THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

SIF: A Framework for Solidity Contract Instrumentation and Analysis

Citation for published version:

Peng, C, Akca, S & Rajan, A 2020, SIF: A Framework for Solidity Contract Instrumentation and Analysis. in *2019 26th Asia-Pacific Software Engineering Conference (APSEC)*. Institute of Electrical and Electronics Engineers (IEEE), pp. 466-473, The 26th Asia-Pacific Software Engineering Conference, Putrajaya, Malaysia, 2/12/19. <https://doi.org/10.1109/APSEC48747.2019.00069>

Digital Object Identifier (DOI):

[10.1109/APSEC48747.2019.00069](https://doi.org/10.1109/APSEC48747.2019.00069)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2019 26th Asia-Pacific Software Engineering Conference (APSEC)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

SIF: A Framework for Solidity Contract Instrumentation and Analysis

Chao Peng
University of Edinburgh
Edinburgh, United Kingdom
chao.peng@ed.ac.uk

Sefa Akca
University of Edinburgh
Edinburgh, United Kingdom
s.akca@sms.ed.ac.uk

Ajitha Rajan
University of Edinburgh
Edinburgh, United Kingdom
arajan@ed.ac.uk

Abstract—Solidity is an object-oriented and high-level language for writing smart contracts that are used to execute, verify and enforce credible transactions on permissionless blockchains. In the last few years, analysis of smart contracts has raised considerable interest and numerous techniques have been proposed to check the presence of vulnerabilities in them. Current techniques lack traceability in source code and have widely differing work flows. There is no single unifying framework for analysis, instrumentation, optimisation and code generation of Solidity contracts.

In this paper, we present SIF, a comprehensive framework for Solidity contract analysis, query, instrumentation, and code generation. SIF provides support for Solidity contract developers and testers to build source level techniques for analysis, understanding, diagnostics, optimisations and code generation. We show feasibility and applicability of the framework by building practical tools on top of it and running them on 1838 real smart contracts deployed on the Ethereum network.

Index Terms—high level languages, software testing, code instrumentation, program analysis

I. INTRODUCTION

Blockchains are the underlying technology for making on-line secure transactions using cryptocurrencies such as Bitcoins and Ethers. Executing, verifying and enforcing credible transactions on blockchains is done using smart contracts, which is code written by the buyer and seller using Turing-complete languages [1]. Solidity is a popular object-oriented and high-level language for writing smart contracts [2], [3] and can be compiled to bytecode for execution on the blockchain network. With the increased use of smart contracts across application domains, there is a crucial need for a unified framework that supports and facilitates Solidity code analysis, understanding, transformation, and development of tools for verification and testing that provide strong security guarantees.

A. Problem

Traditional programming languages are generally supported by a comprehensive framework for code instrumentation, monitoring, optimisation and code generation such as LLVM/Clang [4] for C/C++. This framework support is lacking for smart contract programming languages like Solidity. Existing work and tools for Solidity contracts is primarily based on analysing the bytecode of smart contracts for security vulnerabilities, or translating Solidity code to other languages or intermediate forms over which security analysis is performed [3], [5]–[10]. For instance, Oyente [5] analyses bytecode from

the Solidity compiler and successfully reports the presence of vulnerabilities. However, it lacks the ability to trace and localise bugs in Solidity code. Zeus [10] translates smart contracts to LLVM bitcode but does not support complete Solidity syntax, including `throw` statements, `self-destructs`, virtual functions and assembly code blocks. These techniques, although effective in detecting security vulnerabilities, are not general purpose code instrumentation and analysis tools. As a result, building program analysis tools with a purpose different from detecting security vulnerabilities is difficult using existing techniques. Traceability back to Solidity source code is also a common issue in these techniques. We propose a general purpose instrumentation and analysis framework that addresses this need for Solidity contracts.

B. Goal

We seek to provide Solidity developers and testers with a framework, similar to the widely-used C/C++ instrumentation technique Clang LibTooling [11], to easily and effectively understand, manipulate and analyse Solidity code. To achieve this goal, we design and implement *Solidity Instrumentation Framework (SIF)* with the following capabilities:

- 1) Provide an interface for users to query the Abstract Syntax Tree (AST) of Solidity code.
- 2) Support code instrumentation or transformation using pre-defined helper functions that can modify the AST.
- 3) Support Solidity code generation from AST.

