

# Smart contract with secret parameters

Marin Thiercelin <sup>\*†</sup>    Chen-Mou Cheng <sup>\*</sup>    Atsuko Miyaji <sup>\*</sup>    Serge Vaudenay <sup>†</sup>

**Abstract:** By design, smart contracts' data and computations are public to all participants. In this paper, we study how to create smart contracts with parameters that need to stay secret. We propose a way to keep some of the parameters off-chain, while guaranteeing correctness of the computation, using a combination of a commitment scheme and a zero-knowledge proof system. We describe an implementation of our construction, based on ethereum smart contracts and zk-SNARKS. We also provide a small example and a cost analysis of our approach.

**Keywords:** smart-contract, privacy, decentralized, zero-knowledge, ethereum

## 1 Introduction

The aim of smart contract systems is to secure relationships, over computer networks, without the needs of legal third parties. This idea has been brought to reality by achievements in the fields of distributed systems and cryptography. Existing smart contract systems use a decentralized consensus protocol, which requires all the information and computations involved in the contract to be public. This specificity provides a strong accountability, since everything is logged and verified by the system. However, the public nature of smart contracts can also be a drawback, as many applications involve private information. In this work, we want to look at constructions that make it possible to use smart contracts of which some parameters need to stay secret, while preserving the correctness of the contracts.

## 2 Background

### 2.1 Notations

We present notations used in this paper.

1. We write the security parameter as  $\lambda$ .
2. We say a function  $\epsilon$  is negligible (in  $\lambda$ ), iff  $\forall c \in \mathbb{N}^*, |\epsilon(\lambda)| = \mathcal{O}(\frac{1}{\lambda^c})$ .
3. We say an algorithm is ppt, when it is a randomized algorithm with polynomial asymptotic complexity in  $\lambda$ .
4. We write  $x \leftarrow_s X$ , when we sample  $x$  uniformly in the set  $X$ . We write  $y \leftarrow A(x)$  when  $A$  is a randomized algorithm and  $y := A(x)$  when  $A$  is deterministic.
5. We write  $a := \langle b|c \rangle$  to define  $a$  as the concatenation of  $b$  and  $c$ , and  $\langle b|c \rangle := a$  when we decompose  $a$  into  $b$  and  $c$ .

<sup>\*</sup> Miyaji Laboratory, Osaka University

<sup>†</sup> LASEC, EPFL

6. We write the states of a smart contract as  $\mathbf{s}_i$ . We write  $\mathbf{s}_i[x]$  when we want to access variable  $x$  in the state  $\mathbf{s}_i$ .
7. We say a computation is executed onchain when it is executed by all the nodes of the blockchain. We say that it is executed off-chain when it is executed locally by some node.

### 2.2 Blockchain

A blockchain is a decentralized system, where the decision power is shared by all the parties involved (also called nodes or peers). A blockchain is used by peers to settle on a shared, sequential state history. For a fixed period of time, the nodes broadcast a set of changes, called transactions, that they want to see in the next state. At the end of the period, using a consensus algorithm, the nodes decide on a set of state changes, called a block, and add it to their state history (which forms a chain of blocks, hence the name). Blockchains use cryptography and distributed systems to provide strong guarantees that the blockchain state cannot be corrupted, and it is impossible to change the state history retroactively. Among other things, a blockchain can be used to power a digital currency [1]. Peers sending transactions are identified using their public keys, and digital signatures are used to make sure that only authorized transactions are accepted by the blockchain. Another important property of blockchain is that a new block is computed at almost regular time intervals, thus the block counter can be used as a clock by applications.

### 2.3 Smart contract

As discussed in Introduction, a smart contract is to be used as a contract in an online and decentralized setting. It needs to secure relationships without the use of trusted authorities or legal systems. In practice, smart contract systems are implemented in a programming language on top of a blockchain protocol. The smart

contract is defined as a contract state and a piece of computer code that is included (we say that it is deployed) in the blockchain state. When a transaction is sent to the contract, the code is executed and produces a new contract state, which is included in the blockchain. More interestingly, smart contracts can be programmed to hold and transfer digital assets.

## 2.4 Commitment scheme

A commitment scheme is a cryptographic primitive, which can be used in a protocol to force a party to commit to a secret value before continuing the protocol.

**Definition 1** (Commitment scheme). *A commitment scheme consists of three efficient algorithms:*

- *Setup*: randomized algorithm, takes in the security parameter and outputs system parameters  $\sigma$ .
- *Commit*: randomized algorithm, given a value  $v \in V$ , it returns a commitment  $c$  and a key  $k$ .
- *Open*: given a commitment  $c$  and key  $k$  outputs a value in  $V$  or  $\perp$  (meaning an incorrect opening).

