tilde{\sigma}$ signed used its own key sk .

Step 1. is self-explanatory. Steps 2. and 3. require a bit more effort, in fact in order to build a valid request λ , and so also a valid proof π_s , first of all the user must produce an hash digest for the attribute ck_i , that we call h , then computes two more variables c and s_h defined as

$$\begin{aligned} c &= r \cdot G_1 + h \cdot HS \\ s_h &= (a, b) = (k \cdot G_1, k \cdot \gamma + h \cdot c) \end{aligned}$$

where r and k are fresh randomly generated integers, HS is an hard-encoded point on the curve E , and $\gamma = ck_i \cdot G_1$. These two variables are alleged in the credential request λ produced in `prepare_blind_sign` and are needed to the verifier to assure the authenticity of the user through the proof π_s , which requires as input h, k, r, c . The Non-Interactive Zero Knowledge proof (for short NIZK proof) π_s generated by the function `blind_sign` is computed as follows:

- **Randomization phase.** Three new nonces $w_h, w_k, w_r \in \mathbb{Z}$ are generated, each one related to the input variables h, k, r respectively as we will show soon;
- **Challenge phase.** The protocol creates three commitment values, namely A_w, B_w, C_w defined as follows

$$\begin{aligned} A_w &= w_k \cdot G_1 \\ B_w &= w_k \cdot \gamma + w_h \cdot c \\ C_w &= w_r \cdot G_1 + w_h \cdot HS \end{aligned}$$

Then these variables are used as input of a function φ producing an integer $c_h = \varphi(\{c, A_w, B_w, C_w\})$;

- **Response phase.** In order that the proof can be verified the protocol generates three more variables which are alleged inside the proof itself and link the nonces w_h, w_k, w_r with h, k, r , i.e:

$$\begin{aligned} r_h &= w_h - c_h h \\ r_k &= w_k - c_h k \\ r_r &= w_r - c_h r \end{aligned}$$

So basically the proof π_s contains the three response variables r_h, r_k, r_r and also the commitment value c_h , that can be used for a predicate ϕ which is true when computed on h . Once the verifier receives the request λ , in order to check if the proof is valid it should be able to reconstruct A_w, B_w, C_w by doing these computations,

$$\begin{aligned} \hat{A}_w &= c_h \cdot a + r_k \cdot G_1 \\ \hat{B}_w &= c_h \cdot b + r_k \cdot \gamma + r_h \cdot c \\ \hat{C}_w &= c_h \cdot c + r_r \cdot G_1 + r_h \cdot HS \end{aligned}$$

If the request is correct, then we will have that

$$\varphi(\{c, \hat{A}_w, \hat{B}_w, \hat{C}_w\}) = \varphi(\{c, A_w, B_w, C_w\}) = c_h \quad (1)$$

and verification is thus complete, meaning that the verifier has right to believe that the prover actually owns the secret attribute ck_i associated to the public variable γ and that consequently has produced a valid commitment c and (El-Gamal) encryption s ; in other words,

$$\begin{aligned} \pi_s &= \text{NIZK}\{(ck_i, h, r, k) : \\ &\quad \gamma = ck_i \cdot G_1, \\ &\quad c = r \cdot G_1 + h \cdot HS, \\ &\quad s_h = (k \cdot G_1, k \cdot \gamma + h \cdot c), \\ &\quad \phi(h) = 1\} \end{aligned}$$

At this point the user will now have a blind credential $\tilde{\sigma} = (c, \tilde{a}, \tilde{b})$ issued by the authority, where

$$\begin{aligned} \tilde{a} &= y \cdot a \\ \tilde{b} &= x \cdot c + y \cdot b \end{aligned}$$

The user then will have to un-blind it using its secret credential key, obtaining $\sigma_{ck} = (c, s) = (c, \tilde{b} - ck_i(\tilde{a}))$, which will use to prove its identity when signing a message. The procedure is similar to the one seen before with some extra details:

- **Setup.** As for the the BLS signature, an elliptic point U , associated to the hash of the message to sign, is required as an Unique Identifier (UID) for the signing session;
- **Credential proving.** The user produces two cryptographic objects θ (containing a new proof π_v) and ζ (which is unequivocally associated to U) through `prove_cred_uid`, taking as input its own credential σ , the related secret attribute c_k , the authority public key $v_k = (\alpha, \beta)$ and the session point U .

The new objects θ and ζ are derived as follows:

- as before the user hashes ck into h , and this time generates two random values r and r' ;
- next, it randomizes its credential σ_{ck} into $\sigma'_{ck} = (c', s') = (r' \cdot c, r' \cdot s)$ and then computes two elliptic curve points κ and ν as

$$\begin{aligned} \kappa &= \alpha + h \cdot \beta + r \cdot G_2 \\ \nu &= r \cdot c' \end{aligned}$$

- finally, θ will be the t-tuple $(\kappa, \nu, \pi_v, \phi')$, where π_v is a valid zero knowledge proof of the following form

$$\begin{aligned} \pi_v = \text{NIZK}\{(h, r) : \\ \kappa = \alpha + h \cdot \beta + r \cdot G_2 \\ \nu = r \cdot c' \\ \phi'(h) = 1\} \end{aligned}$$

with ϕ' being a predicate which is true on h ;

- ζ will be instead the elliptic curve point obtained as $h \cdot U \in E$.

