

# Portal: Time-Bound and Replay-Resistant Zero-Knowledge Proofs for Single Sign-On

Jan Lauinger, Serhat Bezmez, Jens Ernstberger, Sebastian Steinhorst  
 Technical University of Munich  
 Munich, Germany  
 name.surname@tum.de

**Abstract**—Latest identity systems rely on public blockchains to enhance user autonomy and reduce tracking from conventional identity providers. At the same time, identity systems integrate novel technologies such as zero-knowledge proofs (ZKPs) to improve data privacy and data compliance. We show that a naive verification of ZKPs at smart contracts enables replay attacks: Attackers can replay ZKPs at arbitrary times without having access to the private inputs that are required for the computation of the ZKP. To solve this problem, we construct a transaction sequence which verifies time-bound and replay-resistant ZKPs at smart contracts. Our construction introduces an additional but constant fee of 0.14\$ per verification of a ZKP on the public blockchain Ethereum. With our new construction, we propose Portal, a novel identity system for decentralized single sign-on.

**Index Terms**—Zero-knowledge Proofs, Smart Contracts, Decentralized Resolution, Single Sign-On

## I. INTRODUCTION

**Motivation:** Almost every service of today’s web manages *users* based on an identifiable session and requires a mechanism to authenticate *users* beforehand. The *user* authentication uniquely identifies every *user* of the system and guarantees that the session is unique to one *user*. To avoid each web service from implementing their own identity and authentication system, OpenID Connect (OIDC), as the latest Single Sign-On (SSO) protocol, was standardized in 2014 [1]. The SSO paradigm delegates *user* authentication at a web service towards a third-party Identity Provider (IdP), which handles the unique identification of the *user* (cf. case *a* in Figure 1).

Even though delegated authorization and authentication via SSO is very convenient and cheap for *users*, IdPs can track every log-in and data access of a *user*. To solve the misaligned incentives between all parties, recent approaches (e.g. Sign-In with Ethereum (SIWE) [2]) replace the IdP with a public blockchain and provide *users* with new notions of autonomy [3]. Polygon ID [4] employs Zero-knowledge Proof (ZKP) technology to enhance data privacy and data compliance of *users*. Modern identity systems rely on certification ecosystems, where issuers verify and attest to data claims made by *users* [4]. Similarly, recent works [5] rely on assumptions (e.g. existence of trustworthy issuers) which go beyond the requirements of SSO systems. Because, in the trust establishment phase of SSO systems, *users* agree to the IdPs’s terms and conditions which require *users* to honestly create profiles without requesting specific credentials [1].

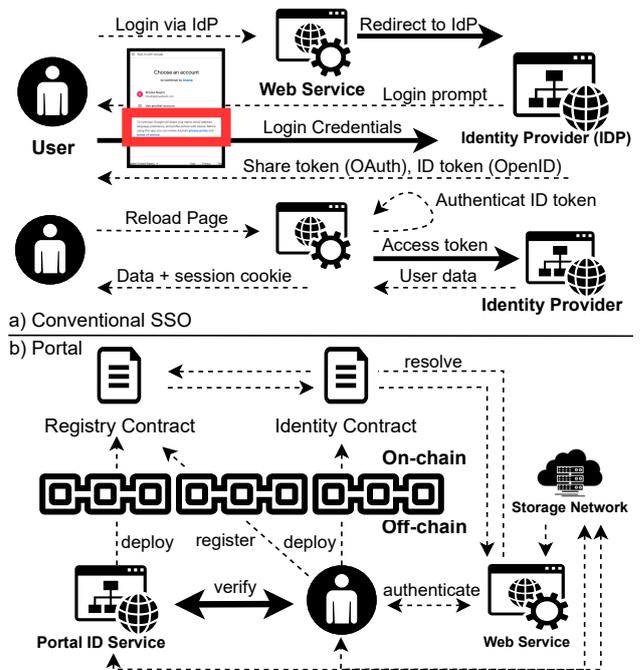


Fig. 1. a) Overview of the Single Sign-On (SSO) delegated authentication and authorization where the *user* agrees to a fixed policy (red box) of the Identity Provider (IdP). Bold arrows indicate *user*-to-IdP interactions which track *user* activities. b) Simplified view of the *Portal* identity system, where *users* manage data and authenticate towards web services with self custody.

**Challenge:** In this work and according to the requirements found in SSO systems, we investigate the honest creation and management of user data, which does not require any form of third-party attestation. In this scenario, we entirely rely on the interaction between *users* and smart contracts, where smart contracts verify the data claims made by *users*. To create a claim on a data sample, *users* convince smart contracts that the data sample complies with a public statement. If the smart contract successfully verifies the claim, then the smart contract accepts a mapping between a data claim and the address of the *user*. If *users* create claims on private data, then the smart contracts validate ZKPs asserting the claim. Based on accepted claims, *users* can authenticate to any third party.

We find that replay attacks are a concern because blockchain logs transparently expose transaction payloads to adversaries. Thus, any claim can be replayed by re-executing the same contract functionality using previously exposed payloads.