*They must satisfy the following properties.*

- *Correctness*: Let  $\sigma \leftarrow \text{Setup}(1^\lambda)$ ,  $\forall v \in V$ , if  $(c, k) \leftarrow \text{Commit}_\sigma(v)$ , then  $\text{Open}_\sigma(c, k) = v$ .
- *Hiding*: Any ppt adversary, choosing  $v_0, v_1 \in V$  and given a commitment  $c$  of value  $v_b$  with  $b \leftarrow_{\$} \{0, 1\}$ , cannot output a bit  $b' = b$  with probability significantly larger than  $\frac{1}{2}$ .
- *Biding*: Any ppt adversary can output  $(c, k, k')$  such that  $\perp \neq \text{Open}_\sigma(c, k) \neq \text{Open}_\sigma(c, k') \neq \perp$  only with negligible probability.

## 2.5 Proof system

A proof scheme is used in a protocol with 2 parties, the prover and the verifier, where the prover wants to convince the verifier that a statement is true.

**Definition 2** (Non-interactive proof-of-knowledge). *Let  $R(a, w)$  be an efficiently computable boolean function, and  $L = \{a \mid \exists w, R(a, w) = 1\}$  its associated language. If  $R(a, w) = 1$ , we say that  $w$  is a witness of the statement  $a \in L$ . For such  $L$ , a non-interactive proof-of-knowledge system consists of 3 efficient algorithms:*

- *Setup*: randomized algorithm, takes in the security parameter and outputs system parameters  $\eta$ ;
- *Prove*: randomized algorithm, given a statement  $a$  and witness  $w$ , it returns a proof  $\pi$ ; and
- *Verify*: given a statement  $a$  and a proof  $\pi$ , it returns a bit.

*They must satisfy the following properties.*

- *Completeness*:  $\forall a \in L$ , if  $\eta \leftarrow \text{Setup}(1^\lambda)$ ,  $R(a, w) = 1$  and  $\pi \leftarrow \text{Prove}_\eta(a, w)$ , then  $\text{Verify}_\eta(a, \pi) = 1$ .

- *Soundness*:  $\forall a \notin L$ , any ppt adversary produces a proof  $\pi'$  such that  $\text{Verify}_\eta(a, \pi') = 1$  only with negligible probability.

**Zero-knowledge proof** In our construction, we will require as an additional property that the proof be zero-knowledge. Informally, the zero-knowledge property means that any efficient adversary, given a correct proof  $\pi$ , only learns that  $a \in L$  is true, and specifically, doesn't learn anything about the witness  $w$ .

## 2.6 Related work

In 2016, the Hawk team proposed a construction for privacy-preserving smart contracts [2]. Their aim was to achieve transaction privacy, where the transaction data (source, destination, value, etc.) is not disclosed. They use a special party, called the manager, to run the computation, as well as a zero-knowledge proof system to enforce correctness.

Another line of work has been using zero-knowledge proofs to introduce privacy in blockchain transactions; for example, the zerocash protocol [3] extends the bitcoin protocol, with hidden transactions, where the amount, sender and receiver are not shared with anyone else.

In the Zexe paper [4], an extension of the zerocash protocol is proposed, which allows programmable coins. Zexe successfully realizes a primitive they named *decentralized private computation*, that can be useful in applications such as private digital currency exchanges.

## 2.7 Our contribution

We propose a construction that allows smart contracts be used for applications with private parameters. Our solution works directly with existing smart contract systems. We implement our construction with ethereum smart contracts in a software library. The rest of this paper is organized as follows. In Section 3, we formally define the problem we want to solve and give a small illustrative example. In Section 4, we propose a construction for a smart contract with private parameters that satisfies the requirements of Section 3. In Section 5, we describe an implementation of our construction using ethereum smart contracts. In Section 6, we analyze the performance of our implementation. Finally in Section 7, we discuss the drawbacks of our construction and evoke some possible directions to address them.

## 3 Problem definition

iiiiiii **HEAD** In this paper, we study applications defined by an initial state  $\mathbf{s}_0$ , and a transition function *IdealTransition*. A special party, the owner, chooses a secret value *secret*, which is used as a parameter by the transition function. When other parties, the users, input the application with some data *data*, the next state of the application is computed as  $\mathbf{s}_{i+1} = \text{IdealTransition}_{\text{secret}}(\mathbf{s}_i, \text{data})$ . As we said in the introduction, these applications can't be implemented using smart contracts alone, as their public nature would

leak the value *secret*. In this paper, we looked at a relaxed problem, where we allow an implementation to enter an intermediate state  $s'_i$ , and eventually reach the correct state  $s_{i+1}$ . We can break down *IdealTransition* into subroutines  $f, g, h, k$ , as in algorithm 2.

---

**Algorithm 1** *IdealTransition<sub>secret</sub>(s<sub>i</sub>, data)*

---

```

 $s'_i := k(s_i, data)$ 
 $z := h(s'_i)$ 
 $y := f(secret, z)$ 
return  $(s'_i, y)$ 

```

---

Subroutines  $f, g, h, k$  have distinct roles:

- $k$  includes all the computations independent of *secret*, and produces an intermediate contract state  $s'_i$
- $h$  extract the input in a format expected by  $f$
- $f$  is a deterministic function of the value *secret* and the value  $z$  produced by  $h$
- $g$  links everything and computes the next state

Now we can formally define the security we want to achieve:

**Definition 3.** We say that a protocol  $P$  securely realizes *IdealTransition<sub>secret</sub>*, if it satisfies the following requirements:

1. No efficient adversary should gain any knowledge on *secret*, other than the values  $f(secret, h(s'_i))$ , with  $s'_i := k(s_i, data)$ , of each previous state transition
2. For any user action with input data at state  $s_i$ ,  $P$  eventually gets to a new state  $s_{i+1} = \text{IdealTransition}_{secret}(s_i, data)$

Example 2 shows a simple application involving a secret parameter.