Building the proof π_v requires similar steps as seen for π_s , in fact we create three commitment values A_w, B_w, C_w defined as

$$\begin{aligned} A_w &= \alpha + w_h \cdot \beta + w_r \cdot G_2 \\ B_w &= w_r \cdot c' \\ C_w &= w_h \cdot U \end{aligned}$$

where w_h, w_r are fresh generated nonces; then we set the challenge as the computation of $c_h = \varphi(\{\alpha, \beta, A_w, B_w, C_w\})$ with the related responses

$$\begin{aligned} r_h &= w_h - hc \\ r_r &= w_r - rc \end{aligned}$$

The values of r_h and r_r are stored inside π_v which will be then sent through θ (together with ζ) from the prover to the verifier. In order to check that the user has legitimately generated the proof and at the same time is the owner of the credential the following steps must be made:

1. extracting κ, ν, π_v (which is (r_h, r_r, c_h)) from θ ;
2. reconstructing the commitments A_w, B_w, C_w as

$$\begin{aligned} \widehat{A}_w &= c_h \cdot \kappa + r_r \cdot G_2 + (1 - c_h) \cdot \alpha + r_h \cdot \beta \\ \widehat{B}_w &= r_r \cdot c' + c_h \cdot \nu \\ \widehat{C}_w &= r_h \cdot U + c_h \cdot \zeta \end{aligned}$$

3. checking either that

$$\begin{aligned} \varphi(\{\alpha, \beta, \widehat{A}_w, \widehat{B}_w, \widehat{C}_w\}) = \\ \varphi(\{\alpha, \beta, A_w, B_w, C_w\}) = c_h, \quad (2) \end{aligned}$$

that $c' \neq O$, the point at infinity, and that

$$e(\kappa, c') = e(G_2, s' + \nu) \quad (3)$$

Actually the predicate ϕ in the definition of π_v can be thought as performing steps 2. and 3. and, if any of these fails the protocol will abort returning a failure, otherwise verification passes and the user can finally produce the signature.

Proof of the verification algorithms.

We now show for the proof π_s that actually by using the responses r_h, r_k and r_r , together with c_h and the other parameters inside λ , i.e. $s = (a, b), c$ and γ , using also the hard coded point HS , it is possible to reconstruct the commitments A_w, B_w, C_w :

$$\begin{aligned} \widehat{A}_w &= c_h \cdot a + r_k \cdot G_1 = c_h k \cdot G_1 + r_k \cdot G_1 \\ &= (c_h k + r_k) \cdot G_1 = (c_h k + w_k - c_h k) \cdot G_1 \\ &= w_k \cdot G_1 = A_w \end{aligned}$$

$$\begin{aligned} \widehat{B}_w &= c_h \cdot b + r_k \cdot \gamma + r_h \cdot c \\ &= c_h \cdot (k \cdot \gamma + h \cdot c) + r_k \cdot \gamma + r_h \cdot c \\ &= (c_h k + r_k) \cdot \gamma + (c_h h + r_h) \cdot c \\ &= (c_h k + w_k - c_h k) \cdot \gamma + (c_h h + w_h - c_h h) \cdot c \\ &= w_k \cdot \gamma + w_h \cdot c = B_w \end{aligned}$$

$$\begin{aligned} \widehat{C}_w &= c_h \cdot c + r_r \cdot G_1 + r_h \cdot HS \\ &= c_h \cdot (r \cdot G_1 + h \cdot HS) + r_r \cdot G_1 + r_h \cdot HS \\ &= (c_h r + r_r) \cdot G_1 + (c_h h + r_h) \cdot HS \\ &= (c_h r + w_r - c_h r) \cdot G_1 + (c_h h + w_h - c_h h) \cdot HS \\ &= w_r \cdot G_1 + w_h \cdot HS = C_w \end{aligned}$$

Regarding the second proof π_v , we have to prove both the identities (2) and (3) hold. We will focus only on the latter since the former requires a similar approach on what we have done for π_s , but with different parameters involved (κ, ν , etc.). The left-hand side of the relation can be expressed as

$$\begin{aligned} e(\kappa, c') &= e(\alpha + h \cdot \beta + r \cdot G_2, c') \\ &= e(x \cdot G_2 + hy \cdot G_2 + r \cdot G_2, \tilde{r} \cdot G_1) \\ &= e((x + hy + r) \cdot G_2, \tilde{r} \cdot G_1) \\ &= e(G_2, G_1)^{(x+hy+r)\tilde{r}} \end{aligned}$$

using the substitution $c' = \tilde{r} \cdot G_1$, with $\tilde{r} \in \mathbb{F}_p$ since we know that $c' \in E$. For the right-hand side we have instead

$$\begin{aligned} e(G_2, s' + \nu) &= e(G_2, r' \cdot s + r \cdot c') \\ &= e(G_2, r'(\tilde{b} - ck_i y(\tilde{a})) + r \cdot c') \\ &= e(G_2, r'(x \cdot c + y \cdot b - ck_i y \cdot a) + r \cdot c') \\ &= e(G_2, r'(x \cdot c + y \cdot b - ck_i y \cdot a + r \cdot c)) \end{aligned}$$

The second argument of the pairing can be rewritten as

$$\begin{aligned} r'(x \cdot c + y \cdot b - ck_i y \cdot a + r \cdot c) &= \\ r'(x \cdot c + y(k \cdot \gamma + h \cdot c) - (ck_i)yk \cdot G_1 + r \cdot c) &= \\ r'(x \cdot c + yh \cdot c + yk(ck_i) \cdot G_1 - yk(ck_i) \cdot G_1 + r \cdot c) &= \\ r'(x + yh + r) \cdot c & \end{aligned}$$

So, at the end

$$\begin{aligned} e(G_2, s' + \nu) &= e(G_2, r'(x \cdot c + y \cdot b - dy \cdot a + r \cdot c)) \\ &= e(G_2, r'(x + yh + r) \cdot c) \\ &= e(G_2, (x + hy + r)\tilde{r} \cdot G_1) \\ &= e(G_2, G_1)^{(x+hy+r)\tilde{r}} \end{aligned}$$

and (3) is finally proved. \square

Security considerations.

As mentioned in Coconut [Sonnino et al., 2018], BLS signatures and the proof system obtained with credentials are considered secure by assuming the existence of random oracles [Koblitz and Menezes, 2015], together with the decisional Diffie-Hellman Problem (DDH) [Boneh, 1998], the external Diffie-Hellman Problem (XDH), and with the Lysyanskaya-Rivest-Sahai-Wol Problem (LRSW) [Lysyanskaya et al., 1999], which are connected to the Discrete Logarithm. In fact, under these assumptions, we have that our protocol satisfies unforgeability, blindness, and unlinkability.

Reserves can be made about the maturity of pairing-based elliptic curve cryptography despite various efforts to measure its security and design curve parameters that raise it, it is reasonable to consider this as a pioneering field of cryptography in contrast to well tested standards.

