



## SmartInspect: Solidity Smart Contract Inspector

Santiago Bragagnolo, Henrique S C Rocha, Marcus Denker, Stéphane Ducasse

### ► To cite this version:

Santiago Bragagnolo, Henrique S C Rocha, Marcus Denker, Stéphane Ducasse. SmartInspect: Solidity Smart Contract Inspector. IWBOSE 2018 - 1st International Workshop on Blockchain Oriented Software Engineering, Mar 2018, Campobasso, Italy. 10.1109/IWBOSE.2018.8327566 . hal-01831075

**HAL Id: hal-01831075**

**<https://inria.hal.science/hal-01831075>**

Submitted on 5 Jul 2018

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# SmartInspect: Solidity Smart Contract Inspector

Santiago Bragagnolo\*, Henrique Rocha\*, Marcus Denker\*, and Stéphane Ducasse\*

\* Inria Lille - Nord Europe

Villeneuve D'ascq, France

{santiago.bragagnolo, henrique.rocha, marcus.denker, stephane.ducasse}@inria.fr

**Abstract**—Solidity is a language used for smart contracts on the Ethereum blockchain. Smart contracts are embedded procedures stored with the data they act upon. Debugging smart contracts is a really difficult task since once deployed, the code cannot be re-executed and inspecting a simple attribute is not easily possible because data is encoded. In this paper, we address the lack of inspectability of a deployed contract by analyzing contract state using decompilation techniques driven by the contract structure definition. Our solution, SmartInspect, also uses a mirror-based architecture to represent locally object responsible for the interpretation of the contract state. SmartInspect allows contract developers to better visualize and understand the contract stored state without needing to redeploy, nor develop any ad-hoc code.

**Index Terms**—Blockchain, Inspecting, Solidity, Smart Contracts

## I. INTRODUCTION

Blockchain technology has attracted a lot of attention recently [1]. A blockchain is a distributed database, managed by a peer-to-peer network that stores a list of blocks or records. Ethereum [2], and BitCoin [3] are examples of blockchain technologies. Blockchains can be used for many applications such as cryptocurrency, digital wallets, ad-hoc networks, remote transactions, among other uses [1]–[7]. One notable application of blockchain is the execution of smart contracts [8].

Smart contracts are what embedded procedures are for databases: programs executed in the blockchain to manage and transfer digital assets. When used in platforms like Ethereum, the contract language is Turing-complete [9]. Therefore, smart contracts can be used in many different scenarios. For example, there are smart contracts employed to subcurrency [10], and outsourced computation [1]. Solidity [10] is the predominant programming language used to specify smart contracts on the Ethereum blockchain platform.

Smart contracts define a data structure as well as the operations used to interact with these data [10]. Far from a typical database, where the primary representation is data, and the available operations are about the structure and the content, the principal element of an Ethereum database is not just the data. In addition, the database stores the behavior provided to interact with these data, and to trigger other behaviors, by sending messages to other contracts. Ethereum is a database that works as a stored environment of contract instance (objects). Compiled versions of the contract instances are then published as part of transactions to the blockchain.

One of the challenges faced by developers of smart contracts is finding and fixing bugs. Indeed, contracts are opaque in the sense that once deployed in the blockchain it is difficult to

access the value of a given contract attribute. In this paper, we focus on inspecting a smart contract state as a first step to support contract debugging.

The difficulty to inspect contract data is not a widely known problem. Although there are many tools for traditional databases to access its stored data, Ethereum and Solidity provide no such tool to inspect contract information. On the other hand, there are two practices that we can use to access contract data: (i) introducing getter methods, which requires the redeployment of the contract (if it is already running in the blockchain) and a possible data conversion (if the type is not supported as return); and (ii) using the API to acquire raw data and applying an ad-hoc decoding of the content.

Both described practices are tedious time-consuming tasks for a developer. By contrast, nowadays any programming language offers simpler ways to inspect data. Developers use such inspection to access run-time data during development or maintenance activities. Similarly, developers could benefit from smart contract run-time inspection to verify the currently stored data. Moreover, from a business perspective, companies could use contract inspection to help clients better understand the information that is actually stored in the contract. In fact, the UTOCAT<sup>1</sup> company deemed this interaction with clients as an important scenario, regarding the complexity to explain and understand Ethereum technology possibilities.

As a solution, we propose SmartInspect, an inspector based on pluggable property reflection. The main idea is that the binary structure of the contract is decompiled using a memory layout reification. The memory layout reification is built from the Solidity source code. Our SmartInspect architecture is based on decompilation capabilities encapsulated in mirrors [11]. Such mirrors are automatically generated from an analysis of Solidity source code. This approach allows us to access unstructured information from a deployed contract in a structured way. Therefore, our SmartInspect approach can introspect the current state of a smart contract instance without needing to redeploy it nor develop additional code for decoding.

The remainder of this paper is organized as follows. Section II starts with an example of a smart contract. In Section III, we detail the main problem our proposed approach aims to address. Section IV describes our proposed solution, SmartInspect. Section V shows a preliminary evaluation of the SmartInspect approach by comparing to other practices.

<sup>1</sup><https://www.utocat.com/en/>, verified 2018-02-22.

Section VI provides a brief discussion on contract inspection. In Section VII, we describe the related work. Finally, Section VIII presents our conclusions and outlines possible future work ideas.

## II. SMART CONTRACT BY EXAMPLE

In this section, we present an example of a smart contract written in Solidity (Section II-A), and we also describe a client application in Pharo Smalltalk to interact with it (Section II-B). We use this example throughout the paper to explain the opacity problem.