**Contribution:** For claims on private data, we show that replay attacks can be prevented. To do so, we introduce a new transaction sequence which unequivocally binds the proof computation to a specific *user* and time (cf. Section IV-C). Instead of using a verifier-chosen nonce that binds a proof presentation to a verification session [5], [6], our transaction sequence relies on the blockchain Proof of Stake (PoS) randomness as the verifier-chosen nonce. Our transaction sequence achieves an efficient cost structure as it does not require additional contracts that prevent replay attacks (e.g. access control smart contracts [7]). Based on this contribution, we propose a novel identity system, called *Portal*, which supports on-chain and off-chain validations of ZKPs on *user* data during *user* authentication (cf. bottom part of Figure 1). In summary,

- Our new transaction sequence secures on-chain ZKP verifications against replay attacks (cf. Section IV-C).
- We propose *Portal*, an alternative SSO solution with enhanced privacy and control.
- We open-source<sup>1</sup> our proof of concept of *Portal* and evaluate operation costs (cf. Section VI).

In systems with strong know your customer (KYC) requirements, where users cannot be trusted to responsibly operate claims, we want to highlight that *Portal* can and should be used with third-party attestations.

## II. PRELIMINARIES

### A. Secure Hash Functions

A secure hash function implements an algorithm, where

- $\mathbf{h.Hash}(m) \rightarrow (h)$  takes as input a message string and outputs a constant size hash string  $h$ .

and guarantees three properties: *Preimage-resistance* ensures that given  $h$ , and attacker cannot find  $m$  if  $h = \mathbf{h.Hash}(m)$ . *Second preimage-resistance* ensures that given  $m_1$  an attacker cannot find  $m_2$  such that  $\mathbf{h.Hash}(m_1) = \mathbf{h.Hash}(m_2)$  holds, with  $m_1 \neq m_2$ . *Collision-resistance* ensures that finding  $m_1 \neq m_2$  with  $\mathbf{h.Hash}(m_1) = \mathbf{h.Hash}(m_2)$  is infeasible.

### B. Public Key Cryptography (PKC) & Digital Signatures

PKC systems provide users with complementing key pairs, a *public key* and a *private key*, where the *private key* is never disclosed. Using PKC, we define a digital signature scheme on a message string  $m$  with the algorithms, where

- $\mathbf{pk.Setup}(1^\lambda) \rightarrow (sk, pk)$  uses a security parameter to output a PKC private key  $sk$  and public key  $pk$ .
- $\mathbf{pk.Sign}(sk, m) \rightarrow (\sigma)$  takes as input the secret key and a message string  $m$ , and outputs the signature  $\sigma$ .
- $\mathbf{pk.Verify}(pk, m, \sigma) \rightarrow \{0, 1\}$  takes as input the public key, the message message string, and a signature, and outputs either a 1 if the signature verification succeeds. Otherwise, the output is a 0.

### C. Commitment Schemes

We define commitment schemes with algorithms, where

- $\mathbf{cs.Commit}(x) \rightarrow (c, w)$  takes as input the data  $x$ , generates a witness (e.g. randomness), and outputs a commitment string  $c$  and the witness  $w$ .

<sup>1</sup><https://github.com/jplaiui/portal>

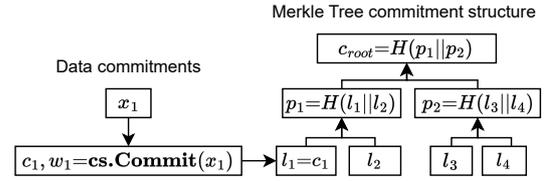


Fig. 2. Binary Merkle Tree (MT) commitment structure on a set of *data items*  $x_i$ , with  $i \in \{0, \dots, N\}$ . The depicted MT has a depth  $D=2$ , leaves  $l_1, \dots, l_{2^D}$ , parents  $p_1, p_{2^{*2}D-1-2}$ , a root  $c_{root}$ , and depends on the hash function  $H$ . The root  $c_{root}$  represents the commitment string and the witness  $w$  consists of the internal witnesses  $w_i$ , with  $i \in \{0, \dots, N\}$  and a Merkle path  $f_{path}(x_i)$  that depends on the committed *data items*. In this figure, the witness comprises the set of tuples  $w = [(w_1, [l_2^2, p_2^2] = f_{path}(x_1))]$ , where  $l_2^2$  indicates that  $l_2$  is the second concatenation when computing  $p_1$ .

- $\mathbf{cs.Open}(w, x, c) \rightarrow \{0, 1\}$  uses the witness to verify if the committed data matches the commitment string. In case of a match, the algorithm outputs 1, and 0 otherwise.

Commitment schemes are *hiding* if the commitment string  $c$  does not leak any information of  $x$  to an adversary with access to  $c$ . Commitment schemes are *binding* if there exists an unequivocal mapping between  $x$ ,  $w$ , and  $c$ , such that an adversary cannot find a second valid opening yielding  $1 = \mathbf{cs.Open}(w', x', c)$ , with  $x' \neq x$ ,  $w' \neq w$ . In this work, we rely on Merkle Tree (MT) commitments [8], [9] (cf. Figure 2).