In addition to considerations on the maturity of EC, the future growth of quantum-computing technologies may be able to overcome the Discrete Logarithmic assumptions by qualitatively different computational means. Reflow then may be vulnerable to quantum-computing attacks, as well hard to patch, because the pairing-based design sits at its core with the adoption of ATE / Miller loop pairing of curves in twisted space, a practice that is not covered by research on quantum-proof algorithms and will eventually need more time to be addressed; however this is all speculative reasoning on what we can expect from the future.

The Reflow implementation we are presenting in this paper and that we have published as a Zenroom scenario ready to use is based on the BLS383 curve [Scott, 2017] that in the current implementation provided by the AMCL library has shown to pass all lab-tests regarding pairing properties, a positive result that is not shared with the slightly different BLS381-12 curve adopted by ETH2.0. Debating the choice of BLS381 is well beyond the scope of this paper, but is worth mentioning that our lab-tests have proved also the BLS461 curve to work in Reflow: it is based on a 461 bit prime and hence upgrades our implementation to 128 bit security [Barbulescu and Duquesne, 2019] against attacks looking for discrete logs on elliptic curves [Lim and Lee, 1997].

At last the complexity and flexibility of Reflow in its different applications, its optional use of fingerprint lists, multiple UIDs saved from aggregation and other features covering the different applications also represent a security risk in the technical integration phase. We believe that the adoption of Zenroom and the creation of a Zencode scenario addresses well this vulnerability by providing an easy to use integrated development environment (apiroom.net) and a test-bed for the design of different scenarios of application for Reflow that can be deployed correctly granting end-to-end encryption and data minimization according to privacy by design guidelines [Hoepman et al., 2019].

4 Implementation

In this section we illustrate our implementation of Reflow keygen, sign and verify operations outlining for each:

- the communication sequence diagram when necessary
- the Zencode statements to operate the transformation

Zencode is actual code executed inside the Zenroom VM, behind its implementation the algorithms follow very closely the mathematical formulation explained in this article and can be reviewed in the free and open source code published at Zenroom.org. The Reflow Zencode scenario implementation is contained inside the Zenroom source code.

To execute an example flow on one's own computer it is enough to visit the website ApiRoom.net and select the Reflow scenario examples, pressing play it will execute zencode and show input and outputs in JSON format.

Credential Setup

The Credential Setup sequence which has to be executed only once to enable participants to produce signatures. This sequence is briefly illustrated with a diagram and it consists in the creation of keypair for both the Issuer, who will sign the participant credentials, and the participant who will request an issuer credential.

Listing 1: Issuer Keygen

```
Given I am 'The Issuer'
When I create the issuer key
Then print my 'keys'
```

Executed by the Zencode utterance:

When I create the issuer key

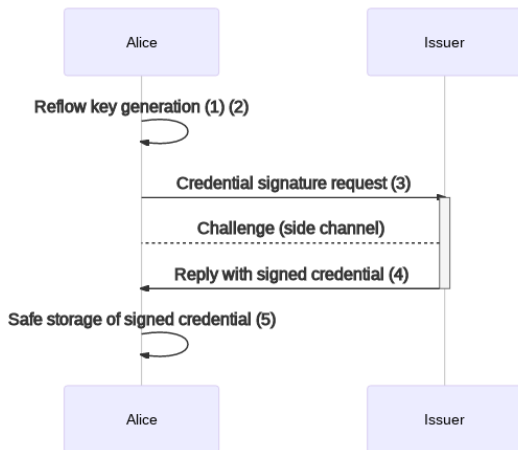
It will create a new *issuer key* that can be used to sign each new *credential request*. Its public member *.verify* should be public and know to anyone willing to verify the credentials of signers. It can be safely separated from the secret key and exported with:

Listing 2: Issuer verifier

```
Scenario credential: publish verifier
Given that I am known as 'The Issuer'
and I have my 'keys'
When I create the issuer public key
Then print my 'issuer public key'
```

Credential signature: Generate a credential request and have it signed by an Issuer, as well generate all the keys used to sign documents. This procedure will generate private keys that should not be communicated, as well BLS keys whose public section should be communicated and can be later aggregated to create a seal (signature session) and verify one.

Figure 5: Keygen process sequence diagram



will be executed by the Signer to aggregate one or more *credential signature* into the secret *keys* structure. The *credentials* consist of a σ_{ck} that should be stored locally and secretly as it will be later required to create the zero-knowledge proof components θ and ζ that authenticate the Reflow signature against rogue-key attacks.

Figure 5 illustrates how this process unfolds in a sequence of communication. An interactive exchange takes place between the Signer and the Issuer that verifies the possession of the secret Reflow key and signs a credential to witness this condition.

The public key of the Issuer which is used to sign the credential should have been public and known by the Signer at the beginning of the Keygen process.

This code is executed in multiple steps by the Zencode utterances:

- ① **When I create the reflow key and I create the credential key**
 will create inside the *keys* data structure two new secret keys: the *credential* and the *bls* keys both consisting of *BIG* integers. Secret keys are randomly chosen uniformly in \mathbb{F}_n , where n is the order of the group \mathbb{G}_1 .
- ② **When I create the credential request**
 will use the *secret keys* to create a new *credential request* consisting of a λ containing a zero-knowledge proof π_s of the authenticity of Signer i owning keys (sk_i, PK_i) and requesting the Issuer to sign its secret attribute $ck_i \in \mathbb{Z}$. The Issuer may request additional communications held on a *side channel* in order to establish proof-of-possession and the authenticity of the Signer, for instance a challenge session in which a random message is demonstrably signed by the secret credential key using an El-Gamal signature [ElGamal, 1985] and verified by the public key.
- ③ **Given I have a 'credential request' When I create the credential signature and I create the issuer public key**
 will be executed by the Issuer to sign the credential after the proof-of-possession challenge is positive. It consists of a $\tilde{\sigma}$ signed used the Issuer's own secret key sk after verifying the proof π_s inside the received λ .
- ④ **Given I have a 'credential signature' When I create the credentials**

4.1 Signature

The Reflow signature consists into a multi-signature that can be indicted by anyone who selects multiple participants. It requires no interaction between participants and it results into what we call a *Seal* collecting all the signatures.