### A. Poll Smart Contract

In this example, the contract manages a poll where users are allowed to vote a single time. Only the contract owner is allowed to modify the list of voters. The poll is managed with a contract because it is used for management decisions that rely on the veracity of the information.

The following listing contains the code of this contract:

Listing 1. Solidity Poll Contract Example

```
1 pragma solidity ^0.4.16;
2
3 contract Public3StatesPoll {
4     /* Type Definition */
5     enum Choice { POSITIVE, NEGATIVE, NEUTRAL }
6     struct PollEntry { address user; Choice
7         choice; bool hasVoted; }
8
9     /* Properties */
10    PollEntry[] pollTable;
11    address owner;
```

This contract defines two user types Choice (line 5) and PollEntry (line 6). A Choice models the answers to the poll (whether the vote was *positive*, *negative* or *neutral*). A PollEntry is a record representing a vote, *i.e.*, the voting user, the selected option, and if he/she has voted or not. Note that to refer to the user we need an account address (using the primitive type address) that refers to an external account.

The contract stores internally a poll table (an array of PollEntry) (line 9) and an address to the contract's owner account (line 10). The poll table is an empty array where the contract owner will eventually store the poll information (*i.e.*, the array will have an entry for each user that is allowed to vote). The contract owner's address is used for security checks.

```
11 /* Constructor */
12 function Public3StatesPoll () {
13     owner = msg.sender;
14 }
```

Lines 11-14 define the contract constructor. This constructor is executed when the contract is deployed in the blockchain. It keeps track of the user who owns the smart contract for future reference.

```
15 function isRegistered (address voterAccount)
16     returns (bool) {
17     return (voterIndex (voterAccount) > -1);
18 }
```

```
18
19 function voterIndex (address voterAccount)
20     returns (int) {
21     for (uint x = 0; x < pollTable.length;
22         x++) {
23         if (pollTable[x].user == voterAccount)
24             return int(x);
25     }
26     return -1;
27 }
```

We define the helper function voterIndex (lines 19-26), which returns the index of the voter in the poll table. We also created the function isRegistered (lines 15-17) to determine whether the user was registered to vote by using the voterIndex function. Since array indexes in Solidity are unsigned integers (uint), we need to explicitly convert it to a regular integer (line 22).

```
27 function addVoter(address voterAccount)
28     returns (uint) {
29     assert( owner == msg.sender );
30     assert( !isRegistered(voterAccount) );
31     pollTable.push(PollEntry(voterAccount,
32         Choice.NEUTRAL, false));
33     return pollTable.length -1;
34 }
35
36 function vote (Choice choice) {
37     assert( isRegistered(msg.sender) );
38     uint index = uint(voterIndex(msg.sender));
39     assert( !pollTable[index].hasVoted );
40     pollTable[index].choice = choice;
41     pollTable[index].hasVoted = true;
42 }
43
44 function votesFor(Choice choice) returns
45     (uint) {
46     uint votes = 0;
47     for (uint x = 0; x < pollTable.length;
48         x++) {
49         if (pollTable[x].hasVoted &&
50             pollTable[x].choice == choice)
51             votes = votes +1;
52     }
53     return votes;
54 }
55
56 function allParticipantsHaveVoted () returns
57     (bool) {
58     for(uint x = 0; x < pollTable.length;
59         x++) {
60         if (!pollTable[x].hasVoted) return
61             false;
62     }
63     return true;
64 }
65 } //end of contract
```

The rest of the contract defines the following functions:

- *addVoter* (lines 27-32). This function registers a voter into the poll table. It tries to assert<sup>2</sup> that the caller is the contract owner and the voter is not already registered.

<sup>2</sup>The *assert* function checks for a condition and throws an exception if such condition is not met. In Solidity, exceptions undo all changes made in the invoked method.

- *vote* (lines 34-40). This function assigns the given choice to the entry related to the calling user. The user must be registered and not voted yet.
- *votesFor* (lines 42-49). It returns the number of users that voted for the given choice.
- *allParticipantsHaveVoted* (lines 51-56). It returns true if all the registered users have voted.

## B. Client Side

Once our poll contract is deployed in the blockchain, we need a client application to interact with it. For example, we can implement a web application providing a user interface or a web service as a means to invoke the functions in the contract and vote or get the poll results. The following listing illustrates the code of a Poll class implemented in Pharo for our client application that will act as a façade to our contract:

Listing 2. Client side for the Voter Smart Contract

```
1 Object subclass: #Poll
2   instanceVariableNames: 'deployedContract'
3   package: #PollContract.
```

Lines 1-3 declare a Poll class with a deployedContract instance variable. This instance variable refers to a proxy to the deployed contract. This is a common implementation used for Ethereum clients coded in other languages.

```
4 Poll class>> config
5   ^{#fromAccount →self systemAccount.
6     #gas →30000. #etc}.
7
8 Poll class>> deployNewContract: src
9   accounts: accounts
10  connection: conn
11
12  deployedContract := conn deploy: source
13    configuration: self config.
14  accounts do: [ :account |
15    deployedContract addVoter: account
16      configuration: self config ].
```