Tools. We demonstrate the general purpose nature of the SIF framework in supporting code instrumentation and analysis by using it to implement 7 different tools or utilities for Solidity contracts. Four of these tools query the Solidity abstract syntax tree (AST) to provide information on the contract code and three others instrument and modify the code using the AST. A brief description of the tools is provided below,

Function Listing lists all the function definitions with function names, return lists, parameter lists and which contract they belong to. Useful in summarising and reviewing Solidity contracts.

Call Graph Generator produces function call graphs depicting calling relationships between functions in a Solidity contract. Nodes represent functions and edges represent the calls relation. Useful in understanding and reviewing smart contracts.

Control Flow Graph Generator produces a graphical representation of control flow within a Solidity contract. Control flow graphs are very useful in static analysis and program optimisations.

AST Diff is a syntactic diff tool for comparing Solifity contracts at the AST level. It ignores comments, white spaces and blank lines.

SIF Rename allows the user to easily rename existing identifiers. All definitions and references to the specified identifier will be changed to the new name provided.

Fault Seeder allows developer to artificially seed bugs in smart contracts. The tool currently supports 3 common bug/vulnerability types. The seeded bugs can be used to assess effectiveness of testing or verification tools in uncovering vulnerabilities.

Assertion Analyser examines all AST nodes of the smart contract and determines if the node is susceptible to vulnerabilities, such as division by zero, overflow and underflow. Assertions will be inserted by the tool at contract locations susceptible to these vulnerabilities. Assertions will be checked during execution of the modified contract.

We evaluate SIF and the 7 tools built on top of it on a collection of 1838 real smart contracts that are running on the Ethereum network. Our results show that SIF and the utility tools are easy to use, highly automated and new tools can be easily implemented with the helper functions.

II. RELATED WORK

Several static analysis and vulnerability detection tools for smart contracts have emerged in the last few years. We present and discuss existing work over Solidity contracts for vulnerability detection, code generation, query and instrumentation.

a) Vulnerability detection: Many of the existing techniques rely on bytecode analysis to check for potential vulnerabilities [3], [5]–[7]. Bytecode is the compiled hexadecimal format of smart contracts. Solidity contract verification at the bytecode level is sufficient to detect vulnerabilities. However, it is challenging to trace the problem back to Solidity source code, that the developer can then fix. Some techniques translate Solidity code to the F* languages or represent code in an intermediate form, such as LLVM IR (Intermediate Representation) or XML [8]–[10]. Vulnerability analysis is then done over the intermediate representation. The techniques do not support translation of the intermediate form back to Solidity, which results in loss of traceability in source code.

b) Solidity code generation from AST/Intermediate form: Currently, translation from intermediate form or AST back to Solidity has limited support. We are only aware of one tool, Soltar [12], that supports translation of Solidity AST back to Solidity code. However, Soltar does not provide capabilities supported by SIF to query and modify the AST. Additionally, Soltar is not maintained and does not support Solidity versions 0.4.3 onwards. SIF, presented in Section III, is able to handle the latest Solidity version 0.5.3, and older versions.

An existing tool, Zeus [10], supports translation of Solidity code to LLVM bitcode but does not have the capability

to generate Solidity back from bitcode. Additionally, complete Solidity syntax is not supported (`throw` statements, `self-destructs`, virtual functions and assembly code blocks) and there is no provision for code instrumentation. Porosity [13] and an Ethereum smart contract decompiler, JEB [14], translate bytecode back to Solidity code. However, support for Solidity code query, modification and instrumentation is not available in these tools.

c) Query and instrumentation of Solidity code: There exists no general purpose framework or tool that provides capabilities for users to query, modify and instrument Solidity code. These capabilities are crucial for program understanding, analysis, optimisation and verification techniques [15]–[18]. The SIF framework is novel in providing these capabilities and in providing interfaces that allow users to easily define the kind of query, modification and instrumentation needed. SIF's capabilities have allowed us to query and instrument Solidity code to produce call graphs, control flow graphs, function listing, Fault Seeder, Assertion Analyser along with scope for building many custom utilities and diagnostics. We discuss SIF's design and the tools built on it in the next Section.