Both the credential to sign and the signatures in the reflow seal are non traceable. The credentials are zero-knowledge proof of possessions bound to the seal identifier; the public keys of participants are blinded and known only by the one who indicts the signature, who may or may not be among participants.

The signatures are non-interactive and made in a multi-party computation setup: any participant with a credential to sign can do so without interacting with others, just by adding the reflow signature to the reflow seal.

Every signature in a seal produces a *fingerprint* unique to the seal and the signer consisting in a ζ that is also non traceable. Using fingerprints a check may be enforced to avoid double signing and invalidate the reflow seal.

When a seal is created a session is opened for multiple participants to add their signature, anyone opening such a process needs to have an *array of public keys* of all desired participants.

Listing 3: Create reflow seal

```
Scenario reflow: create a new signing session (seal)
Given I have a 'reflow public key array' named 'public keys'
and I have a 'string dictionary' named 'Event'
When I aggregate the reflow public key from array 'public keys'
and I create the reflow identity of 'Event'
and I create the reflow seal with identity 'reflow identity'
Then print the 'reflow seal'
```

When a seal is signed there are two different steps to be performed: one by the participant publishing the signature and one by anyone aggregating the signature to the seal. This is necessary since we are in a multi-party computation session where each step forward in the process can be performed independently and without requiring anyone to disclose any secret.

Listing 4: Create a reflow signature

```
Scenario reflow: participant signs a seal by producing a signature
Given I am 'Alice'
and I have the 'credentials'
and I have the 'keys'
and I have a 'reflow seal'
and I have a 'issuer public key' from 'The Issuer'
When I create the reflow signature
Then print the 'reflow signature'
```

Listing 5: Add a signature to a seal

```
Scenario reflow
Given I have a 'reflow seal'
and I have a 'issuer public key' in 'The Authority'
and I have a 'reflow signature'
When I aggregate all the issuer public keys
and I verify the reflow signature credential
and I check the reflow signature fingerprint is new
and I add the reflow fingerprint to the reflow seal
and I add the reflow signature to the reflow seal
Then print the 'reflow seal'
```

The session may be considered terminated at any moment by anyone: the seal can be validated in any moment as-is and it will simply return false validation if signatures are wrong or incomplete.

Listing 6: Compare reflow identities

```

Scenario reflow: compare two different objects using reflow IDs
Given I have a 'reflow seal'
Given I have a 'string dictionary' named Event
When I create the reflow identity of 'Event'
When I rename the 'reflow identity' to 'EventIdentity'
When I verify 'EventIdentity' is equal to 'identity' in 'reflow seal'
Then print the string 'The seal really belongs to this Event'

```

Listing 7: Verify a seal

```

Scenario reflow: verify a seal contains all valid signatures
Given I have a 'reflow seal'
When I verify the reflow seal is valid
Then print the string 'SUCCESS'
and print the 'reflow seal'

```

At this point the seal may be marked as closed as the session is completed successfully and all correct signatures have been collected without duplicates. In case this optional status will be attributed to the seal then is also possible to remove the *fingerprints* array from it by reducing its size.

4.2 Material Passport

The Reflow material passport consists in a REA model track-and-trace [Laurier and Poels, 2012] listing. It aggregates Reflow seals by means of collecting their signed UID (unique identifier). The list of fingerprints in a material passport is not to be confused with that of a seal: it is the list of all UID being traced in one seal, which adds its own UID resulting from the hash to point of the new message.

The complexity of this implementation requires the adoption of a small ontology, the Valueflows vocabulary [Foster et al., 2017], that briefly distinguishes between roles of the nodes represented and organizes their interaction.

A seal may still be signed by one or more participants which are the *Agents* in the ValueFlows REA model. Agents may be signing a new *Event* or *Process* alone or together with more agents. All signatures will still be blinded and non-traceable to satisfy Reflow's privacy requirements, while the UID composing the material passport may be listed in a fingerprint section to authenticate the track-and-trace graph and perform deeper verification of all nodes.

Listing 8: Create a new material passport

```

Scenario reflow
Given I am 'Alice'
and I have the 'keys'
and I have the 'credentials'
and I have a 'issuer public key' in 'The Authority'
and I have a 'reflow identity'
and I have a 'string dictionary' named 'EconomicEvent'
When I create the material passport of 'EconomicEvent'
Then print the 'EconomicEvent'
and print the 'material passport'

```

What follows is the representation of the output, which shows how a material passport is formatted, keeping the same variable names used in the mathematical formulas and indicating the type of members ("ecp" stands for elliptic curve point and a 2 is added to indicate those on a twisted curve) and their size in bytes.

Listing 9: Example material passport data

```

material_passport = {
  proof = {
    kappa = ecp2[192] ,
    nu = ecp[49] ,

```

```

    pi_v = {
        c = int[32] ,
        rm = int[32] ,
        rr = int[32]
    },
    sigma_prime = {
        h_prime = ecp[49] ,
        s_prime = ecp[49]
    }
},
seal = {
    SM = ecp[49] ,
    fingerprints = [1] <table 1>,
    identity = ecp[49] ,
    verifier = ecp2[192]
},
zeta = ecp[49]
}

```

As hinted by the optional presence of a list of fingerprints, as well optional record of the agent's credential *Proof* inside the material passport, this implementation is extremely flexible and opens up a range of possible solutions in the field of privacy-enhanced graph authentication, allowing import and export operations between federated graph data bases, also considering each data silo may be an *Issuer* publicly identified and granting agent credentials.

The verification of a material passport implicitly verifies that the the hash to point of the objects's contents result in the same *reflow identity*, then executes a cryptographic verification that the given material passport is the valid signature of that identity:

Listing 10: Verify a material passport

Scenario reflow

```

Given I have a 'string dictionary' named 'EconomicEvent'
and I have a 'issuer public key' in 'The Authority'
and I have a 'material passport'
When I verify the material passport of 'EconomicEvent'
Then print the string 'Valid Event material passport'

```

To avoid confusion is important to keep in mind that identities are UIDs and fingerprints are zetas.