The class method<sup>3</sup> `deployNewContract:accounts:connection:` (lines 8-15) receives the contract source code, a list of user accounts that are allowed to vote and a connection to the blockchain. It first deploys the contract in the blockchain using the contract source code, and then calls the `addVoter:configuration:` function of the new contract for each of the given user accounts. When we call functions from a deployed contract, we need to provide configuration information as well.

```
16 Poll>> config: usr
17   ^ {#fromAccount →usr account.
18     #gas →30000. #etc}.
19
20 Poll>> user: usr votes: aValue
21   deployedContract vote: aValue configuration:
22     (self config: usr).
```

<sup>3</sup>A class method is comparable to a static method in the Java jargon. In Pharo method call with multiple parameters place arguments in between the method name. Hence `this.fooBar(arg1, arg2)` is expressed as `this foo: arg1 bar: arg2`.

```
22
23 Poll>> isFinished
24   ^ deployedContract allParticipantsHaveVoted
25     configuration: (self class config).
26
27 Poll>> results
28   ^ { #POSITIVE . #NEGATIVE . #NEUTRAL }
29     collect: [ :value | value →
30       deployedContract votesFor: value
31       configuration: (self class config) ].
```

The method `config:` (lines 16-18) provides configuration data with the user's account. The method `user:votes:` (lines 20-22) invokes the function `vote:configuration:` from the contract using the user's configuration. Likewise, the method `isFinished` (lines 23-25) invokes the function `allParticipantsHaveVoted()` of the contract. The method `results` (lines 26-28) invokes repeatedly `votesFor()` for each of the contract choices and returns a map relating each choice to the number of people that voted for it.

It is noteworthy that the Poll class works as a thin layer over the remote contract performing remote calls to it. In the scenario of a real business application, this layer may define the complete process of a large business and its logic.

## III. THE PROBLEM: CONTRACT OPAQUENESS

Contrary to traditional SQL databases such as Oracle or PostgreSQL which have a multitude of tools (e.g., DBeaver, Navicat, SQL Maestro, Toad, PgAdmin, etc.) to access the database schema and the *actual* data stored in a given column or row, Ethereum/Solidity does not provide any tool to inspect contract state in the referential model of the application. Since the contract is an arbitrary data type, the offered API to interact and inspect is both restricted and at a low-level of abstraction.

Contract state is *read-only* in the sense that unauthorized clients cannot interact with it. Finally, contract state is *opaque*: since it is encoded there is no simple way for a software developer to know the actual value of a contract specific attribute effectively stored in the blockchain. The remote architecture of deployed solutions should also be taken into account.

### A. Contract Remote Structure

A contract is stored in the blockchain database and its object representation can be accessed in the application client layer. The Ethereum platform employs proxies based on the contract ABI<sup>4</sup> for covering the gap introduced by the physical location of the object (which is stored remotely in the blockchain). For example, when we invoke the method `vote: configuration:` in a Pharo client, the method call will be sent to a proxy that will connect through RPC (Remote Procedure Call) to the remote ABI object. The result of this method is a transaction receipt hash and, if applicable, the client will also receive any returned values of the method call. It is noteworthy that method calls to blockchain objects may have a transaction cost related with them.

<sup>4</sup>Application Binary Interface (ABI) is the Ethereum standard to interact with contracts. This standard encodes contract data according to its specification.

By proxying the remote contract, a client application can use the contract methods just as any other object. Moreover, the client can activate methods that will be executed elsewhere in the blockchain (by paying the transaction cost if applicable) and it will return values that we can use (as any other values for other methods and objects). This simple way to interact with the contract is unsatisfactory, since the objects cannot be inspected. Since there is no simple way to access contract properties, it makes debugging session tedious or even impossible.

### B. Opaqueness problem example

We use the poll contract example (Section II) to illustrate the opaqueness problem. Let's suppose that a new user arrives a couple of days before the poll expiration date. When he tries to vote, the client system executes a routine that calls the contract's vote function and reports whether the call was successful (Listing 3). More specifically, the routine calls our method `user:votes:` defined in the Poll class (Listing 2, lines 16-17) that uses a proxy to remote call the vote function in the deployed contract.

Listing 3. Client voting routine

```
1 UserSession >> vote: aValue
2 transactionReceipt := poll user: user votes:
  aValue.
3 transactionReceipt
4 onSuccess: [ :t | self informToUser ];
5 onError: [ :e | self informError: e ].
```

The call will fail because the user was not a registered voter. We can see in the client code that the only point where there is a setup of users in the contract is during its deployment (Listing 2, lines 8-15). Therefore, the contract will not encounter the user and it will throw an exception. In the client, the details that caused the exception will be hidden, it will only know that the invoked method failed. Moreover, since the error is being thrown by the remote object, inspecting the contract code could identify the problem. However, a regular user does not have access to the contract code, only the people in charge of the contract have such access. For this reason, the user's only option is to submit a bug report stating that he cannot vote.

Listing 4. Contract vote function highlighting the asserts

```
1 function vote (Choice choice) {
2   assert( isRegistered(msg.sender) );
3   uint index = uint(voterIndex(msg.sender));
4   assert( !pollTable[index].hasVoted );
5   pollTable[index].choice = choice;
6   pollTable[index].hasVoted = true;
7 }
```

The person in charge (let's call him Bob) of solving this issue, will go and check the contract code (Listing 4), and deduce that there are two possible reasons for the code to fail: (i) the user is not authorized, i.e., he/she was not registered into the contract; or (ii) the user already voted.

However, Bob cannot know for certain what caused the issue without analyzing the contract data. In this specific case, since we know the problem was caused by the unregistered voter, the

easiest solution would be to change the contract state manually by calling the function `addVoter` to add the new user.