III. SIF OVERVIEW

In this section, we present our generic framework, SIF, implemented in C++. SIF enables Solidity code query and instrumentation and works at the AST level. The framework also provides the capability for generating Solidity code back from the AST. SIF supports the entire Solidity syntax, up to version 0.5.3. The workflow of SIF is illustrated in Figure 1. SIF starts from the AST of Solidity code, produced by the Solidity compiler. It then accepts user instructions on queries and/or modifications needed. The framework then gathers the desired query information or performs modifications to the AST and finally generates Solidity code from the AST.

In the rest of this Section, we describe the design SIF, how it can be used for query and instrumentation along with instructions needed from the user, a concrete illustration of the query and instrumentation capabilities using seven tools that we built on top of SIF. The framework with its source code and user guide is available at <https://github.com/chao-peng/SIF>. To allow users to try the framework without having to download and build the source code, we have provided an online version of SIF at <https://wandbox.org/permlink/PnaL6bO9zipKRuKu>.

A. SIF Design

Operations of SIF are divided into 3 phases. Phase 1 focuses on representing AST nodes as C++ classes with methods to retrieve and modify information of the node. Phase 2 interacts with user defined query and/or instrumentation functions, and traverses the AST to perform the desired operation. Phase 3 generates Solidity code from the AST. We discuss each of these phases in more detail in the rest of this Section. We use an example `struct` definition named `Request` shown in Listing 1, containing a data element and a method, to illustrate the different phases.

Listing 1: Struct example

Fig. 1: SIF work flow

Fig. 2: An AST example highlighting a struct definition

```

1 struct Request {
2   bytes data;
3   function (bytes memory) external callback;
4 }

```

Phase 1: AST Representation. The Solidity compiler generates the AST from Solidity code in two formats: plain text and JSON (JavaScript Object Notation, a structured data format). In an attempt to make our tool for code instrumentation generic and easy to use, we use C++ classes as our intermediate AST representation. Given the Solidity AST, SIF first traverses the AST and for each node, it instantiates a class of that node type with all the associated information. Figure 2 shows a struct definition appearing in an example AST. The classes contain information about the AST node in their data fields, and provide methods to query and modify the data. Listing 2 shows the C++ class representation of a struct definition node in the AST. The C++ class provides methods for getting name of the struct, setting it to a different name, querying number of fields in the struct, getting a particular field, adding, removing and updating a field. These methods facilitate query and modification of the AST node.

Listing 2: Representing structs using a C++ class

```

1 class StructDefinitionNode : public ASTNode {
2 public:
3   StructDefinitionNode() : ASTNode(); // Constructor
4   string source_code(); // produce source code
5   string get_name(); // get the name
6   void set_name(new_name); // update the name
7   int num_field(); // get the number of fields
8   void add_field(new_field); // add a new field
9   void remove_field(index); // remove a field
10  void update_field(index, _new); // update a field
11  ASTNode get_field(index); // get a field by
12 private:
13  string name; // name of the struct
14  vector<ASTNode> fields; // list of struct fields
15 };

```

Phase 2: Query and Instrumentation. SIF interacts with users for AST query and instrumentation through a function *visit* that

it declares. SIF traverses the AST starting from the root node in a depth-first fashion using the *visit* function. The user can implement queries and modifications, if any, that are needed within the *visit* function. For example, the user can change the name of the example struct *Request* from Listing 1 to *DirectRequest* by implementing it in the *visit* function, as shown in Listing 3. Each time SIF visits an AST node, the *visit* function is called to process operations defined by the user. In this example, the *visit* function first determines whether the current node is a struct definition and then checks whether the struct name matches *Request*. If there is a match, the struct name is changed to *DirectRequest*. The methods, *get_name()* and *set_name()*, used in this implementation are helper methods provided by SIF to facilitate writing of AST queries and instrumentations.

Listing 3: Change the name of a struct in SIF

```

1 void visit(ASTNode* node) {
2   if (node->get_node_type() == StructDefinition) {
3     StructDefinition* sd = (StructDefinition*) node;
4     if (sd.get_name() == "Request")
5       sd.set_name("DirectRequest");
6   }
7 }

```

Phase 3: Solidity code generation. For each type of AST node classes, SIF defines code templates that allow corresponding Solidity code to be generated. Listing 4 illustrates the code template for the *StructDefinitionNode*. In the listing below, the template generates Solidity code for a struct definition by first printing the struct keyword, followed by the name of the struct, a left brace indicating the start of struct element definitions. The template then iterates through the list of struct element definitions. Once an element definition is visited, the source code of that definition is appended to the struct source code. Once all the element definitions are completed, a right closing brace is added, indicating the end of the struct definition. If the name of the struct were changed by the user, as shown in Listing 3, the new struct name will appear in the generated Solidity code.