The simple signature and verification of a single object is not enough to implement a material passport, we need to authenticate the *track and trace* of it eventually validating all the complex history of transformations that lead to the current object. In order to do so our implementation exploits the aggregation property of elliptic curve points by making a sum of all object identities present in the track and trace. Following the ValueFlows logic then, every Event or Process identity will result from the sum of all identities of previous Agents, Events and Processes. By recalculating this sum and checking its signature then one can authenticate the integrity of the track and trace.

For example, assuming we have an array of seal objects found in material passports of a track and trace result of one level depth:

Listing 11: Example material passport data

```

Seals = [2] { {
    SM = ecp[49] ,
    fingerprints = [1] { ecp[49] },
    identity = ecp[49] ,
    verifier = ecp2[192]
}, {
    SM = ecp[49] ,
    fingerprints = [1] { ecp[49] },
    identity = ecp[49] ,
    verifier = ecp2[192]
}
}

```



```
} }
```

Then we can calculate a new identity using aggregation (sum value) with Zencode statements:

Listing 12: Aggregate reflow seal identity

```
Given I have a 'reflow seal array' named 'Seals'  
When I create the sum value 'identity' for dictionaries in 'Seals'  
Then print the 'sum value'
```

Which will output an elliptic curve point corresponding to the sum of all identity elements found in the "Seals" array. The fact that the verification happens to a sum that cascades down aggregating further depth level of parent objects allows to modularize the track-and-trace implementation making its depth configurable.

5 Evaluation

The goal of this section is to produce benchmarks of the implementation of the cryptographic flow, in realistic conditions of use, so in a way that is similar to how a software solution would use the software.

Our approach was opposite to testing single algorithms in a sandbox or a profiler, instead we tested Zenroom scripts, written in Zencode, that include also loading and parsing input data from streams or file system, as well as producing output as a deterministically formed JSON object.

Platforms

The three target platforms used for the benchmarks were:

- **X86-64** on Intel(R) Core(TM) i5-5300U CPU @ 2.30GHz, running Ubuntu 18.04 (64 bit)
- **ARM 32bit** on **Raspberry Pi 4** board, running Raspberry Pi OS (32bit, kernel version: 5.10)
- **ARM 32bit** on **Raspberry Pi 0** board, running Raspberry Pi OS (32bit, kernel version: 5.10)

The the X86-64 machine runs on a Lenovo X250 laptop, with 8GB of RAM. The Raspberry 0 and 4 have been chosen as benchmarking platforms because of their similarities to, respectively, very low-cost IoT devices (5 USD) and very low cost mobile phones (sub 100 USD).

Builds

Being a self-contained application written in C, Zenroom can be built for several CPU architectures and operative systems, both as command line interface (CLI) application and as a library. The configurations we chose for this benchmark are:

- Two binary Zenroom CLI, compiled using the GCC tool-chain in native 32bit ELF binaries, once for the X86-64 platform and once for the ARMv7-32bit. We refer to these builds as **Zenroom CLI**.
- One mixed library, trans-compiled to WASM using the Emscripten tool-chain [Zakai, 2011], then built into an NPM package along with a JavaScript wrapper. After trans-compilation, the WASM library is converted to base64 and embedded inside the JavaScript wrapper, which unpacks it at run-time. The library runs in browsers (using the native WASM support, currently present in Chromium/Chrome, Firefox and Safari), as well in node JS based application. We refer to this build as **Zenroom WASM**.

CLI and WASM use cases

The use case for Zenroom CLI is typically a server-side application or micro-service. The WASM library can run in the browser, in the front end of a web application or in a

server-side application (running on node JS) or a mobile application.

Zenroom can also be built as a native library for Android and iOS (on ARM 32bit, ARM 32bit or X86), but no benchmarks of the Reflow cryptographic flow have been performed using Zenroom as a native library. In previous benchmarks, we noted similar performances for the native CLI and the native library versions of Zenroom, we therefore assume that performances would be comparable between the CLI builds we are using and the native library builds.

Testing

Benchmarking of Zenroom CLI and Zenroom WASM required using two different tools. Both tools performed chained execution of Zencode scripts that would simulate the whole cryptographic flow. Furthermore, the tool allowed to configure the amount of participants in the flow and the repeat the run of selected scripts in order to gather more data.

- The testing of Zenroom CLI was executed using a single, self-contained [bash script](#). The script outputs the duration of the execution of each Zenroom scripts, along with the memory usage and size of output data, in a CSV formatted summary. The script also saves to files both the data and scripts for quality control reference.
- The testing of Zenroom WASM required building a [JavaScript application](#), running in node JS. The application outputs the duration of the execution of each Zenroom scripts, along with the memory usage, it calculate averages and returns output in a JSON formatted summary.

Benchmarks for the cryptographic flow were run on both the CLI and the WASM builds of Zenroom, for each platform, for a total of six different data collections.

Scripts description

Our benchmarks provide measurements for the all the steps in the signature flow: for each of them we reported the time of execution (expressed in seconds) and the size of output (expressed in bytes) in different conditions with a progressive number of participants (5, 10, 50, 100, 1000 signatures).

The Zencode scripts used in the flow and for benchmarking can be divided in three categories:

- Scripts with **variable execution times** executed by anyone, marked with (A): they include the creation of the Reflow seal (Session start), the aggregation of the signatures (Collect sign) and the verification of the signatures on Reflow seal. Their duration is proportional to the amount of participants and can differ greatly. These scripts are the most calculation intensive and expected

to typically run server-side. The (A) stands for "Anyone" as they don't need an identity or a keypair to run. These scripts are typically expected to run server side.

- Scripts executed by each participant, marked with (P): they include participant's setup and signing, their duration is not correlated with the amount of participants. These scripts are expected to typically run in the browser or on mobile devices.
- Scripts executed by the issuer, marked with (I): they include issuer's setup and signing, their duration is not correlated with the amount of participants. These scripts are expected to typically run server-side.
- Every script of the groups (P) and (I) have acceptable performances on any platform when using Zenroom compiled as CLI.
- With a few exceptions, we find a similar result when running scripts of the (P) and (I) groups, for Zenroom compiled in WASM. With this setup, only three scripts run over the 1 second limit, only on Raspberry Pi 0.
- The scripts in the group (A), when using Zenroom CLI, can be executed within one second for an amount of participants between 100 and 500 on the X86 machine and for a similar amount on a Raspberry Pi 4, while on a Raspberry Pi 0 the limit is between 10 and 50 participants.