**Example 1** ("Higher or Lower?" game). Here we have two parties, the owner and the player. The game is parameterized by a set  $S \subset \mathbb{N}$ , and a number of query  $t$ . The owner chooses  $s \leftarrow S$  and let the player make some guesses, and on each guess  $g \in S$ , returns whether  $s$  is a number higher, lower or equal to  $g$ . After  $t$  guesses, the player makes a final guess, and wins if it is a correct guess. We've described this game in Figure 1.

If we can build a smart contract with private parameters that satisfies the requirements, then we can realize the game in the following way. Set the secret parameter *secret*, to the value to guess  $s$ . We let the state of the contract hold a counter of queries  $cnt = 1$ . We define  $f, g, h, k$  as in Figure 2.

This smart contract can act as a "Higher/Lower" oracle, and designate the winner, with a potential reward. ===== In this paper, we study applications defined with two phases. In an initial phase, a

special party, called the owner, chooses a parameter *secret* and an initial state  $s_0$ . In the second phase, other parties, the users, can participate by choosing an input *data*, and update the application state:  $s_{i+1} = \text{IdealTransition}_{secret}(s_i, data)$ . Smart contracts can't implement these applications trivially, as their public nature would leak the value of *secret*. To overcome this, we look at a relaxed problem, where a smart contract realizing such applications is allowed to enter an intermediate state  $s'_i$ , and we only require that it eventually (after some fixed amount of time) enters the correct state  $s_{i+1}$ . Before formally defining our requirements, we break down the specification of *IdealTransition* into subroutines  $f, g, h, k$  as shown in Algorithm 2.

---

**Algorithm 2** *IdealTransition<sub>secret</sub>(s<sub>i</sub>, data)*

---

```

 $s'_i := k(s_i, data)$ 
 $z := h(s'_i)$ 
 $y := f(secret, z)$ 
return  $g(s'_i, y)$ 

```

---

Each subroutine has a different role.

- $k$  includes all the computations independent of *secret* and produces an intermediate state  $s'_i$ .
- $h$  extracts (part of) the intermediate state in a format expected by  $f$ .
- $f$  is a deterministic function of the value *secret* and the value  $z$  produced by  $h$ .
- $g$  links everything and computes the next state.

Now we formally define what we want to achieve.

**Definition 4.** We say that a protocol  $P$  securely implements *IdealTransition* if it satisfies the following requirements.

1. No efficient adversary should be able to extract any knowledge on *secret* from the transcript of  $P$ , other than what can be deduced from the values  $f(secret, h(k(s_i, data)))$ , for all previous transitions with input data and state  $s_i$ .
2. For any transition with input data at state  $s_i$ , the application eventually enters the state  $s_{i+1} = \text{IdealTransition}_{secret}(s_i, data)$ .

To illustrate, we provide a small game in Example 2, which can be implemented using a system satisfying Definition 4.

**Example 2** (Higher or Lower? game). Here we have two parties, the owner and the player. The game is parameterized by a set  $S \subset \mathbb{N}$ , and a maximum number of guesses  $t$ . The owner chooses  $s \leftarrow S$  and lets the player make some guesses. On each guess  $g \in S$ , the owner returns whether  $s$  is a number higher, lower, or equal to  $g$ . After  $t$  guesses, the player makes a final guess and wins if it is correct, as shown in Figure 1.




Figure 1: "Higher or Lower?" game

In Example 2, we can set *secret* to be the value to guess  $s$ . We set the initial state to hold a counter of queries  $\text{cnt} := 1$  and write the transition function with subroutines  $f, g, h, k$  as in Figure 2. If the implementa-

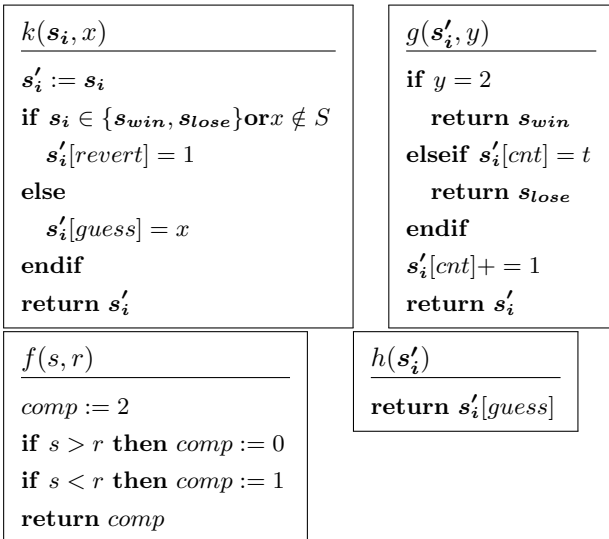


Figure 2: Subroutines for the Higher/Lower game

tion satisfies Definition 4, then we are assured that only the value  $f(s, q_i)$  is leaked for each guess  $q_i$ , i.e., we only leak whether the user's guess was higher or lower than the secret value. 7bdf9183b9bc96f1b4246ace02cee780fb297a28