Listing 4: Source code template of structs

```

1 string StructDefinitionNode::source_code() {
2   string source = "struct_" + name + "\n";
3   Iterate the list of element definitions:
4     source += element.source_code() + "\n";
5   return source + ";\n";
6 }

```

Fig. 3: Using SIF to build tools

B. Using SIF

SIF provides an interface in the form of predefined functions - *before*, *visit* and *after*, to perform user-defined AST queries and instrumentations. We have presented in Section III-A how the *visit* function can be used for this purpose. The *before* function is called by SIF in advance of traversing the AST, and can be used for data initialisation. The *after* function is called when visiting the AST is completed and helps summarise information gathered from the AST. In the following Section, we demonstrate how these three functions can be used in custom utility tools built with SIF.

C. Tools using SIF

To illustrate the framework's extent of support for building custom queries and instrumentations, we implemented seven utility tools for Solidity contracts, as shown in Figure 3. The seven tools were inspired by tools built over other mature code instrumentation frameworks, such as Clang LibTooling for C/C++. We discuss and present each of these tools in the rest of this Section.

a) *Function Listing*: Given a Solidity file with multiple contracts and function definitions, this tool outputs a list of functions that appears in the file with function names, parameter lists, return lists and the contracts in which they are defined. It is worth noting that in Solidity, functions are considered members of contracts and appear as children of contract definition nodes in the AST representation. To implement the Function Listing tool, we use the predefined *visit* function within SIF. We check whether the visited node is a contract definition node. For contract definition nodes, we check all the children nodes to see if they are of type function definition. If function definition nodes exist, we record information on its name, parameters, and return values using SIF's helper methods for this node type. Once

SIF finishes traversing the AST, the recorded information on functions definitions is printed. Listing 5 shows the list of function definitions printed by this tool for the AztraToken contract in Etherscan¹.

Listing 5: Function definitions summarised by Function Listing

```

1 [In AztraToken] AztraToken() returns ()
2 [In AztraToken] _transfer(address _from, address _to
  , uint _value) returns ()
3 [In AztraToken] transfer(address _to, uint256 _value
  ) returns ()
4 [In AztraToken] transferFrom(address _from, address
  _to, uint256 _value) returns (bool success)
5 [In AztraToken] approve(address _spender, uint256
  _value) returns (bool success)
6 [In AztraToken] burn(uint256 _value) returns (bool
  success)
7 [In AztraToken] burnFrom(address _from, uint256
  _value) returns (bool success)
8 [In AztraToken] mintToken(address target, uint256
  mintedAmount) returns ()
9 [In AztraToken] freezeAccount(address target, bool
  freeze) returns ()
10 [In AztraToken] transferOwnership(address newOwner)
   returns ()

```

b) *CG Generator*: CG Generator illustrates the calling relationships between functions as a call graph (CG). The tool is implemented using the *visit* function first to process two types of AST nodes: function definition and function call. The tool maintains a map containing the callee function with the caller functions it is associated with. After traversing the full AST, the graph is drawn by the *after* function using a graph drawing tool, Graphviz [19], with function names as nodes and edges based on the relations recorded in the map. Figure 4 presents the call graph generated by our tool for the AztraToken contract.

c) *CFG Generator*: Control flow graphs (CFG) are used to illustrate control flow within a program [16] and is fundamental to many static analysis and compiler optimisation techniques. The tool CFG Generator produces control flow graphs for functions or subroutines within Solidity contracts. Figure 5 shows the CFG generated by our tool for the *uint2str* function, shown in Listing 6, that converts an unsigned integer variable to a string. The *uint2str* function is found in contract *Item* from Etherscan².