Each script was executed 50 times, on each platform and each configuration, and the average execution time was extracted.

Findings

- The execution time of the scripts in the (A) group, grows linearly with the amount of participants.
- On a X86-64 machine, the benchmarks execute on Zenroom CLI have a comparable duration with their counterparts run on WASM. On the other hand, the execution times differ greatly on the ARM 32bit based Raspberry Pi 0 and 4 machine, when comparing Zenroom CLI and Zenroom WASM executions. While investigating the cause of the differences is beyond the purpose of this paper, we speculate this difference signals different levels of optimization, between the of WASM interpreters built for X86-64 on one hand, and ARM 32bit on the other.
- The most numerically outstanding benchmark is the execution time of the issuer's keygen script on Raspberry Pi 0, in comparison with the other platforms. Once more, investigating the reason for this discrepancy is beyond the scope of this paper: we do anyway speculate that the ECP2 pairing performed in the script, may justify the large difference.

Benchmark conclusions

The ultimate goal of this benchmarks is to assess what parts of cryptographic flow are suitable for execution on each platform, both in terms of execution times as well as RAM usage. Concerning the **RAM usage**, the highest value recorded is in the 5 megabytes range, which we consider to be well within the acceptable range for any of the platforms used in the benchmarks. Future analysis will investigate the possibility to execute the less resource intense parts of the flow, on ultra low power chips such as the **Cortex M4**, for which a port of Zenroom is available.

Concerning the **execution time**, if we choose that the longest execution time acceptable for any script of the flow, is up to 1 second, we find that:

- The scripts in the group (A), when using Zenroom WASM, can be executed within one second, for an amount of participants between 100 and 500 on the X86 machine and up to 100 on a Raspberry Pi 4, while on a Raspberry Pi 0 the limit is between 2 and 5 participants.


6 Benchmarks

Following are the benchmarks of the benchmarks for the (A) group of scripts (whose execution times change based on the amount of participants), running in CLI binaries, grouped by platform.

Zenroom CLI X86-64

Table 1: Execution timings on **CLI X64** in *seconds per participant*


<i>S/P</i>	2	10	50	100	1000	5000
(A) Session start	0.0188	0.0253	0.0410	0.0623	0.1056	0.5039
(A) Collect sign	0.0729	0.0730	0.1031	0.1373	0.7435	3.6061
(A) Verify sign	0.0405	0.0515	0.0996	0.1619	1.3974	7.0460



Zenroom CLI Raspberry Pi 4

Table 2: Execution times on **CLI Raspberry 4** in *seconds per participant*


<i>S/P</i>	2	10	50	100	1000
(A) Session start	0.0516	0.0602	0.0674	0.0750	0.2762
(A) Collect sign	0.2087	0.2273	0.3031	0.3942	2.1849
(A) Verify sign	0.1049	0.1393	0.2978	0.4831	4.0688



Zenroom CLI Raspberry Pi 0

Table 3: Execution times for Zenroom CLI Raspberry Pi 0 in seconds per participant

<i>S/P</i>	2	10	50	100	1000
(A) Session start	0.2760	0.2936	0.3413	0.3725	1.3097
(A) Collect sign	1.0439	1.1458	1.3688	1.7669	9.6403
(A) Verify sign	0.5753	0.7199	1.3381	2.1971	17.8956



Zenroom CLI - all platforms


Following are the benchmarks of the benchmarks for the (P) and the (I) groups of scripts (whose execution times don't change based on the amount of participants), with a comparison of all the platforms, running in CLI binaries.

Table 4: Execution timings on CLI of scripts in *seconds per platform*

<i>script/platform</i>	(P) Keygen	(P) Request	(P) Pub-Key	(P) Aggr. Cred.	(P) Sign Session
X86-X64	0.0137	0.0391	0.0196	0.0219	0.0521
Raspberry Pi 4	0.0373	0.0883	0.0475	0.0579	0.1456
Raspberry Pi 0	0.2492	0.5165	0.3017	0.3469	0.6603

Table 5: Execution timings on CLI of scripts in *seconds per platform*

<i>script/arch</i>	(I) Keygen	(I) Pub-Key	(I) Sign Req.
X64	0.0129	0.0223	0.0524
Raspberry Pi 4	0.0404	0.0597	0.1405
Raspberry Pi 0	0.2467	0.3766	0.7687

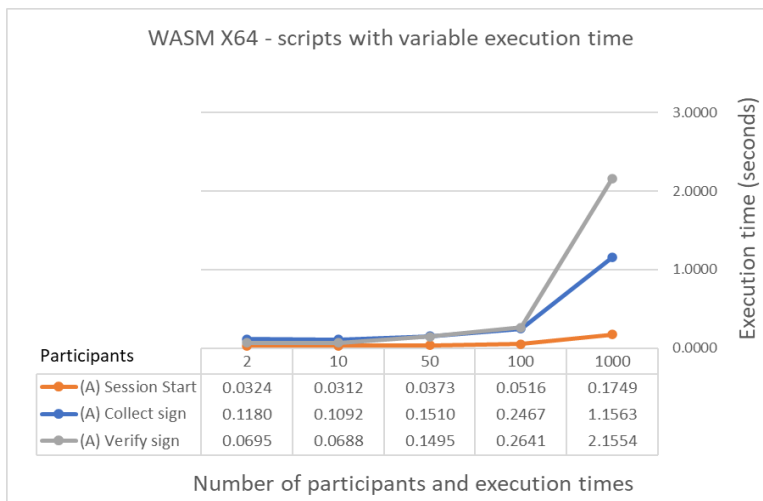


Following are the benchmarks of the benchmarks for the (A) group of scripts (whose execution times change based on the amount of participants), running in WASM libraries, grouped by platform.