Bob will face many difficulties to find the issue. As we can see, we did not define the contract properties as `public`<sup>5</sup> (Listing 1, lines 9-10). Therefore, if Bob wants to find the nature of the error, he will need the contract instance's current state to inspect it. Bob has two possibilities then: (i) re-instantiate the contract to add a new function to return the data (i.e., create a getter method); or (ii) develop a costly, ad-hoc decoder for reading the binary content of the contract.

If Bob takes the first possibility, he will add a new function to analyze the content of the contract. Since Solidity functions cannot return arrays or structs, Bob will need to adapt its function accordingly to acquire the poll data. Moreover, the contract will have to be redeployed, creating a new instance of it. Therefore, Bob will be able to analyze the new instance data with his function, but not the previous one (which was the one that presented the issue). Besides, there are the transactional costs to redeploy the contract to be considered, as well as the inconvenience to ask the users to vote again (since it is a new instance). In our example, where there is only a few days left to close the poll, it would not be feasible for Bob to ask all users to vote again.

The second possibility is for Bob to spend time into creating an add-hoc decoder. The main advantage on this possibility is that Bob does not need to redeploy the contract. The decoder could access the complex binary slots of the contract's related storage and converted them into the desired content that Bob is trying to analyze. Since the Solidity documentation for its binary encoding is incomplete, Bob will have a difficult time to create the decoder. Moreover, this decoder is a one time solution, as it is designed for a specific data in a particular contract, i.e., Bob will not be able to easily reuse this solution to another contract. In our example, Bob might not have the time required to create such decoder before the poll expiration date.

This scenario illustrates a simple aspect of the impact of the opacity of contracts. The general concern is that the developer should be able to understand the value of a given contract attribute.

### C. Challenges

There are many challenges to pierce through the opaqueness problem and reveal contract information.

- **Binary and incomplete specification.** From the technical aspects we only have the Ethereum API to access a binary representation of the contract. The first challenge we faced is an *incomplete* specification of the contract encoding performed by the Solidity compiler [10].

<sup>5</sup>In Solidity, when a contract attribute (state variable) does not specify its visibility, it assumes the default "internal". Internal members can only be accessed from within the current contract or contracts deriving from it. It is noteworthy that everything inside a contract can be visualized by external users. Restricted visibilities (e.g., internal, private) only prevents other contracts from modifying the information [10].

- **Inconsistent specification of hash computation.** Another challenge related to the specification is the hash computation for dynamic types. Static types use text as input for the hash calculation. However, dynamic types follow a different standard that it is not clearly specified in the documentation. For dynamic types, it is necessary to use binary data packed specifically for each type. For example, to access an array it is necessary to pack the index number and the array position (offset) into a binary representation to obtain the correct hash. Other dynamic types would require a different input to get its hash.
- **Packed and ordered data.** We also highlight the challenges on decoding types, as the compiler packs as much data as possible into contiguous memory. Therefore, we need to know the specific types in the correct order to acquire the contract data, and that is not an easy task when we have an incomplete specification.

We acknowledge as a problem the challenges and difficulties of analyzing our own objects which are deployed in the Ethereum blockchain database. To solve this problem we propose an inspector that allows the user to perceive a clean representation of the object he/she is dealing with.

#### IV. SMARTINSPECT: CONTRACT INSPECTOR

SmartInspect is a local pluggable mirror-based reflection system for remotely deployed objects on a reflection-less system (the Ethereum platform). The goal of SmartInspect is to allow the inspection of known contracts based on its source code focusing on the debugging properties of interactivens and distribution [12]. This reflective approach allows a user to see the contents of any contract instance of the given source code, without needing to redeploy, nor develop any ad-hoc code.

SmartInspect is implemented in Pharo and it is publicly available as part of the SmartShackle tool suite.<sup>6</sup>

##### A. The Basics

The general idea of the Smart Inspector is to decompile the storage layout encoded by the Ethereum API (Figure 1). The decompilation employs a local pluggable mirror-based reflection architecture for remotely deployed objects on an Ethereum network.

Figure 2 shows the process to inspect contract, the pluggable reflective architecture generates a mirror for a given contract's source code by using its AST (Abstract Syntax Tree). Then, we use this mirror to extract information from a remote contract instance deployed in the blockchain (which is encoded as a binary memory layout). The contract data we gathered is exposed in four different formats: (i) data proxy object (REST), (ii) Pharo widget user interface, (iii) JSON, and (iv) HTML.

This approach allows us to access remote structureless information in a structured way. Our solution meets most of the desirable properties that are important for remote debugging namely: interactivens and distribution [12].

<sup>6</sup><https://github.com/RMODINRIA-Blockchain/SmartShackle>, <http://smart-shackle.io/>, verified 2018-02-22.




Fig. 1. Static storage with the related code definition.




Fig. 2. SmartInspect building process.

##### B. Discovering the Memory Layout

First we needed to decompile the binary representation of a deployed contract (i.e., a contract ABI) to discover the information inside the contract instance. That would resolve the opaqueness problem we described earlier, since we would be able to understand contract attributes.

The Ethereum API provides only one way to access memory layout of a contract: `getStorageAt` calls. This call gives access to a tree where information is encoded into slots accessible through contiguous indexes, for statically allocated memory (static types), and accessible by Keccak hash for dynamically allocated memory (variable sized arrays and mappings) [10]. It was a big challenge to decompile the memory layout because the Solidity documentation is incomplete. We had to reverse engineer some of the encoding performed by the compiler by ourselves.