### D. Zero-knowledge Proof Systems

A general-purpose ZKP system allows a *prover* to convince a *verifier* of knowing a secret witness  $w$  which satisfies a statement expressed via a computation circuit  $\mathcal{C}$ . The *verifier* relies on an polynomial time algorithm to verify if  $w$  is a valid proof of the statement and learns nothing beyond the validity of the statement. A ZKP system achieves the properties of (i) *completeness*, where an honest *prover* with a valid witness convinces an honest *verifier*, (ii) *soundness*, where a cheating *prover* without a valid witness cannot convince an honest *verifier*, and (iii) *zero-knowledge*, where a cheating *verifier* learns nothing beyond the validity of a proven statement. We use a ZKP system with the algorithms, where

- $\mathbf{zk.Setup}(1^\lambda, ccs_{\mathcal{C}}) \rightarrow (pk, vk)$  takes as input a security parameter and a compiled constraint system expressing a circuit  $\mathcal{C}$ , and outputs the *prover* and *verifier* keys  $pk, vk$ .
- $\mathbf{zk.Prove}(ccs_{\mathcal{C}}, w, pk) \rightarrow \pi$  takes as input the compiled constraint system, a private witness, and the prover key  $pk$  and outputs a proof  $\pi$ .
- $\mathbf{zk.Verify}(w_{pub}, vk, \pi) \rightarrow \{0, 1\}$  takes as input a public witness  $w_{pub}$ , the verifier key  $vk$ , and the proof  $\pi$  and outputs a 1 if  $\pi$  combined with  $vk$  successfully verify against  $w_{pub}$ . Otherwise a 0 is returned.

If a ZKP system computes  $\mathbf{cs.Open}$  (cf. Section II-C) as  $\mathcal{C}$  while taking the witness  $w$  as a private input, then a commitment opening maintains input privacy. For example, computing an MT inclusion proof against a commitment  $c_{root}$  requires the circuit  $\mathcal{C}$  to derive  $c_{root}'$  based on the secrets  $x_1$  and  $w$ , and check if  $c_{root}' = c_{root}$  (cf. Figure 2).

## E. Blockchains & Smart Contracts

Public blockchains are open computer networks anyone can join, which run a *consensus* protocol to agree upon a common and correct *state*  $s_t$  at time  $t$ . The *state* maintains two types of accounts; the externally owned account (EOA) and the *smart contract*. An EOA is controlled by a PKC key pair and is updated if a *user* owning the key pair sends signed transactions to the blockchain. A *smart contract* is an executable program at a unique address that can be invoked by transactions. The execution of *smart contracts* is measured in *gas* and must be paid by a medium called cryptocurrency. Transactions are stored in a *mempool* until a new state update is proposed via a new block of transactions. Blockchain nodes verify new blocks by comparing local *state* updates with the digests found in new blocks. If the verification succeeds, transactions are locally applied such that the network reaches a new global state.

*Blockchains* achieve the properties of *safety* which provides *state* integrity according to past *states*, *liveness* where every transaction is eventually included in the *state*, and *consistency* where every node eventually has the same view of the *state*. Blockchain transactions are *non-repudiable* as signatures of transactions unambiguously identifies *users*.

## III. SYSTEM MODEL

### A. Notations

**Key pairs** are the *public* and *private* keys of a PKC system. **Addresses** are derived from a *user's* *public* key and exist as 42-character hexadecimal strings appended with '0x'.

**Wallets**  $W$  generate and maintain *key pairs* and, with that, control the *address*  $W_{addr}$  corresponding to the *key pairs*.

**Data items** are key-value pairs, where the key string is a descriptor of the value instance that expresses the data.

**Statements**  $\phi = \text{"key-op-comp"}$  are strings that express relations between a value *comp* and a data item with  $\text{key=key}$ . *Statements* use at least one *key*, one operator *op* (e.g.  $>, <, \neq, =, \geq, \in$ , etc.) and one comparison value *comp* (e.g. threshold).

**Claims** exist as *public claims*  $\text{claim}^{pub} = \{d, \phi, t\}$  and as *private claims*  $\text{claim}^{priv} = \{d, \phi, L, e_{id}, t\}$ . *Public claims* include the *data item*  $d$ , a *statement*  $\phi$ , and a timestamp  $t$ . If the *data item* of  $\text{claim}^{pub}$  is stored externally, then  $d$  is set to a location identifier  $d=L$ . *Private claims* include a *data item*  $d$ , a *statement*  $\phi$ , a location identifier  $L$ , an event identifier  $e_{id}$ , and a timestamp  $t$ . In  $\text{claim}^{priv}$ , the value instance of  $d$  is a commitment string  $c$  (e.g.  $d[\text{"age"}] : c$ ) and the location identifier points to a circuit storage address as  $L=L_{p_{\Pi}}$ .

**Circuits** are tuples  $p_{\Pi} = \{\mathcal{C}, \phi, ccsc, w_{pub}, pk, vk, L_C\}$ , where the compiled constraint system  $ccsc$  expresses a provable representation of a circuit  $\mathcal{C}$  that implements the assertions expressed by the *statement*  $\phi$ . To assert *statements*, the circuit  $\mathcal{C}$  evaluates private inputs  $w$  to a representation which can be compared against public inputs  $w_{pub}$ . The prover and verifier keys  $pk, vk$  are created by running the setup algorithm  $\mathbf{zk.Setup}$  of a proof system  $\Pi$ . If the verification call of the circuit  $\mathcal{C}$  is deployed as a *smart contract*, then the locator  $L_C$  is set to the *address* of the *circuit contract*. Otherwise,  $L_C = \text{null}$ .

**Transactions** are tuples  $tx = \{\sigma, d_{pl}, t_{addr}, g_{used}\}$  with a signature  $\sigma$  from the transaction sender, a data payload  $d_{pl}$ , a *gas* value  $g_{used}$  and a destination *address*  $t_{addr}$ . *Transactions* are used to invoke and pay for *smart contract* calls at an address  $t_{addr}$  and provide non-repudiation of the transaction sender.