Zenroom WASM X86-64

Table 6: Execution timings on WASM X64 in *seconds per participant*


<i>S/P</i>	2	10	50	100	1000
(A) Session start	0.0324	0.0312	0.0373	0.0516	0.1749
(A) Collect sign	0.1180	0.1092	0.1510	0.2467	1.1563
(A) Verify sign	0.0695	0.0688	0.1495	0.2641	2.1554



Zenroom WASM Raspberry Pi 4

Table 7: Execution timings on WASM Raspberry Pi 4 in seconds per participant


<i>S/P</i>	2	10	50	100	1000
(A) Session start	0.1065	0.1092	0.1269	0.1459	0.5432
(A) Collect sign	0.3897	0.4170	0.5741	0.7654	4.3594
(A) Verify sign	0.1880	0.2510	0.5621	0.9494	8.1078



Zenroom WASM Raspberry Pi 0

Table 8: Execution timings on WASM Raspberry Pi 0 in seconds per participant

<i>S/P</i>	2	10	50	100	1000
(A) Session start	0.8688	0.8711	1.0392	1.2257	4.7067
(A) Collect sign	2.7272	2.8665	3.9580	5.1851	29.0912
(A) Verify sign	1.3818	1.7820	3.8503	6.3604	53.8938



Zenroom WASM - all platforms


Following are the benchmarks of the benchmarks for the (P) and the (I) groups of scripts (whose executions don't change based on the amount of participants), with a comparison of all the platforms, running in CLI binaries.

Table 9: Execution times on **Zenroom WASM** of scripts in *seconds per platform*

<i>script/platform</i>	(P) Keygen	(P) Request	(P) Pub-Key	(P) Aggr. Cred.	(P) Sign Session
X86-X64	0.0213	0.0490	0.0262	0.0304	0.0745
Raspberry Pi 4	0.169	0.3365	0.2095	0.173	0.2705
Raspberry Pi 0	0.6275	1.291	0.7185	0.811	1.9435

Table 10: Execution times on **Zenroom WASM** of scripts in *seconds per platform*

<i>script/platform</i>	(I) Keygen	(I) Pub-Key	(I) Sign Req.
X86-X64	0.0204	0.0317	0.0746
Raspberry Pi 4	1.01	0.276	0.6585
Raspberry Pi 0	9.958	1.101	2.6355



RAM usage

The benchmarking of the RAM has been performed on Zenroom CLI only. The RAM usage showed negligible differences among the X86 and the arm builds. The tables show the results of the X86 benchmarks.

Table 11: Execution timings on **RAM Usage** in *KiloBytes* per *participant*

<i>KB/Participant</i>	2	10	50	100	1000	5000
(A) Session start	500 KB	507 KB	534 KB	589 KB	1289 KB	4543 KB
(A) Collect sign	520 KB	548 KB	665 KB	817 KB	1651 KB	4368 KB
(A) Verify sign	507 KB	532 KB	653 KB	801 KB	1660 KB	4695 KB

Table 12: Execution times on **RAM Usage** in *KiloBytes*

(P) Keygen	(P) Request	(P) Pub-Key	(P) Aggr. Cred.	(P) Sign Session
498 KB	512 KB	484 KB	497 KB	541 KB

Table 13: Execution times on **RAM Usage** in *KiloBytes*

(I) Keygen	(I) Pub-Key	(I) Sign Req.
488 KB	491 KB	504 KB

7 Conclusion

This article and the referenced free and open source implementation made in Zencode successfully provides an easy, flexible and well performant way to realize anonymous signatures made by zero-knowledge proof credential holders and aggregated in multi-party computation decentralized environments.

The privacy-preserving features of Reflow signatures are protected from rogue-key attacks by enforcing credential issuance, at the price of requiring one or multiple authorities. From this approach it comes the feature of generating signature "fingerprints" that are contextual and become traceable only with the consensus of participants (for instance by replicating the signature). These fingerprints are also useful to avoid double-signing and to strengthen the track-and-trace features of the most advanced material passport use-case scenario.

The use-cases we observed haven't required the adoption of multiple authorities, however it is possible to enhance our scheme by porting a threshold credential issuance mechanism based on Lagrange interpolation as described in Coconut's paper [Sonnino et al., 2018].

Generally speaking the Reflow cryptographic system can be adapted to authenticate many sorts of graph data sets; this aspect goes well beyond the state of the art and its application may come useful to an upcoming generation of non-linear distributed ledger technologies.

7.1 Future directions

This cryptographic scheme is named after the Reflow project and currently follows its path to serve the use-case of circular economy projects for which the integrity of accounting, the portability of data and the privacy of users are of paramount importance. We are proud of the priority this cryptography research gives to use-cases dealing with environmental responsibility and will carry on this ethos through its future developments.

Standardisation is an important point in order to establish federated and decentralised use of Reflow across a variety of contexts, so we have high incentives in integrating this signature in W3C standards as well submitting it to the attention of the Object Management Group.

As hinted in the overview, the discourse around disposable identities is also very interesting and deserves further interaction as it may open up more privacy preserving possibilities of development in the field of health-care and transportation as well public services and accounting traceability for KYC/AML practices.

The possibility to aggregate Reflow signatures may also be exploited to serve need-to-know schemes where different actors may be required to sign different parts of a document without knowing the whole until its final publication, a scenario present in peer-review processes that may improve their impartiality.

Being at the core of the Reflow project development of a federated database of information that serves the piloted circular-economy use-cases, the main deployment of the Reflow crypto model is inside the Zenpub software component and will refine its interaction to GUI and dashboard components through a GraphQL interface and ActivityPub event protocol, as well leveraging the federating capabilities of this scheme to a stable and replicable implementation.

The developments of next-generation DLT applications as Holochain [Harris-Braun et al., 2018] and Hashgraph [Green, 2019] are also very interesting, in particular Holochain is known have the ValueFlows REA approach in course of implementation and may soon provide a platform where to further distribute the graph data beyond the federated server provision currently implemented in ZenPub.