There are two key restrictions in memory access: types and order. Each memory slot stores up to 32 bytes. As general policy, the compiler tries to pack as much data as possible for basic types. For example, two booleans and one `int128` occupy 18 bytes from a slot, one byte for each boolean plus 16 bytes






Fig. 3. Memory Layout Representation: (a) Static, (b) Dynamic.

for the *int128*. If we add another type that can not fit into the current slot, then the compiler places it in a new memory slot.

In the case of structs, they always start in a new slot and their data may take as many slots as needed. Any data after the struct will be encoded to start in a new slot, no matter if it could fit into the remaining struct slot. We show some of these encoding details for static types in Figure 3a.<sup>7</sup> Then, transversally, all the static representation of the different variables is available contiguously in the first slots of the binary representation, from 0 to *N*.

Dynamically allocated data types (e.g., arrays, mappings) are encoded in other hashed addresses, as shown in Figure 3b.<sup>8</sup> Therefore, the decompilation has to dive into a sometimes contiguous, sometimes indexed slots, with arbitrary allocations of space that may depend on the size of the type or in the next and/or previous type, to be able to read the stored content.

Therefore, we were successful in addressing our concern to introspect a deployed contract data, as we resolved the contract opaqueness problem by decompiling the binary representation and decoding the memory layout.

### C. Building the Mirror

After we decoded the memory layout, we still needed to apply our solution to any contract for a general reusable

<sup>7</sup>The figure shows contract attributes definitions with the arrows pointing to the memory representation. In this particular case, we first defined two booleans and one *int128*, which occupy 18 bytes from the first memory slot. Next, we define another *int128*, one *uint* (256 bits), and an *int8* that will occupy the second, third, and fourth memory slots respectively. The struct encoding places it in the fifth and sixth slot. Finally, even though the last variable (*int8*) could fit in the struct slot, the encoding places it in the next available slot (the seventh).

<sup>8</sup>We decided to represent the slot as contiguous data to facilitate its representation. In this case, we first defined three static variables (two booleans and one *int128*) to show they are placed in the first available slots. Then, we defined two dynamic types: one variable size array, and a mapping (which is similar to a hash table). Unlike static types, they are not encoded into the first memory slots but placed elsewhere in the memory.

solution. We employ a mirror-based architecture [11] that mimics the structure of any contract for us to access the memory layout that we can decode. A mirror works like an independent meta-programming layer which splits the concern of reflection capabilities into a mirror object.

First, we require the contract source code as input to start building the mirror. Then, we parse the source to create an AST (Abstract Syntax Tree). By interpreting the AST, we are able to know every type declared in the contract in the correct order. As described earlier (Section IV-B), we need know the types in order to decoded the memory layout and access the contract data.

Aiming at a general solution, we model configurable mirror objects that allow us to interact with deployed contract instances of the same configuration (usually meaning the same contract deployed in the blockchain). Our approach builds a composite mirror object, called *ContractMirror*, whose each component knows how to decode, in order, the contract state (Figure 4, the left side of the diagram). For each variable or struct a corresponding elementary mirror is added to the composite.

At this point we have a mirror with the structural representation of the AST that knows, in order, each contract property with its related type. Now our approach builds a representation of the memory layout to access the stored data. By using static code analysis provided by the AST, we can find the exact place of storage of every contract property. Moreover, we map the contract properties to its corresponding memory slot. Therefore, the mirror uses this mapping to gather the contract data for inspection. Figure 4 (on the right part) shows the mirror's memory mapping as a class diagram.

### D. Inspection Example

Getting back at our problem example (Section III-B), where the person in charge (Bob) needed to verify the data in a deployed poll contract instance. Back then, Bob had only two possibilities to address an user's issue, and both were




Fig. 4. Contract Mirror UML Diagram

not feasible. Now, let's suppose that Bob just learned about SmartInspect, which gives him a third possibility, inspect the contract data using our tool. Since Bob is the person in charge he has access to both the contract source code and its deployed binary representation. Bob executes SmartInspect on its poll contract and he is presented with the data from the pollTable property (Figure 5). Finally, Bob can see that the user was not registered, and he can now easily fix the problem by executing the *addVoter* function on his client application. This simple example illustrates the importance of contract inspection.




Fig. 5. SmartInspect Pharo User Interface Screenshot

## V. PRELIMINARY EVALUATION

In this section, we present a preliminary evaluation of the SmartInspect approach. The goal of this evaluation was to investigate whether SmartInspect implements the necessary and

desirable features for an inspector. We used the following four characteristics used for remote debugging by Papoulias [12]: Interactiveness, Distribution, Security, and Instrumentation. We also analyzed other five characteristics that are important for a blockchain remote inspector: Privacy, Pluggability, Consistency, Reusability, and Unrestricted types. We detail the nine characteristics as follows:

- *Interactiveness*: the inspector shows the object's state in real time. A lack of interactiveness could be a problem in blockchain platforms because of the contract's state may change during inspection and the user would be presented with outdated information.
- *Distribution*: the inspector can be extended for other technologies. Ideally, a debugger or inspector should rely on a middleware that is extensible. For smart contracts, the inspector should be extensible over different smart contract languages and blockchain technologies.
- *Security*: since remote debugging access a target through a network, it is important to ensure security from both ends. On the target side, the inspector should not have unrestricted access to its device; this is already ensured by the blockchain platform.
- *Instrumentation*: the inspector can alter the semantics of a process to assist in debugging. Basically, this is the mechanism to halt the process and inspect it at that point (e.g., breakpoints and watchpoints). This characteristic is not possible in blockchain platforms, as we cannot modify the deployed contract code to halt a function in the middle of its execution in the blockchain.
- *Privacy*: inspection should not breach or compromise data privacy by exposing data to unauthorized users. When considering smart contracts, a lack of privacy is dangerous



as it could be exploited by malicious users to acquire illicit advantages and resources.