## 4 Construction

In this section, we propose a construction that meets the requirements of Definition 4. We will do so in an incremental fashion: first we give constructions that satisfy the requirements using some assumptions, which will be refined by adopting more realistic assumptions at each step.

### 4.1 Naive approaches

#### 4.1.1 Regular smart contract

In regular smart contract systems, all parameters (including *secret*) are directly included in the binary of the smart contract. In that case, when a transaction gives input *data* to the contract, the blockchain nodes have all the information necessary to compute the next state  $s_{i+1} := g(s'_i, f(\text{secret}, h(s'_i)))$  with  $s'_i = k(s_i, \text{data})$ , so we say that  $f$  is computed on-chain. The incorruptible nature of regular smart contracts automatically satisfies Requirement 2. However, Requirement 1 is satisfied if we make the assumption that no other parties will look at the value of *secret*, which is accessible publicly in the binary of the contract. This assumption is really strong and unrealistic, especially if knowing *secret* can provide some financial gains (e.g., win a reward in Example 2).

#### 4.1.2 Off-chain computation

If we want to protect against malicious users, then the parameter *secret* can't be included directly in the contract. A natural approach is to have the owner keep the value *secret* locally and publish the smart contract without *secret*. When a user sends a transaction with data *data*, the contract gets to the new state  $s'_i := k(s_i, \text{data})$ . The owner monitors the state changes of the contract. When the smart contract gets to a state  $s'_i$ , the owner computes  $z := h(s'_i)$  and  $y := f(\text{secret}, z)$  locally (off-chain) and sends  $y$  to the smart contract. The new state is then computed as  $s_{i+1} := g(s'_i, y)$ . Here Requirement 1 is satisfied by design. However, Requirement 2 is satisfied if we assume that the owner is honest. This assumption is arguably less strong (than no other parties will peek), as in many cases the owner will be a company with a reputation and regulations and hence will be less tempted to cheat the protocol. However, this assumptions reduces the application to a client-server application, with the smart contract only acting as a middleman that relays messages and transfers assets.

### 4.2 Our proposal: off-chain and verified

If we want to protect against both malicious users and owners, then the smart contract (and the nodes of the underlying blockchain) should be able to verify that the value  $y$  returned by the owner really corresponds to  $f(\text{secret}, h(s'_i))$ . To achieve that, we use a commitment scheme and a zero-knowledge proof system for the language:

$$L_f = \{a =: \langle z | c | y \rangle \mid \exists w =: \langle sk | k \rangle, \text{Open}_{\sigma}(c, k) = sk \cap f(sk, z) = y\}.$$

Our construction is parameterized by two functions  $l_{\text{fake}}$  and  $l_{\text{timeout}}$ , used to penalize a malicious owner, and a timeout value  $T$ . We assume that the *Setup* algorithm of both the commitment scheme and proof system was run honestly, and the system parameters  $\sigma$  and  $\eta$  are known by all parties.

**Initialisation:** On deployment, the owner chooses its secret parameter  $secret$ , computes a commitment  $(c, k) \leftarrow \text{Commit}_\sigma(secret)$ , includes  $c$  in the code of the smart contract, and keeps  $k$  locally.

**Start transition:** When a user sends a transaction with data  $data$ , the contract gets to the new state  $\mathbf{s}'_i := k(\mathbf{s}_i, data)$ .

**Off-chain computation:** The owner monitors the state changes of the contract. When the smart contract get to a state  $\mathbf{s}'_i$ , the owner computes  $z := h(\mathbf{s}'_i)$ ,  $y := f(secret, z)$ , and  $\pi \leftarrow \text{Prove}_\eta(\langle z|c|y \rangle, \langle secret|k \rangle)$  to prove that  $\langle z|c|y \rangle \in L_f$ .

**Finish transition:** The owner then sends  $y$  and  $\pi$  to the smart contract. The smart contract computes  $z' := h(\mathbf{s}'_i)$ , reconstructs  $a' := \langle z'|c|y \rangle$  from the value  $c$  it stored and the  $y$  it received, and computes  $b := \text{Verify}_\sigma(a', \pi)$ . If  $b = 1$ , then the smart contract accepts  $y$  as the value  $f(secret, z')$ , and the next state is computed as  $\mathbf{s}_{i+1} = g(\mathbf{s}'_i, y)$ . Otherwise, the smart contract goes to the next state  $\mathbf{s}_{i+1} := l_{\text{fake}}(\mathbf{s}'_i)$ .