**Circuit contracts**  $C^C$  verify ZKPs on-chain and emit events  $e_{id}$  according to the outcome of a ZKP verification. *Circuit contracts* expose the *sample* and *verify* methods. If a *transaction* calls the *sample* method, then  $C^C$  associates a PoS randomness as a nonce with the wallet address of the *user* in a map  $m[W_{addr}].nonce$ . The randomness is used during the *verify* method which verifies a ZKP.

### B. System Roles

**Users** hold *wallets*, deploy *identity contracts*, and register the *address* of the *identity contract* at the *registry contract* after passing an authenticity verification at the *identity service*. *Users* individually manage *claims* and *attestations*, and authenticate themselves at *third-party services* by linking or presenting data. *Users* count as *issuers* in the context of signing and sharing credentials towards other *users*.

**Identity services** deploy and maintain *registry* and *circuit contracts* and connect *users* to the *Portal* identity system. We envision non-profit organizations to take the role of the *identity service* and assume that *identity services* have the expertise to create secure ZKP circuits which evaluate *claims* of *users*.

**Third-party services** (e.g. web services) authenticate *users* based on the *Portal* identity system and trust *issuers*.

**Blockchain networks** provide decentralized and verifiable computation and storage through *smart contracts* and manage *registry*, *identity*, and *circuit contracts*.

**Storage networks** provide decentralized, fault-tolerant, and high-availability storage of data at locations  $L$  and are used to store larger data objects such as circuit parameters  $p_{\Pi}$ .

### C. Threat Model

We assume that transactions sent to blockchain nodes are secured via Transport Layer Security (TLS) such that the TLS properties of message confidentiality, integrity, and authenticity hold. We assume that (i) honest *users* are able to resolve the correct state  $s_t$  of the blockchain at time  $t$ , that (ii) collision resistant hash functions are used in the blockchain PoS protocol to determine the block randomness [10], and (iii) active, adaptive, and probabilistic polynomial time (PPT) adversaries that are able to perform machine-in-the-middle (MITM) attacks and intercept communication traffic. Adversaries are not able to block traffic indefinitely and cannot modify intercepted traffic. Adversaries have access the *mempool*, can access transaction payloads by observing blockchain logs, and replay transactions  $tx$  or ZKPs of a *user*.

## IV. CONSTRUCTING TIME-BOUND AND REPLAY-RESISTANT ZKPS

### A. ZKP Verification at Smart Contracts

As the initial setup, we assume access to a circuit tuple  $p_{\Pi}$ , which has been instantiated by a trusted party  $p_0$ . The party  $p_0$  derives the solidity verification code of  $\mathcal{C}_1 \in p_{\Pi}$

$\text{assertClaim}(d, p^{\text{MT}}, W_{\text{addr}}, n; \text{root}^{\text{MT}}, W_{\text{addr}}, n, \phi):$ 1. assert: $n \stackrel{?}{=} n; W_{\text{addr}} \stackrel{?}{=} W_{\text{addr}}; 1 \stackrel{?}{=} f_{\phi}(d)$ 2. return: $1 \stackrel{?}{=} \text{cs.Open}(p^{\text{MT}}, d, \text{root}^{\text{MT}})$
--

Fig. 3. ZKP circuit to verify a *data item*  $d$  of a private claim against a MT commitment  $\text{root}^{\text{MT}}$ . The MT has a depth of 5 and a path  $p^{\text{MT}}$  as the private witness. The circuit has 9.29K constraints and evaluates  $d$  against  $\phi = "d[\text{age}] > -18"$  using the function  $f_{\phi}$ . The semicolon ; separates private inputs (left of ;) from boldly formatted public inputs (right of ;).

for the creation and deployment of a circuit contract  $C^{C_1}$  (cf. steps 1.4, 1.5 of Figure 4).  $\Pi$  uses a ZKP system which compiles the circuit  $C_1$ . The circuit  $C_1$  performs an address and nonce check, asserts a private *data item* against a statement  $\phi$ , and checks if the *data item* computes to a public commitment string (cf. **assertClaim** logic of Figure 3). Now, a *user* as party  $p_1$  is able to compile transactions with a payload that contains the bytes of a ZKP  $\pi$ , and call the deployed contract  $C^{C_1}$  for an on-chain verification of  $\pi$ .

### B. Binding ZKP Computations to the PoS Randomness

In the following we define a transaction sequence where a user  $p_1$  compiles the transaction  $tx_1$  to call the *sample* method of the contract  $C^{C_1}$ . Upon receiving  $tx_1$ ,  $C^{C_1}$  associates the latest PoS randomness  $r$  with the *user's* wallet address by depositing both parameters into the map  $m[W_{\text{addr}}]_{\text{nonce}}$ . Initially the randomness is concatenated with a state string to represent the nonce as  $\text{nonce} = s_t.\text{prevrandao} || "-0"$ . After  $C^{C_1}$  samples the nonce, *users* fetch and use the deposited nonce to compute a ZKP  $\pi$  using the circuit  $C_1$ . To prevent replay attacks and ensure time-bound proofs (cf. Section IV-C), the ZKP circuit  $C_1$  takes in and compares both the *user's* wallet address and the nonce as private inputs and public inputs. Notice that binding values (e.g. the nonce) to a ZKP computation via public inputs is secure [6]. In a transaction  $tx_2$ ,  $p_1$  calls the *verify* method of  $C^{C_1}$ , which upon a successful verification of  $\pi$ , sets the nonce to  $m[W_{\text{addr}}]_{s_t.\text{prevrandao} || "-1"}$  and emits an event with an identifier  $e_{id}$ . If party  $p_1$  presents  $e_{id}$  towards any *third-party service*, then the *third-party service* can use  $e_{id}$  to resolve and verify a successful on-chain ZKP verification via smart contract logs (cf. steps 2.1-2.8 in Figure 4).