- *Pluggability*: the inspector can be used on existing objects without the need to re-instantiate the objects or the system. For contracts, this means we can inspect existing deployed contracts without any dependency on the contract side, or the need to redeploy the contract. An unpluggable approach has the disadvantage of requiring the redeployment of a contract, which has non-trivial transactional costs.
- *Consistency*: the representation used by the inspector must reveal the information in a consistent manner, i.e., the inspection must reflect the current state of the deployed contract.
- *Reusability*: the inspector can be reused for different contracts. Lack of reusability would require a developer to spend time redefining the inspection for each individual contract.
- *Unrestricted Types*: the inspection can handle all types of objects. In contrast, a type-restricted inspector supports only a subset of data types (e.g., primitive types, static types).

We analyzed SmartInspect according to the characteristics of inspection tools just presented. Even though we wanted to compare SmartInspect against related approaches, as far as we know, there are no other inspectors available for Solidity smart contract. Therefore, we compared our approach against two other practices to access contract data: Getter methods, and ad-hoc Decoder (Table I).

TABLE I  
RELATED TECHNIQUES COMPARISON

| Characteristic     | SmartInspect | Getter  | Decoder |
|--------------------|--------------|---------|---------|
| Interactiveness    | Yes          | Partial | Partial |
| Distribution       | Yes          | No      | No      |
| Security           | Yes          | Yes     | Yes     |
| Instrumentation    | No           | No      | No      |
| Privacy            | Yes          | No      | Yes     |
| Pluggability       | Yes          | No      | Yes     |
| Consistency        | Yes          | Yes     | Yes     |
| Reusability        | Yes          | No      | No      |
| Unrestricted Types | Yes          | No      | Yes     |

As we can see from Table I, SmartInspect’s only characteristic flaw is related to *instrumentation* for remote debugging. However, this is a limitation imposed by the Ethereum blockchain technology rather than a design flaw in our approach.

Getter methods are a simple solution, since they are cheap to implement and easy to test. The developer does not need to know the memory layout of a contract to create getter methods. However, if the developer forgets to make a getter for a given attribute, he/she will need to re-deploy the contract and, most often, lose the data from the previous instance. Solidity does not support the return of many complex types (e.g, structs, mappings) on its functions. Therefore, a developer might need

to adapt his/hers data or function to provide access to a complex type. Moreover, the easy access to the data may cause a loss of privacy, since getter methods are a public part of the contract binary encoding.

Another practice is the ad-hoc Decoder that uses the Ethereum API on the memory slots. This is a complex task since it demands a deep understanding of the memory layout of each contract a developer plans to inspect. It also requires a developer to know the type of each attribute and code the ad-hoc decoder accordingly. Its advantages are that it allows access to data without loss of privacy and without the need to redeploy the contract. In fact, SmartInspect uses this concept of decoding memory layouts as a part of its inspection process.

## VI. DISCUSSION

In this section, we discuss our evaluation (Section VI-A) and the possible benefits for inspecting smart contracts (Section VI-B).

### A. Evaluating the Inspector

In our preliminary evaluation, we compared SmartInspect against two practices that can be used to access contract data (Section V). We acknowledge that we need other inspectors for a better comparison, since practices are a less polished solution than a fully designed approach for inspection. However, as far as we know, there is no other inspector tool available for Ethereum blockchain.

In the future, we plan to improve the evaluation by comparing with other inspectors, and performing a user driven evaluation.

### B. Benefits from Inspecting a Contract

There are several benefits for inspecting a smart contract on a blockchain platform. We highlight the following:

- **Easier to Understand.** Blockchain provides a mechanism that can be used to build trust between entities without a middleman [2], [3]. In the blockchain environment, smart contracts became a popular way to transfer digital assets among such entities [8], [9]. From a beginners perspective, it is better to work with concrete examples to understand a new concept. Thus, an inspector provides a simple way to access the contract state, which facilitates its understanding.
- **Find Bugs.** Contract inspection can help developers to find bugs more easily (as we illustrated in Section III-B).
- **Transparency.** Supporting inspection of contracts can increase transparency and improve overall trust among entities dealing with blockchain. For instance, it is possible for two entities to show the current state of their contracts to each other promoting transparency in their interactions.
- **Encourage the adoption of Contracts.** By allowing contract inspection, we can promote trust and transparency on blockchain platforms to companies and institutions. This will encourage more people to adopt smart contracts for their business or academic activities.

## VII. RELATED WORK

We organized the related work into three groups: (i) inspecting and debugging, (ii) reverse engineering, and (iii) blockchains and smart contracts.

### A. Inspecting and Debugging

Chis et al. [13] performed an exploratory research to better understand what developers expect from object inspectors, and based on that feedback they propose a novel inspector model. The authors interviewed 16 developers for a qualitative study, and a quantitative study conducting an online survey where 62 people responded. Both studies were used to identify four requirements needed in an inspector. Then they propose the Moldable Inspector, which indicates a new model to address multiple types of inspection needs. We followed the lessons taken by the Moldable Inspector when creating SmartInspect. We deem noteworthy the multiple views aspect, as SmartInspect can present its inspected data in four different views (REST, Pharo U.I., JSON, and HTML).