**Timeout:** If the contract is still in state  $\mathbf{s}'_i$  after  $T$  blocks have been generated, the user sends a *Timeout* transaction. When the contract receives it, it goes to a state  $\mathbf{s}_{i+1} := l_{\text{timeout}}(\mathbf{s}'_i)$ .

Figure ?? describes the structure of the smart contract, and Figure ?? describes the protocol as a whole.

## Security

**Proposition 1.** *For any rational, ppt adversary, we can select  $l_{\text{fake}}$  and  $l_{\text{timeout}}$  such that our protocol satisfies Definition 4.*

*Proof.* Definition 4 has two requirements.

### Requirement 1:

In this construction, the only values that are published by the owner are  $c$ ,  $y$ , and  $\pi$ . The hiding property of the commitment scheme guarantees that no information on  $secret$  is leaked by  $c$  against any ppt adversary. The zero-knowledge property of the proof scheme guarantees that no information on  $w := \langle secret|k \rangle$  is leaked against any ppt adversary. The only released information on  $secret$  is the value  $y := f(secret, h(k(\mathbf{s}_i, data)))$ ; therefore it satisfies Requirement 1.

### Requirement 2

If the owner follows protocol, then the completeness property of the proof system guarantees that the  $\text{Verify}$  procedure will return 1. Since  $b = 1$ , then  $\mathbf{s}_{i+1} := g(\mathbf{s}'_i, y)$ , it follows that Requirement 2 is satisfied. If the owner doesn't follow protocol, then we have to look at several cases.

1. If the owner simply doesn't respond, then the user sends a *Timeout* transaction after  $T$  blocks have been generated, and the the contract goes to state  $\mathbf{s}_{i+1} := l_{\text{timeout}}(\mathbf{s}'_i)$ .
2. If the owner sends  $y' \neq f(secret, z)$  and  $\pi'$ , then there are two cases:

- (a)  $\pi'$  is not a valid proof. Then the soundness property of the proof system guarantees that  $\text{Verify}_\eta(\langle z|c|y \rangle, \pi) = 0$ , and  $\mathbf{s}_{i+1} := l_{\text{fake}}(\mathbf{s}'_i)$ .
- (b)  $\pi'$  is a valid proof. Then the completeness property guarantees that  $\text{Verify}_\eta(\langle z|c|y \rangle, \pi) = 1$ , and  $\mathbf{s}_{i+1} := g(\mathbf{s}'_i, y')$

In all these three cases, the final state  $\mathbf{s}_{i+1}$  is different from the value returned by  $\text{IdealTransition}$ , hence Requirement 2 is not satisfied. To mitigate case 1 and 2a, we can set  $l_{\text{fake}}$  and  $l_{\text{timeout}}$  to include appropriate financial penalties for the owner in the new state, in order to deter a malicious but rational owner. In the case 2b, if  $\pi'$  is valid, then the owner knows a witness  $w := \langle secret'|k' \rangle$  for the statement  $\langle z|c|y' \rangle \in L_f$ . The definition of  $L_f$  implies  $\text{Open}_\sigma(c, k') = secret' \cap f(secret', z) = y'$ . Since  $y' \neq f(secret, z)$ , we have  $secret' \neq secret$  ( $f$  is deterministic). We also have  $\text{Open}_\sigma(c, k) = secret$  from the correctness of the initial commitment. It implies the owner knows  $(c, k, k')$  with  $secret = \text{Open}_\sigma(c, k) \neq \text{Open}_\sigma(c, k') = secret'$ , so the binding property of the commitment guarantees that this happens with negligible probability for a ppt adversary.  $\square$

As a result, our construction satisfies Definition 4 under the assumption that the *Setup* algorithms of the commitment scheme and proof system are executed correctly, for all rational ppt adversaries.

## 5 Implementation

In Section 4.2, we described a construction to build a smart contract that satisfies Definition 4. We made assumptions on the type of adversary, and we assumed a trusted setup for the commitment and proof system. Since our goal is to have a construction that can be implemented on existing smart contract systems, we have limited the assumptions we made about the capabilities of the contracts. To simplify, we will restrict our scope to applications where the secret parameter  $secret$  is a single 128-bit unsigned integer, and  $f$  is a function taking two 128-bit unsigned integers and returning another 128-bit unsigned integer.

### 5.1 Tools

Before we explain our implementation, we will introduce some of the tools we use.

#### 5.1.1 Ethereum smart contracts

We implement the on-chain part of our construction using ethereum smart contracts. Ethereum [5] is a commonly used public and programmable blockchain. It provides a set of computer instructions that can be used to write smart contracts, as well as a decentralized computing unit, the ethereum virtual machine (EVM), on top of which contracts are executed. Ethereum has a base currency, the ether (ETH), and uses a concept of accounts to associate a public key to an amount of

ether controlled by the key. In ethereum, smart contracts have their own accounts: users can transfer ether to contracts, and a contract can transfer ether from its account to users (or to other contracts). To avoid attackers using all of the resources of the EVM, ethereum uses a concept of gas. The gas is a fee that is included with every transaction. All resources (memory, computation, etc.) used for a transaction have a cost, which is taken away from the gas and given to the ethereum nodes. If a transaction runs out of gas, the computations are reversed, but the gas used is not returned.