### C. Security Analysis

**Theorem 1.** *If a party  $p_1$  with access to*

- a smart contract  $C^{C_1}$
- a secure proof system  $\Pi_{\pi}$
- a secure signature scheme  $\Pi_{\sigma}$
- a secure hash function  $\Pi_H$

*performs the sequence of computations*

- 1)  $p_1$  compiles and signs a transaction  $tx_1$  with  $\Pi_{\sigma}$ . Sign
- 2)  $p_1$  calls  $C^{C_1}.\text{sample}$  with  $tx_1$  such that  $C^{C_1}$  generates the *prevrandao* randomness  $r$  using  $\Pi_H$ . Hash and stores  $m[W_{\text{addr}}^{p_1}]_r || "-0"$  at timestamp  $t_1$
- 3)  $p_1$  fetches  $r$  from  $m[W_{\text{addr}}^{p_1}]$
- 4)  $p_1$  computes  $\pi = \Pi_{\pi}.\text{Prove}(ccs_{C_1}, w, pk)$

- 5)  $p_1$  compiles and signs a transaction  $tx_2$  with  $\Pi_{\sigma}$ . Sign, where  $\pi \in tx_2.d_{pl}$
- 6)  $p_1$  calls  $C^{C_1}.\text{verify}$  with  $tx_2$  and  $C^{C_1}$  sets  $m[W_{\text{addr}}^{p_1}]_r || "-1"$  at timestamp  $t_2 > t_1$

*under the assumptions that*

- $C^{C_1}$  runs on a blockchain which guarantees liveness, consistency, safety

*we say that the proof  $\pi$  is resistant against replay attacks performed by a malicious PPT adversary such that  $\pi \in tx_2'$  is never accepted by  $C^{C_1}$ . And, we say that computing  $\pi$  is bound by the time  $t_1$  and cannot be accepted after  $t_2$ .*

**Proof 1.** At time  $t_1$ , the adversary  $\mathcal{A}$  cannot change  $tx_1$  of  $p_1$  as unforgeability of transaction signatures holds. But,  $\mathcal{A}$  is capable of registering the same nonce of  $p_1$  twice at  $C^{C_1}$  with  $tx_1'$ .  $C^{C_1}$  maps the nonce of  $\mathcal{A}$  at the address  $m[W_{\text{addr}}^{\mathcal{A}}]$ . At time  $t > t_1$ ,  $\mathcal{A}$  uses the blockchain logs to access  $tx_2$  of  $p_1$  and, with that,  $\pi$ . If  $\mathcal{A}$  replays  $\pi$  in a transaction  $tx_2'$  and calls  $C^{C_1}.\text{verify}$ , then the verification of circuit  $C_1$  fails because of the following issue. The proof  $\pi$  has been computed with the address  $m[W_{\text{addr}}]$  as private input and  $C^{C_1}$  asserts that  $\pi$  is verified against the owner of  $tx_2$ .  $\mathcal{A}$  has signed  $tx_2'$  such that  $C^{C_1}$  asserts  $\pi$  with the public input  $W_{\text{addr}}^{\mathcal{A}}$  taken from  $tx_2'$ , which fails.

Further,  $\mathcal{A}$  tries to replay a previously accepted proof  $\pi^{\mathcal{A}}$  (sampled and proven with  $tx_1^{\mathcal{A}}$  and  $tx_2^{\mathcal{A}}$ ). At time  $t > t_1, t_2$ ,  $\mathcal{A}$  cannot replay  $\pi^{\mathcal{A}}$  because, even though a reoccurring nonce is negligible (collision resistance of PoS randomness),  $C^{C_1}$  prevents overwriting an existing map entry at  $m[W_{\text{addr}}^{\mathcal{A}}]$  if a nonce has already been set. Thus, our scheme is *replay-resistant* and *time-bound* as  $\pi$  can only be computed at  $t_2$  after randomness has been sampled with  $tx_1$  at time  $t_1 < t_2$ .

## V. Portal IDENTITY SYSTEM

### A. System Goals

**Sybil resistance** prevents an adversary to register an arbitrary amount of pseudonymous identities.

**Decentralized resolution** guarantees that the storage and computation of user data remain publicly verifiable, trustless, and available towards a resolving *third-party service*.

**On-chain & off-chain verification of private data** allows *users* to (i) present data to a *third-party service*, where the data has been verified at smart contracts or (ii) interactively convince a *third-party service* of a data verification.

**Cost-efficiency** optimizes operation costs for *third-party* and *identity services* and enables scalability of *Portal* with cheap maintenance costs for the *identity service*.

### B. Architecture

*Portal* requires two new contracts, where

**Registry contracts**  $C^{\text{reg}}$  maintain the map  $m[W_{\text{addr}}]_{C_{\text{addr}}^{\text{id}}}$  linking registered *wallet addresses* and *addresses of identity contracts*.  $C^{\text{reg}}$  exposes a *register* method which requires the transaction payload to include an *identity service* signature on a new *identity contract address*. Further, for the identification of circuits,  $C^{\text{reg}}$  maintains a map  $m[\text{name}^C]_{L_{p_{\Pi}}}$  which associates location identifiers of circuit parameters  $L_{p_{\Pi}}$  with circuit names  $\text{name}^C$ .