Listing 6: Function to generate a control flow graph using CFG Generator

```

1 function uint2str(uint i) internal pure returns (
  string) {
2   if (i == 0) return "0";
3   uint j = i;
4   uint len;
5   while (j != 0) {
6     len++;
7     j /= 10;
8   }
9   bytes memory bstr = new bytes(len);
10  uint k = len - 1;
11  while (i != 0) {
12    bstr[k--] = byte(48 + i % 10);

```

¹The contract AztraToken is available at <https://etherscan.io/address/0x6962E259a8f9633C4494764628A7984cCEd58e10>

²The contract Item is available at <https://etherscan.io/address/0x5f896c654a08323dbe16aded331c461ccaeb370>

Fig. 4: Call Graph of the Smart Contract AztraToken

Fig. 5: Control Flow Graph of the Function uint2str from the Smart Contract Item

```

13     i /= 10;
14 }
15 return string(bstr);
16 }

```

CFG Generator traverses the AST and clusters nodes in a straight line code sequence into basic blocks. When a node introducing control flow is encountered, the most recent basic block ends and is linked to the condition node with the control flow. The true branch of this control flow is linked to the first basic block of the then branch for an if statement or the loop body if it is a loop. The false branch is linked to the first basic block of the else branch for an if statement or the first basic block immediately after the loop body if it is a loop. Basic blocks and links between them are maintained and recorded in a data structure as the AST is traversed. Once the AST traversal is completed, the *visit* function passes the data structure to Graphviz [19] to generate a graph depicting basic blocks and control flow. The data structure and resulting

CFG is shown in the Listing 7 and Figure 5, respectively.

Listing 7: Data structure produced by CFG Generator

```

1 Node [0] -> Node [1];
2 Node [1] -> Node [2] label=True Branch;
3 Node [1] -> Node [3] label=False Branch;
4 Node [2] -> Node [3];
5 Node [3] -> Node [4];
6 Node [4] -> Node [5] label=True Branch;
7 Node [4] -> Node [6] label=False Branch;
8 Node [5] -> Node [4];
9 Node [6] -> Node [7];
10 Node [7] -> Node [8] label=True Branch;
11 Node [7] -> Node [9] label=False Branch;
12 Node [8] -> Node [7];

```

d) *AST Diff*: AST Diff takes two Solidity contracts and compares them at the AST level to report syntactic differences. The tool ignores differences in comments, white spaces, and empty lines. Implementation using the *visit* function starts from the root node of one smart contract. Every node visited within this smart contract is compared with nodes in the other contract. Differences in data fields within the node are reported. Extra or missing nodes are also shown.

e) *SIF Rename*: This tool allows users to change the names of functions, variables and contracts by locating and replacing all occurrences. The tool accepts three additional inputs: identifier type, old name and new name. The *before* function first parses the command line options for these additional inputs. As SIF traverses the AST, the *visit* function checks if the currently visited node matches the type and name of the identifier provided by the user. If a match is found, the old name gets replaced with the new one. The replacement is also performed within other nodes that reference the provided identifier. For instance, if a user wants to rename a particular contract, SIF Rename will rename the contract within the contract definition AST node. It will also rename references to the contract within other nodes such as `member access` and `using for directive` that may reference the given contract.

Fault Seeder and Assertion Analyser These tools focus on 3 types of security vulnerabilities commonly found in smart contracts, presented in Table I. Fault seeder instruments the code by injecting vulnerabilities. Assertion Analyser uses AST query and instrumentation capabilities to insert assertions in the contract.

f) *Fault Seeder*: For each type of vulnerability illustrated in Table I, the Fault Seeder tool creates a new code block containing that vulnerability and injects it into the Solidity

Type of Vulnerability	Assertions Added by Assertion Analyser	Fault seeded by Fault Seeder
Division by Zero	<code>require(c != 0); a = b / c;</code>	Statement with division by zero is inserted.
Unsigned Overflow	<code>a = b + c; assert(a >= b && a >= c);</code> <code>a = b * c;</code> <code>(b != 0 && c != 0)? assert(a >= b && a >= c): assert(a == 0);</code>	Arithmetic operation resulting in overflow is inserted.
Unsigned Underflow	<code>a = b - c; assert(b >= a && b >= c);</code>	Arithmetic operation resulting in underflow is inserted.
Signed Overflow / Underflow	<code>a = b + c;</code> <code>assert((c >= 0 && a >= b) (c < 0 && a < b));</code> <code>a = b - c;</code> <code>assert((c >= 0 && a <= b) (c < 0 && a > b));</code> <code>a = b * c;</code> <code>(b != 0 && c != 0)? assert((a / b == c) && (a / c == b)): assert(a == 0);</code>	Arithmetic operation resulting in overflow / underflow is inserted.