**Solidity** The ethereum community also produced high-level programming languages that can be compiled to EVM programs. We use the Solidity language to define our contracts. Solidity has a syntax close to object-oriented programming languages: contracts have a internal variables (the state) and methods (computations), which can be executed by specific transactions (we say that someone calls the method). Solidity contracts can emit **events**, which notifies the users that something happened. Deployed contracts are identified by an address (a unique 256-bit string) and can call each other's methods using these addresses. Solidity contracts are defined in source files with extension `.sol`.

### 5.1.2 zk-SNARKS

As a proof system, we use zk-SNARKS [6], which are non-interactive zero-knowledge proofs-of-knowledge. Zk-SNARKS are secure in the common reference string (CRS) model, which means that they require a trusted setup. They have short proofs and fast verifications, which induces smaller transactions and reduced ethereum gas costs. More importantly, we can create zk-SNARKS to prove that a computation was correctly executed according to a fixed circuit, while not revealing part of the inputs (called the witness). We can even produce proofs for tinyRAM (a random access machine with a limited instruction set) programs [7]. This allows us to produce zero-knowledge proofs for high-level computations, (e.g., *Open* and *f* in our construction).

### 5.1.3 Zokrates

In Section 4.2, we assumed that the contract is able to compute  $b := \text{Verify}_\sigma(a', \pi)$ . Since *Verify* is not a native instruction of the EVM, we need to include its code as a part of the smart contract. We will use the Zokrates [8] toolbox for this purpose. Zokrates provides:

- A programming language to define computations with inputs that are specified as either **public** or **private**. Programs are written in files with extension `.zok`. Integers are limited to the type `field`, which represents values in  $\mathbb{Z}_p$  for some prime  $p$  specific to the zk-SNARKs. We use `field` to represent 128-bit unsigned integers, as  $p > 2^{128}$ .
- A *compiler* that reduces a Zokrates program to a zk-SNARK.

- A *setup* algorithm that produces proving and verification keys.
- A *prover* algorithm that takes in the proving key and the inputs and produces an output, along with a zero-knowledge proof. The statement to prove includes the **public** inputs and the output, and the witness includes the **private** inputs.
- An *export* algorithm that takes in the verification key and produces an ethereum smart contract, written in Solidity, that can verify the proofs.

### 5.1.4 Hash-Based Commitment

Our construction also needs a commitment scheme, with the specificity that the *Open* algorithm is part of the computation verified by the proof system, so we need to write the check  $\text{Open}(c, k) = \text{secret}$  as a Zokrates program. Zokrates has a limited library, but it includes the SHA256 hash function, which can be used to build a hash-based commitment scheme. We then define our commitment scheme as in Figure ???. The commitment  $c$  is a 256-bit string, and the key  $k$  is a 384-bit string. In Zokrates, we represent  $c$  with type `field[2]`, and  $k$  with type `field[3]`, by taking blocks of 128-bits for each `field`.

## 5.2 Library description

We have implemented our construction as a library [9], where all generic parts of the construction are implemented by the library, so the developer only has to implement the parts that are specific to the target application. The library includes smart contracts, written in Solidity, and Javascript programs to help the owner and users of the contract.

### 5.2.1 Contract Architecture

We want our solution to abstract the way  $f$  is computed, so that the developer has little more to do than defining  $f, g, h, k$  and  $l_{\text{fake}}, l_{\text{timeout}}, T$ . We first break the code of the contract into three contracts. A first contract, called the **requester** contract, is responsible for the computations of  $g, h, k, l_{\text{fake}}$  and  $l_{\text{timeout}}$ . A second contract, called the **holder** contract, is responsible for the interaction with the owner to compute  $f$ . The third contract, called the **verifier** contract, verifies the proofs sent by the owner. The three contracts are deployed together and know of each other's address.

**Requester contract** The requester has four methods: **start**, **callback**, **wrong-proof**, and **timeout**.

- **start(data)** is called by the user, which executes the pre- $f$  computation (i.e.,  $k$  and  $h$ ) and obtains a 128-bit unsigned integer  $z$ . It calls **holder.requestComputation(z)**, obtains a value  $id$ , and emits an event **Start(id)**.
- **callback(id, y)** is called by the **holder** with the value  $y := f(\text{secret}, z)$ , executes the post- $f$  computation  $g$ , and emits an event **End(id, result)** for some **result** generated by  $g$ .

- `wrong_proof(id)` is called by the holder, realizes  $l_{\text{fake}}$ , and emits an event `WrongProof(id)`.
- `timeout(id)` is called by the user, realizes  $l_{\text{timeout}}$ , and emits an event `Timeout(id)`.

**Verifier contract** The verifier has one method, `verifyTx`, that executes the Verify algorithm and returns a bit.

**Holder contract** The holder contract has an internal variable `c`, initialized with the commitment of the *secret* parameter chosen by the owner. The holder has two methods: `requestComputation` and `answerRequest`.