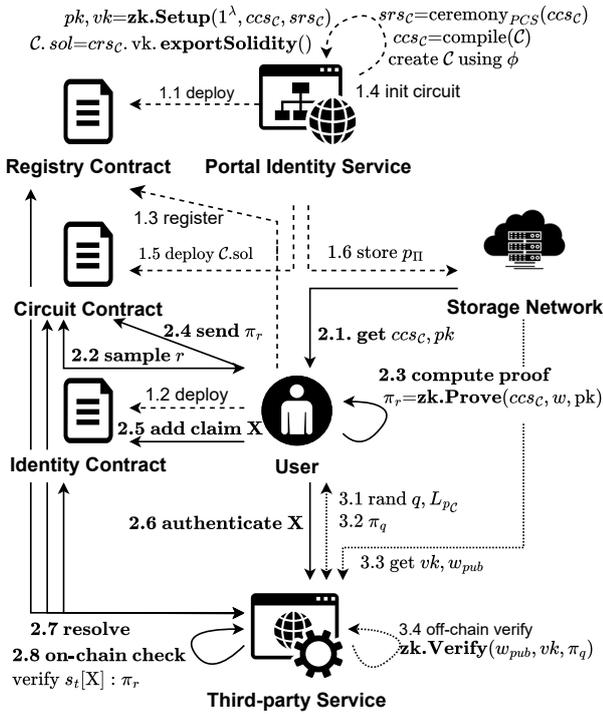


Fig. 4. *Portal* architecture in the context of managing a *private claim*. The system deployment, *user* registration, and the circuit pre-processing is indicated with dashed arrows (1.1-1.6). The on-chain verification of *private claims* at time  $t_1$ , and *private claim* presentation towards a *third-party service* is depicted with solid lines (2.1-2.8). The live verification at time  $t > t_1$  of a *private claim* is indicated with dotted lines (3.1-3.4).

**Identity contracts**  $C^{id}$  maintain *claims*, *attestations*, and *revocations* with the maps  $m[\text{name}^{\text{claim}}]\text{claim}$ ,  $m[\text{name}^{\text{att}}]a$ , and  $m[a_{id}]\text{rev}$ , where  $a_{id}$  is an attestation identifier. The unique strings  $\text{name}^{\text{claim}}$ ,  $\text{name}^{\text{att}}$  represent claim and attestation names.

The registration of a new *user* in the *Portal* system depends on two transactions. The first transaction deploys the *identity contract*  $C^{id}$  of the *user*. In the same way as the *registry contract*, the constructor of the *identity contract* sets the deploying party as the owner of the contract. Only the owner of  $C^{id}$  is able to call methods which modify the state of  $C^{id}$ . The compilation of the second transaction requires the *user* to obtain a signature  $\sigma_{C^{id}_{addr}}$  from the *identity service* on the *identity contract address*. Before signing any  $C^{id}_{addr}$ , the *identity service* verifies and deduplicates *users*, such that *sybil resistance* holds in the *Portal* system. *Users* use the second transaction to invoke the *register* method at the *registry contract*  $C^{reg}$ , which checks the signature validity of  $\sigma_{C^{id}_{addr}}$  before including the *user's* wallet address and  $C^{id}_{addr}$  into the map  $m[W_{addr}]C^{id}_{addr}$ . If the *user* shares the wallet address  $W_{addr}$  with any *third-party service*, then the *third-party service* is able to resolve  $C^{id}_{addr}$  via the map  $m[W_{addr}]C^{id}_{addr}$  such that *decentralized resolution* holds.

Once *users* are registered, *users* can create *private claims* by following the transaction sequence which prevents replay attacks (cf. Section IV). Further, *users* can decide to partake in a live verification of *private data*, where a proof system is deployed between the *user* and the *third-party service* (cf. steps 3.1-3.4 in Figure 4). The live verification

ensures that *private claims* are not validated by smart contracts at timestamps in the past. The data verification modes of *Portal* ensure support for on-chain and off-chain verification of *private data*.

*Third-party services* resolve and verify *user* data through a *Portal* plugin, which performs a signature challenge before every data verification. Similar to the SIWE sign-in challenge [2], our signature challenge demands the *user* to compute a signature on a randomly sampled nonce using the wallet *key pair*, where our plugin samples the nonce.

*Notice:* Replaying publicly accessible data cannot be prevented which is why this work solves replay attacks of claims made on *private data*. However, if an adversary operates claims on *public data* and cheats (upload false claim), then the blockchain properties guarantee that adversaries remain accountable once misbehaviour is detected. In this case, the reputation of a *user* declines.

## VI. EVALUATION

### A. Implementation

The *Portal* proof of concept was conducted locally using the *Ganache*<sup>2</sup> test network (v7.8.0) as the public blockchain. We rely on the solidity compiler *solc* v0.8.20 as the PoS block randomness *prevrando* is available in all versions above v0.8.18. We develop a *Portal* Golang System Development Kit (SDK) to deploy and maintain *Portal* at every party and use the official Ethereum repository *go-ethereum*<sup>3</sup> including *abigen* v1.10.16 to interact with smart contracts. We convert transaction costs into US dollars based on the rate 2084.42\$ per 1 *eth* (November 2023) and select the default *Ganache* gas price  $gas_{\text{price}} = 2\text{gwei}$ . We select the Golang *gnark* (v0.9.1) repository [11] as the ZKP system and configured (i) the *plonk* backend with a universal setup to verify ZKPs on-chain. To prove and store *private claims* efficiently, we benchmark the ZKP circuit  $C_1$  (cf. Figure 3), which evaluates *data items* of  $\text{claims}^{\text{priv}}$  as private input against a MT commitment as the public input. We use the MiMC hash function [12] to compress the MT data. We open-source the *Portal* code with the *smart contracts* and simulation scenarios in the repository<sup>4</sup>.