Papoulias [12] gives a deep analysis on remote debugging. As discussed by the author, remote debugging is specially important for devices that cannot support local development tools. The author identifies four important characteristics for remote debugging: interactiveness, instrumentation, distribution, and security. Based on the identified properties, Papoulias proposed a remote debugging model, called Mercury. Mercury employs a mirror based approach and an adaptable middleware. We used Papoulias research as an inspiration to create SmartInspect, specially relying on mirror for the remote inspection.

Salvaneschi and Mezini [14] propose a methodology, called RP Debugging, to debug reactive programs more effectively. The authors discuss that reactive programming is more customizable and easier to understand than its alternative the observer design pattern. The authors also present the main problems and challenges to debug reactive programs, and the main design decisions when creating their methodology. Although our inspector is from a different application domain, the RP Debugging design served as inspiration to plan our own inspecting approach.

### B. Reverse Engineering

Srinivasan and Reps [15] developed a reverse engineering tool to recover class hierarchical information from a binary program file. Their tool also extracts composition relationships as well. They use dynamic analysis to obtain object traces and then they identify the inheritance and composition information among classes on those traces. The authors experiments show that their recovered information is accurate according to their metrics. The author's tool contrasts with SmartInspect as, we use static analysis and they use dynamic analysis.

Caballero et al. [16] propose an approach to reverse engineer protocols by using dynamic analysis on program binaries. As stated by the authors, this approach differs from others that extract protocol information purely from network traces. The authors argue the importance to extract the protocol information, specially when there is no access to its specification, for

network security applications. They used 11 programs that implemented five different protocols for their evaluation. The authors' technique also contrasts with SmartInspect since they use dynamic analysis.

Fisher et al. [17] propose a multi phase algorithm to process *ad-hoc data* without human interaction. The authors describe *ad-hoc data* as semistructured information that does not have tools easily available. Basically, their algorithm reverse engineer the *ad-hoc data* into a domain specific language, which is used to generate a set of tools such as parsers, printers, query engine, and others. The authors evaluate the performance and correctness of their approach by using different benchmarks.

Lim et al. [18] designed an analysis tool called File Format Extractor for x86 (FFE/x86). They extract information from an executable file to perform several process, including static analysis. Their evaluation consists in applying their tool in three systems: gzip, png2ico, and ping. The authors work is similar to SmartInspect in a way that both approaches use static analysis in one of its steps to enhance the reverse engineering.

### C. Blockchain and Smart Contracts

Dinh et al. [7] describe a benchmarking framework to analyze private blockchain platforms. The authors contrast the different among public blockchain platforms (e.g., Ethereum) and private ones. For instance, private blockchain show more focus towards secure authentication. Although, their framework was designed for private blockchains, they evaluate it using public ones as well. Their evaluation measured four aspects (throughput, latency, scalability, and fault tolerance) in two blockchains (Ethereum and Hyperledger) and the Parity technology. The benchmark framework provided an in dept analysis of blockchain platforms that we used when we designed SmartInspect.

Luu et al. [8] investigates the security problems of executing smart contracts on the Ethereum platform. They also propose solutions to make the contracts more secure. The authors presents several scenarios and the possible malicious exploits for those scenarios. Based on the presented vulnerabilities, they propose solutions to make contracts more secure. The authors also propose a tool, called Oyente, that flags potential security flaws when coding smart contracts. Similarly to SmartInspect, Oyente also uses the contract bytecode to makes its security recommendations. However, our tool uses the memory layout to access the data, and Oyente uses the bytecode for a symbolic execution and security analysis.

Bhargavan et al. [9] proposed a framework to convert contracts to F\* and then improve their security. F\* is a functional programming language proposed by the authors in their work. According to them, F\* was designed to better verify correctness of contracts. Their approach decompiles the contract bytecode into a special F\* code to verify low-level properties; similarly it also compiles a contract source into an F\* version to verify high-level properties. The authors did a preliminary evaluation where 46 contracts were translated to F\*. Their approach also employs decompilation of contract bytecode and parses source code, similar to our SmartInspect.

## VIII. CONCLUSION

In this paper, we present the specific problems of inspecting Solidity smart contracts. Smart contract opaqueness added to the problems of reverse engineering compiler encoding and packing on different entity types makes the inspecting of values in smart contracts almost impossible for regular developers.

Our approach implementation, SmartInspect, is a local pluggable mirror-based reflection system for remotely deployed objects on reflection-less systems (the Ethereum platform). SmartInspect allows the inspection of known contracts based on its source code focusing on the debugging properties of interactiveness and distribution. This reflective approach allows a user to see the contents of any contract instance of the given source code, without needing to redeploy, nor develop any ad-hoc code to decode the memory representation.

We planned the following ideas for future work: (i) extend SmartInspect to inspect other smart contract languages employed in the Ethereum platform besides Solidity (e.g., Serpent, Viper, LLL); (ii) extend SmartInspect to support other blockchain platforms (e.g., Hyperledger); (iii) improve the introspection capabilities to support full debugging on smart contracts; and (iv) improve the evaluation by using SmartInspect on a big contract database, defining metrics for evaluating contract inspectors, comparing it with other inspectors, and performing an extensive user evaluation.

## ACKNOWLEDGMENT

This work was supported by Ministry of Higher Education and Research, Nord-Pas de Calais Regional Council, CPER Nord-Pas de Calais/FEDER DATA Advanced data science and technologies 2015-2020. This research was also supported by UTOCAT.