- `requestComputation(z)` is called by the requester, which generates a unique value `id`, stores `(id, z)` in a table, emits an event `NewRequest(id, z)`, and returns `id` to the requester.
- `answerRequest(id, y, pi)` is called by the owner, which retrieves the tuple `(id, z)` from its table and calls `verifier.verifyTx( $\langle z|c|y \rangle$ , pi)` to check that `y` is correct (i.e.,  $\langle z|c|y \rangle \in L_f$ ). If it is, then the contract calls `requester.callback(id, y)`; otherwise it calls `requester.wrong_proof(id)`.

The **holder** is the same for all  $f$ , so we can make its source code `holder.sol` an internal part of our library.

### 5.2.2 Local computations

Our library also includes programs to generate, deploy, and interact with the smart contracts.

**Setup program** We’ve written a small setup program that is executed by the trusted third party and produces the proving and verification key, for any given Zokrates program.

**Owner program** We require the owner to define  $f$  as a Zokrates program `F.zok` and define  $g, h, k$  and  $l_{\text{fake}}, l_{\text{timeout}}$  as part of the **requester** contract. We wrote a program for the owner with three components, as shown in Algorithm 3. Here `AddCommitCheck` modifies a Zokrates program `F.zok`, which computes  $f(\text{private } \text{secret}, \text{public } z)$ , to a new `F&C.zok` program that has additional inputs (`public c`, `private k`) and checks that  $\text{Open}(c, k) = \text{secret}$  before computing  $f$ . `Link` modifies the three deployed contracts, such that that they know each other’s addresses.

**User program** We defined a user program, as shown in Algorithm 4, used by the user to make calls to the **requester** and receive the corresponding `result` value.

### 5.3 Details

Here we have implemented a **requester** contract that can only access  $f$  in an asynchronous manner through a call back. This breaks the atomic nature of smart contract changes. The developer must take care in writing the contract to make sure that the `Timeout` and `WrongProof` methods revert what needs to be reverted in case of failure. The developer also needs to

---

#### Algorithm 3 Owner program

---

```

procedure GENERATE( $F$ )
   $F\&C \leftarrow \text{AddCommitCheck}(F)$ 
   $p.\text{key}, v.\text{key} \leftarrow \text{Trusted.Setup}(F)$ 
   $v.\text{sol} \leftarrow \text{Zokrates.export}(v.\text{key})$ 
  return  $F\&C, p.\text{key}, v.\text{key}, v.\text{sol}$ 
end procedure

procedure DEPLOY( $r.\text{sol}, v.\text{sol}, \text{secret}$ )
   $r.\text{addr} \leftarrow \text{Deploy}(r.\text{sol})$ 
   $v.\text{addr} \leftarrow \text{Deploy}(v.\text{sol})$ 
   $(c, k) \leftarrow \text{Commit}(\text{secret})$ 
   $d.\text{addr} \leftarrow \text{Deploy}(d.\text{sol}, c)$ 
   $\text{Link}(r.\text{addr}, v.\text{addr}, d.\text{addr})$ 
  return  $(c, k, r.\text{addr}, v.\text{addr}, d.\text{addr})$ 
end procedure

procedure LISTEN( $c, k, s, h.\text{addr}, F\&C, p.\text{key}$ )
   $\text{listener} \leftarrow h.\text{addr}.\text{listen}(\text{NewRequest})$ 
  while true do
     $id, z \leftarrow \text{listener.waitEvent}()$ 
     $y, pi \leftarrow \text{Zokrates.Prove}(F\&C, p.\text{key}, s, z, c, k)$ 
     $h.\text{addr}.\text{answerRequest}(id, y, pi)$ 
  end while
end procedure

```

---



---

#### Algorithm 4 User program

---

```

procedure MAKECALL( $r.\text{addr}, \text{input}, \text{nb\_block}$ )
   $id \leftarrow r.\text{addr}.\text{start}(\text{data})$ 
   $\text{listener} \leftarrow h.\text{addr}.\text{listen}(\text{End}(id, \_))$ 
   $\text{timeout}, \text{result} \leftarrow \text{listener.wait}(\text{nb\_block})$ 
  if  $\text{timeout} = \text{true}$  then
     $r.\text{addr}.\text{timeout}(id)$ 
    return TimeoutError
  else
    return result
  end if
end procedure

```

---

make sure `start` cannot be called a second time before the `callback` of the first call was executed to avoid interleaving of the calls. Another potential vulnerability is that the owner pays the gas used for the execution of `Verify` and `callback`, so malicious users could use that to drain the owner’s account. A solution is to write the **requester** contract in a way that any call to `start` needs to include a small reward in ether, which will be transferred to the owner when `callback` is called to cover the gas fee. The user also needs to check that the owner correctly deployed the contracts, and that the **verifier** contract was correctly generated from the verification key generated by the trusted third party.