### B. Costs Analysis

The evaluation uses a MacBook Pro with the Apple M1 Pro chip and 32 GB of Random Access Memory (RAM). The benchmarks average ten executions of the same experiment.

Table I shows the *Portal* cost analysis, where transaction costs are computed according to  $\text{tx}_{\text{cost}} = gas_{\text{used}} \cdot gas_{\text{price}}$ . We explain the execution times in the range of milliseconds with the local deployment of *Portal*. By deploying *Portal* on the Sepolia<sup>5</sup> testnet, we measured transaction resolution times taking around 150ms and transaction calls taking between 1.3s ( $C_1$  deploy) and 9.4s (sample+verify\_ $\pi$ +claim $^{\text{priv}}$ ). We explain higher execution times of the transactions that deploy  $C_1$  and verify a proof

<sup>2</sup><https://github.com/trufflesuite/ganache>

<sup>3</sup><https://github.com/ethereum/go-ethereum>

<sup>4</sup><https://github.com/jplaiui/portal>

<sup>5</sup><https://www.alchemy.com/overviews/sepolia-testnet>

TABLE I  
Portal BENCHMARKS WITH THE ABBREVIATION BYTE CODE (BC)

Tx / Call	Type	Cost (eth/\$)	Time (ms)	Size (kB)
$C^{reg}$	deploy	4.1e-3/8.6	18	bc:6.5,tx:6.6
$C^{id}$	deploy	6.5e-3/13.5	10	bc:10,tx:10
$C^{C_1}$	deploy	4.9e-3/10.2	<b>385</b>	bc:7.4,tx:12
set_ $C_1$	$C^{reg}$	8.4e-5/0.18	11	tx: 0.46
register	$C^{reg}$	7.4e-5/0.16	51	tx: 0.3
claim <sup>pub</sup>	$C^{id}$	6.4e-05/0.13	3	tx: 0.48
sample	$C^{C_1}$	6.6e-05/0.14	6	tx: 0.1
verify_ $\pi$	$C^{C_1}$	8.4e-4/ <b>1.76</b>	<b>252</b>	tx: 1.20
claim <sup>priv</sup>	$C^{id}$	3.9e-4/0.82	21	tx: 0.68
setup $_{plonk}^{C_1}$	off-chain	-	1029	$p_{\Pi}$ : 7430
prove $_{plonk}^{C_1}$	off-chain	-	195	$\pi$ : 0.552
set/get $_{PPFS}^{PII}$	off-chain	-	631 / 66	7430
get $_{W/n}^{C_1}$	off-chain	-	10/6.2/4.8	42/78/130

of  $C_1$  with the corresponding higher transaction sizes. Compared to other contracts, which initialize empty maps, the byte code of  $C^{C_1}$  stores large cryptographic parameters which increase the transaction size of  $C^{C_1}$ . Except transactions of the type deployment and the transaction to verify a ZKP on-chain, the cost per transaction remains below 1\$. Thus, as claims are verified once and shown multiple times, we consider *Portal* as cost-efficient.

## VII. DISCUSSION

### A. Related Works

The work DecentID [13] introduces a *smart contract* identity system which resolves user data via four different contract types. In contrast to DecentID, *Portal* supports enhanced data privacy through on-chain and off-chain ZKP computations.

The work *zk-creds* [5] proposes the first construction of anonymous zero-knowledge Succinct Non-Interactive Arguments of Knowledge (zkSNARK) credentials. With a verifier-chosen nonce, *zk-creds* prevents credential replays towards the verifier in the off-chain context. By contrast, *Portal* works on chain and applies the PoS randomness to prove zkSNARK claims in unique verification sessions.

The work Zebra [7] introduces a zkSNARK credential scheme with an on-chain ZKP verification at an access control contract. Before a *user* authenticates at an application smart contract with a wallet address  $W_{addr}$ , the *user* posts a ZKP to the access control contract to provide access privileges to  $W_{addr}$ . Instead of relying on additional smart contracts, *Portal* improves the cost-efficiency by solving ZKP replay attacks via a cheap transaction sequence.

The work *zkLogin* [14] constructs a modified OIDC nonce to authenticate transactions in the on-chain context via existing OIDC credentials. *Portal* tries to minimize the reliance on third-party entities (e.g. OIDC providers) and does not answer the question whether session management is handled by an extra provider or the *third-party service*.

### B. Limitations & Future Work

*Portal* runs on the native blockchain network called layer 1 (L1). To optimize transaction costs, we envision deploying *Portal* via scalable layer 2 (L2) networks (e.g. zk-rollups [15]). We expect that our proof-of-concept

TABLE II  
COMPARISON WITH RELATED WORKS.