## REFERENCES

- [1] L. Luu, J. Teutsch, R. Kulkarni, and P. Saxena, "Demystifying incentives in the consensus computer," in *22nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 706–719. [Online]. Available: <http://doi.acm.org/10.1145/2810103.2813659>
- [2] Ethereum Foundation, "Ethereum's white paper," 2014. [Online]. Available: [https://en.wikibooks.org/wiki/LaTeX/Bibliography\\_Management](https://en.wikibooks.org/wiki/LaTeX/Bibliography_Management)
- [3] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system." 2009. [Online]. Available: [bitcoin.org](http://bitcoin.org)
- [4] A. Hari and T. V. Lakshman, "The internet blockchain: A distributed, tamper-resistant transaction framework for the internet," in *15th ACM Workshop on Hot Topics in Networks*, ser. HotNets '16. New York, NY, USA: ACM, 2016, pp. 204–210. [Online]. Available: <http://doi.acm.org/10.1145/3005745.3005771>
- [5] B. Leiding, P. Memarmoshrefi, and D. Hogrefe, "Self-managed and blockchain-based vehicular ad-hoc networks," in *2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, ser. UbiComp '16. New York, NY, USA: ACM, 2016, pp. 137–140. [Online]. Available: <http://doi.acm.org/10.1145/2968219.2971409>
- [7] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "Blockbench: A framework for analyzing private blockchains," in *Proceedings of the 2017 ACM International Conference on Management of Data*, ser. SIGMOD '17. New York, NY, USA: ACM, 2017, pp. 1085–1100. [Online]. Available: <http://doi.acm.org/10.1145/3035918.3064033>
- [6] S. Dziembowski, "Introduction to cryptocurrencies," in *22nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 1700–1701. [Online]. Available: <http://doi.acm.org/10.1145/2810103.2812704>
- [8] L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor, "Making smart contracts smarter," in *2016 ACM SIGSAC Conference on Computer and Communications Security (CCS 16)*. New York, NY, USA: ACM, 2016, pp. 254–269. [Online]. Available: <http://doi.acm.org/10.1145/2976749.2978309>
- [9] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Gollamudi, G. Gonthier, N. Kobeissi, N. Kulatova, A. Rastogi, T. Sibut-Pinote, N. Swamy, and S. Zanella-Béguelin, "Formal verification of smart contracts: Short paper," in *2016 ACM Workshop on Programming Languages and Analysis for Security*, ser. PLAS '16. New York, NY, USA: ACM, 2016, pp. 91–96. [Online]. Available: <http://doi.acm.org/10.1145/2993600.2993611>
- [10] Ethereum Foundation, "Solidity documentation release 0.4.18," Tech. Rep., 2017. [Online]. Available: <https://media.readthedocs.org/pdf/solidity/develop/solidity.pdf>
- [11] G. Bracha and D. Ungar, "Mirrors: design principles for meta-level facilities of object-oriented programming languages," in *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'04), ACM SIGPLAN Notices*. New York, NY, USA: ACM Press, 2004, pp. 331–344. [Online]. Available: <http://bracha.org/mirrors.pdf>
- [12] N. Papoulias, "Remote Debugging and Reflection in Resource Constrained Devices," Theses, Université des Sciences et Technologie de Lille - Lille I, Dec. 2013. [Online]. Available: <https://tel.archives-ouvertes.fr/tel-00932796>
- [13] A. Chiş, O. Nierstrasz, A. Syrel, and T. Gîrba, "The moldable inspector," in *2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*, ser. Onward! 2015. New York, NY, USA: ACM, 2015, pp. 44–60. [Online]. Available: <http://doi.acm.org/10.1145/2814228.2814234>
- [14] G. Salvaneschi and M. Mezini, "Debugging reactive programming with reactive inspector," in *Proceedings of the 38th International Conference on Software Engineering Companion*, ser. ICSE '16. New York, NY, USA: ACM, 2016, pp. 728–730. [Online]. Available: <http://doi.acm.org/10.1145/2889160.2893174>
- [15] V. Srinivasan and T. Reps, "Recovery of class hierarchies and composition relationships from machine code," in *23rd International Conference Compiler Construction (ETAPS/CC 2014), Grenoble, France, April 5-13, 2014. Proceedings*, A. Cohen, Ed. Springer Berlin Heidelberg, 2014, pp. 61–84. [Online]. Available: [https://doi.org/10.1007/978-3-642-54807-9\\_4](https://doi.org/10.1007/978-3-642-54807-9_4)
- [16] J. Caballero, H. Yin, Z. Liang, and D. Song, "Polyglot: Automatic extraction of protocol message format using dynamic binary analysis," in *14th ACM Conference on Computer and Communications Security (CCS'07)*, 2007, pp. 317–329. [Online]. Available: <http://doi.acm.org/10.1145/1315245.1315286>
- [17] K. Fisher, D. Walker, K. Q. Zhu, and P. White, "From dirt to shovels: Fully automatic tool generation from ad hoc data," in *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL '08. New York, NY, USA: ACM, 2008, pp. 421–434. [Online]. Available: <http://doi.acm.org/10.1145/1328438.1328488>
- [18] J. Lim, T. W. Reps, and B. Liblit, "Extracting output formats from executables," in *13th Working Conference on Reverse Engineering (WCRE 2006), 23-27 October 2006, Benevento, Italy, 2006*, pp. 167–178. [Online]. Available: <https://doi.org/10.1109/WCRE.2006.29>