7.2 Acknowledgments

The development and demonstration of the Reflow signature scheme and the material passport has been partially funded by the European project "REFLOW" referenced as H2020-EU.3.5.4 with grant nr.820937; the multi-party signature implementation has been partially funded by Riddle & Code GmbH as private commission.

We are grateful to all our colleagues at Dyne.org for the help and insights they have shared contributing to this work and in particular to Puria Nafisi Azizi, Danilo Spinella, Adam Burns and Ivan Minutillo; as well to Thomas Fuerstner for sharing his visions and passion for crypto development with us; and to prof. Massimiliano Sala and Giancarlo Rinaldo for facilitating Alberto's stage program with the dept. of Mathematics at University of Trento. Our gratitude goes also to our former colleagues in the DECODE project George Danezis and Alberto Sonnino for creating Coconut, a wonderful crypto scheme on which Reflow is grafted. Last not least we are grateful to all the REFLOW project consortium represented by prof. Cristiana Parisi as principal investigator.

References

- Dan Boneh, Sergey Gorbunov, Riad Wahby, Hoeteck Wee, and Zhenfei Zhang. BLS Signatures. Internet-Draft draft-irtf-cfrg-bls-signature-04, IETF Secretariat, September 2020. URL <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-bls-signature-04.txt>.
- Dan Boneh, Manu Drijvers, and Gregory Neven. Compact multi-signatures for smaller blockchains. Cryptology ePrint Archive, Report 2018/483, 2018a. <https://eprint.iacr.org/2018/483>.
- Jerome H Saltzer and Michael D Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, 1975.

Lorna Goulden, Kai M. Hermsen, Jari Isohanni, Mirko Ross, and Jef Vanbockryck. The disposable identities privacy toolkit; enabling trust-by-design to build sustainable data driven value. 2021.

Rob Kranenburg and Gaelle Gars. The cybersecurity aspects of new entities need a cybernetic, holistic perspective. *International Journal of Cyber Forensics and Advanced Threat Investigations*, 2(1):63–68, 2021. ISSN 2753-9997. doi:10.46386/ijcfati.v2i1.36. URL <https://conceptechint.net/index.php/CFATI/article/view/36>.

Lars Marten Luscuere. Materials passports: Optimising value recovery from materials. *Proceedings of the Institution of Civil Engineers - Waste and Resource Management*, 170(1):25–28, 2017. doi:10.1680/jwarm.16.00016. URL <https://doi.org/10.1680/jwarm.16.00016>.

Dyne.org. Reflow os. *constRuctive mEtabolic processes For material fIOWs in urban and peri-urban environments across Europe*, 2020. URL <https://reflowproject.eu/knowledge-hub/>.

Maayke Aimée Damen. A resources passport for a circular economy. Master’s thesis, 2012.

Wim Laurier, Jesper Kiehn, and Simon Polovina. Rea 2: A unified formalisation of the resource-event-agent ontology. *Applied Ontology*, 13(3):201–224, 2018.

Lynn Foster, Elf Pavlik, Bob Haugen, Michael Williams, Paul Mackey, Kalin Maidzhanski, Sam Pospichil, and Jef Vanbockryck. Valueflows vocabulary, 2017. URL <https://valueflo.ws>.

Dan Boneh, Manu Drijvers, and Gregory Neven. Compact Multi-signatures for Smaller Blockchains. In *Advances in Cryptology – ASIACRYPT 2018*, volume 11273 of *Lecture Notes in Computer Science*, pages 435–464. Springer, 2018b. doi:10.1007/978-3-030-03329-3_15.

Alberto Sonnino, Mustafa Al-Bassam, Shehar Bano, and George Danezis. Coconut: Threshold Issuance Selective Disclosure Credentials with Applications to Distributed Ledgers. *CoRR*, abs/1802.07344, 2018. URL <http://arxiv.org/abs/1802.07344>.

Michael Scott. The apache milagro crypto library (version 2.2). 2017.

Jan Camenisch and Markus Stadler. Proof systems for general statements about discrete logarithms. Technical report, 1997.

Neal Koblitz and Alfred J. Menezes. The random oracle model: a twenty-year retrospective. *Designs, Codes and Cryptography*, 77(2):587–610, Dec 2015. ISSN 1573-7586. doi:10.1007/s10623-015-0094-2. URL <https://doi.org/10.1007/s10623-015-0094-2>.

Dan Boneh. The decision diffie-hellman problem. Technical report, Berlin, Heidelberg, 1998.

Anna Lysyanskaya, Ronald L Rivest, Amit Sahai, and Stefan Wolf. Pseudonym systems. Technical report, 1999.

Razvan Barbulescu and Sylvain Duquesne. Updating key size estimations for pairings. *Journal of Cryptology*, 32(4):1298–1336, 2019.

Chae Hoon Lim and Pil Joong Lee. A key recovery attack on discrete log-based schemes using a prime order subgroup. In *Annual International Cryptology Conference*, pages 249–263. Springer, 1997.

Jaap-Henk Hoepman, Shehar Bano, Alberto Sonnino, Eleonora Bassi, Marco Ciurcina, Juan Carlos De Martin, Selina Fenoglietto, Antonio Santangelo, Francisco Sacramento Gutierrez, David Laniado, and Pablo Aragón. Privacy design strategies for the decode architecture. 2019.

Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.

Wim Laurier and Geert Poels. Track and trace future, present, and past product and money flows with a resource-event-agent model. *Information systems management*, 29(2):123–136, 2012.

Alon Zakai. Emscripten: an llvm-to-javascript compiler. In *Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion*, pages 301–312, 2011.

Eric Harris-Braun, Nicolas Luck, and Arthur Brock. Holochain. scalable agent-centric distributed computing. *DRAFT (ALPHA 1)-2/15/2018.[Online]*. Available: <https://whitepaperdatabase.com/wp-content/uploads/2018/08/holochain-HOT-whitepaper.pdf>, 2018.

Oded Green. Hashgraph–scalable hash tables using a sparse graph data structure. *arXiv preprint arXiv:1907.02900*, 2019.

List of Figures

1	Basic credential authentication	2
2	Session Creation	2
3	Sign and Verify	2
4	Valueflows	3
5	Keygen process sequence diagram	9