Paper	Dec. Resolution	On/Off-chain Verify	Extra Contract
<i>DecID</i>	✓	✗ / ✗	✓
<i>zk-creds</i>	✗	✗ / ✓	✗
<i>Zebra</i>	✓	✓ / ✗	✓
<i>zkLogin</i>	✗	✓ / ✓	✗
<i>Portal</i>	✓	✓ / ✓	✗

implementation requires minor adjustments to reach L2 compatibility as existing tooling for L2 deployments exist. With a L2 deployment, *Portal* must be compared towards related works with regard to cost and efficiency. Another item of future work is the security assessment under relaxed assumptions of blockchain properties and the consideration of censorship implications. Concerning decentralizing the *identity service*, we either (i) register *users* based on a multi-party signature issued by multiple *identity services*, or (ii) maintain a list of public keys in the *registry contract*, such that public keys authorize *identity services*. We like to highlight that *Portal* is compatible with data provenance solutions if *users* interact with attesting oracle services [16], [17]. To align *Portal* with standardization efforts, we see OIDC, W3C Decentralized Identity (DID) and Verifiable Credential (VC) as appealing compliance goals.

## VIII. CONCLUSION

In this work, we construct a time-bound and replay-resistant ZKP verification at smart contracts. On top our construction, we present *Portal*, a modern identity system with enhanced privacy and control. *Portal* is designed to satisfy regulatory privacy requirements and provides *third-party services* with a plugin to resolve and verify private or public data claims of *users*. As such, *Portal* serves as the first SSO alternative with conventional usability that gives *users* a choice to pick enhanced control and privacy at small costs.

## ACKNOWLEDGMENT

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of ‘‘Souverän. Digital. Vernetzt.’’. Joint project 6G-life, project identification number: 16KISK002.

## REFERENCES

- [1] C. Mainka, V. Mladenov, J. Schwenk, and T. Wich, ‘‘Sok: single sign-on security—an evaluation of openid connect,’’ in *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2017, pp. 251–266.
- [2] W. Chang, G. Rocco, B. Millegan, N. Johnson, and O. Terbu, ‘‘Erc-4361: Sign-in with ethereum [draft],’’ <https://eips.ethereum.org/EIPS/eip-4361>, October 2021, [Online serial] Accessed: 2023-11-15.
- [3] J. Ernstberger, J. Lauinger, F. Elsheimy, L. Zhou, S. Steinhorst, R. Canetti, A. Miller, A. Gervais, and D. Song, ‘‘Sok: Data sovereignty,’’ *Cryptology ePrint Archive*, 2023.
- [4] P. Labs, ‘‘Polygonid: A blockchain-native identity system,’’ <https://polygonid.com/>, accessed: 2023-03-04.
- [5] M. Rosenberg, J. White, C. Garman, and I. Miers, ‘‘zk-creds: Flexible anonymous credentials from zksnarks and existing identity infrastructure,’’ in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2023, pp. 790–808.

- [6] K. Bagheri, M. Kohlweiss, J. Siim, and M. Volkhov, "Another look at extraction and randomization of groth's zk-snark," in *Financial Cryptography and Data Security: 25th International Conference, FC 2021, Virtual Event, March 1–5, 2021, Revised Selected Papers, Part I 25*. Springer, 2021, pp. 457–475.
- [7] D. Rathee, G. V. Policharla, T. Xie, R. Cottone, and D. Song, "Zebra: Anonymous credentials with practical on-chain verification and applications to kyc in defi," *Cryptology ePrint Archive*, 2022.
- [8] R. Dahlberg, T. Pulls, and R. Peeters, "Efficient sparse merkle trees: Caching strategies and secure (non-) membership proofs," in *Secure IT Systems: 21st Nordic Conference, NordSec 2016, Oulu, Finland, November 2–4, 2016. Proceedings 21*. Springer, 2016, pp. 199–215.
- [9] H. Ritzdorf, K. Wüst, A. Gervais, G. Felley, and S. Capkun, "Tls-n: Non-repudiation over tls enabling-ubiquitous content signing for disintermediation," *Cryptology ePrint Archive*, 2017.
- [10] B. Edgington, "Upgrading ethereum: 2.9.2 randomness." [https://eth2book.info/capella/part2/building\\_blocks/randomness/](https://eth2book.info/capella/part2/building_blocks/randomness/), 2023.
- [11] G. Botrel, T. Piellard, Y. E. Housni, I. Kubjas, and A. Tabaie, "Consensys/gnark: v0.9.0," Feb. 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.5819104>
- [12] M. Albrecht, L. Grassi, C. Rechberger, A. Roy, and T. Tiessen, "Mimc: Efficient encryption and cryptographic hashing with minimal multiplicative complexity," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2016, pp. 191–219.
- [13] S. Friebe, I. Sobik, and M. Zitterbart, "Decentid: Decentralized and privacy-preserving identity storage system using smart contracts," in *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications*. IEEE, 2018, pp. 37–42.
- [14] F. Baldimtsi, K. K. Chalkias, Y. Ji, J. Lindstrøm, D. Maram, B. Riva, A. Roy, M. Sedaghat, and J. Wang, "zklogin: Privacy-preserving blockchain authentication with existing credentials," *arXiv preprint arXiv:2401.11735*, 2024.
- [15] S. Motepalli, L. Freitas, and B. Livshits, "Sok: Decentralized sequencers for rollups," *arXiv preprint arXiv:2310.03616*, 2023.
- [16] J. Ernstberger, J. Lauinger, Y. Wu, A. Gervais, and S. Steinhorst, "Origo: Proving provenance of sensitive data with constant communication," *Cryptology ePrint Archive*, 2024.
- [17] J. Lauinger, J. Ernstberger, A. Finkensteller, and S. Steinhorst, "Janus: Fast privacy-preserving data provenance for tls 1.3," *Cryptology ePrint Archive*, 2023.