## 6 Performance

To analyze the total cost of our solution, we have compared three constructions. In the first one, called *onchain* in the plots, all computations are done on-chain, which corresponds to the construction of Section 4.1.1. In the second one, called *unverified* in the plots,

the computation of  $f$  is done off-chain, but we don't use any verifications on the output, which corresponds to the construction of Section 4.1.2. The last one is our proposed construction, called *zokrates* in the plots, described in Section 4.2 and implemented as described in Section 5. We wrote a dummy **requester** contract and a dummy  $f : (secret, x) \rightarrow secret + x$ . We set up a private test network with a block period of two seconds, and a gas price of  $10^{14}$  wei (1 wei =  $10^{-18}$  ETH). We performed the following experiment: first we used the owner program to generate and deploy the contracts and started the listener. Then we executed the user program to make a call to **start** and to wait for the result. For each action, we measured the time spent (in seconds), as well as the gas cost (in wei) for each party (owner and user), and plot the results in Figure 3.

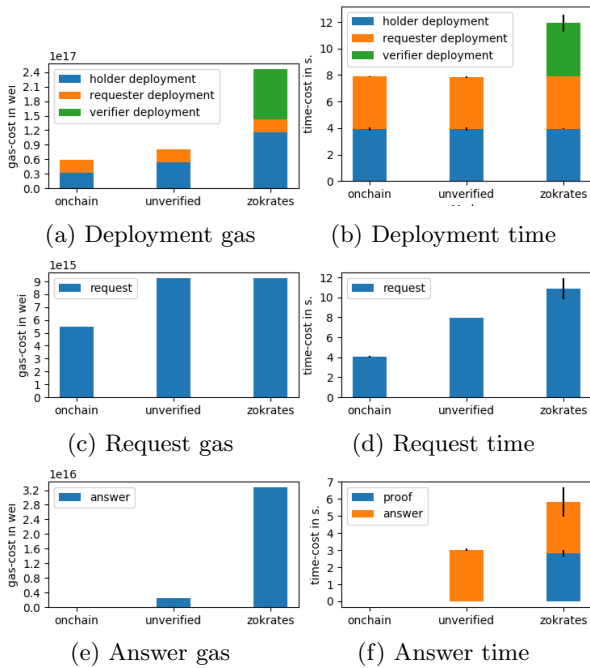


Figure 3: Performance comparison

Costs in Figures 3a, 3b, 3e, and 3f are costs paid by the owner, while costs in Figures 3c and 3d are costs paid by the user. We can see that our construction incurs a significant cost for the owner, among which the costs in Figure 3e can be subsidized by the user through a reward system, as discussed in Section 5.3.

## 7 Future improvements

The main drawback of our construction is that it assumes a trusted setup producing the proving and verification keys. Furthermore, this trusted setup needs to be performed for each new function  $f$  defined by the developer, as the zk-SNARKs we use are circuit-specific. A recent paper [10] has proposed constructions for universal zk-SNARKs, where one pair of keys can be used to verify the computation of any  $f$  (with a bounded circuit size). This means that the trusted setup would

have to be called only once, and the produced keys could be reused. Another solution would be to use zk-STARKS [11], which do not require any trusted setup. We would also like to explore other approaches, such as using secure multi-party computation protocols, obfuscation mechanisms, or secure hardware systems, e.g., trusted execution environments.

## 8 Conclusion

In this paper, we have proposed a construction that reconciles the public and incorruptible nature of a smart contract with the sensitivity of some private application data. We described a model of an ideal contract where some part of the contract's code is parameterized by a secret value that will not be leaked out during the execution of the contract. We then provided a construction that emulates this ideal contract, using a proof system and a commitment scheme. Our construction assumes that we can access a trusted setup for the proof system and is secure against rational ppt adversaries. We provided an implementation of our construction, using ethereum smart contracts and zk-SNARKS, which can be used as a library, providing a foundation for a broad set of applications.

**Acknowledgements** This work is partially supported by CREST (JPMJCR1404) at Japan Science and Technology Agency, enPiT (Education Network for Practical Information Technologies) at MEXT, and Innovation Platform for Society 5.0 at MEXT.

## References

- [1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [2] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou, "Hawk: The blockchain model of cryptography and privacy-preserving smart contracts," 2015.
- [3] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," 2014.
- [4] S. Bowe, A. Chiesa, M. Green, I. Miers, P. Mishra, and H. Wu, "Zexe: Enabling decentralized private computation," 2018.
- [5] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger,"
- [6] J. Groth, "On the size of pairing-based non-interactive arguments," 2016.
- [7] E. Ben-Sasson, A. Chiesa, D. Genkin, E. C. Tromer, and M. Virza, "Snarks for c: Verifying program executions succinctly and in zero knowledge," 2013.
- [8] J. Eberhardt and S. Tai, "Zokrates - scalable privacy-preserving off-chain computations," July 2018.
- [9] M. Thiercelin, "Offchainer library," 2019. <https://github.com/marinthiercelin/offChainer>.
- [10] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. Ward, "Marlin: Preprocessing zksnarks with universal and updatable srs," 2019.
- [11] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, "Scalable, transparent, and post-quantum secure computational integrity," 2018.