TABLE I: Vulnerability types

contract. SIF traverses the contract AST and inserts an extra AST node containing a single vulnerability into the original AST. Solidity code is then generated from the modified AST, and is referred to as a mutated contract. For example, to introduce integer underflow, *Fault Seeder* inserts the code snippet shown in Listing 8 into the original smart contract.

Listing 8: Code snippet to introduce underflow by Fault Seeder

```
1 uint256 minuend = 20;
2 uint256 subtrahend = 250;
3 uint256 result = minuend - subtrahend;
```

g) *Assertion Analyser*: Inserting property checks using assert statements is helpful in verifying program correctness. The *Assertion Analyser* tool inserts pre- and post-conditions for arithmetic operations to help detect overflow, underflow and division by zero vulnerabilities. Within the *visit* function, the *Assertion Analyser* tool first gathers information on the node, and operators and operands if any within. For example, an expression node with arithmetic operations may be prone to an overflow/underflow error. The tool then inserts relevant assert statements, shown in Table I, as extra nodes following the node under analysis. For division by zero vulnerability, SIF inserts a pre condition, before the node under analysis with a division operator, that asserts the divisor expression is greater than zero. Consider the function *uint2str*, shown in Listing 6, with a subtraction operation on line 10. The *Assertion Analyser* tool will insert an assert statement, shown in Listing 9, to check for unsigned integer underflow.

Listing 9: Inserted guard for a subtraction operation by Assertion Analyser

```
1 uint k = len - 1;
2 assert(len >= k && len >= 1);
```

IV. EVALUATION

We evaluate feasibility, ease of use and extent of automation in using SIF and the seven tools built on it over 1838 unique Solidity contracts from the Ethereum network ³.

Additionally, after SIF was released on Github in May 2019, we received various enquiries about using the framework from different parts of the globe. The interest expressed in the framework was encouraging. The feedback helped us fix bugs and improve the usability of the framework and is discussed in Section IV-C.

³Our dataset is available at <https://github.com/chao-peng/SIF>

A. Research Questions

We investigate the following questions in our experiment:

Q1. Solidity Language Support: *Does SIF support all constructs in the Solidity syntax and generate code accordingly?* To answer this question, we first inspect contract code in our dataset and the official Solidity documentation to check the range of syntactic structures present in the dataset. We then use SIF, without any instrumentations, to generate Solidity contracts from ASTs and we compare if the generated code is the same as the original.

Q2. Correctness of AST Query and Instrumentation: *Are the tools built on top of SIF able to query and instrument the AST and produce correct output?* We use a small user group to help answer this question. The user group comprised of ourselves (User-0), 2 PhD students familiar with Solidity language (User-1 and User-2), and a second year undergraduate student (User-3) with basic knowledge of Solidity. We manually check the correctness of the query and instrumentation information over the different contracts.

Q3. Feasibility, Extent of Automation and Ease of Use: *How feasible is the framework for users to write their own tools? Are the tools fully automated and easy to use?* We asked User-1, User-2 and User-3 to evaluate the framework and the 7 tools on the following aspects:

- 1) Ease of implementing new tools** On a scale of 1 (very hard) to 5 (very easy), rate how easy it is to implement new tools using SIF. We asked the users to write a new tool, *Loop Count*, to report the number of loops in a Solidity contract, that utilises the AST query capability of the framework. We also asked the users to create a tool, *Make Signed*, that changes all unsigned integer (uint) types to signed integer (int), utilising both query and instrumentation capabilities.
- 2) Extent of automation of the 7 tools** On a scale of 1 (Completely manual) to 5 (fully automated), rate the extent to which the seven tools can be run automatically.
- 3) Ease of use of the 7 tools** On the scale of 1 (Very difficult to use) to 5 (Very easy to use), rate the ease with which the 7 tools can be used.

B. Experimental Result

We discuss the experimental results in the context of the experiment questions presented earlier in this Section.

1) *Solidity language support*: The 1838 Solidity contracts in our dataset contain a wide variety of Solidity constructs. We manually checked against the Solidity documentation as well as the source code of the official Solidity compiler, and confirmed that our dataset covers all syntactic elements in Solidity. Our technique was able to analyse the AST and generate source code for all 1838 smart contracts from their ASTs automatically. The generated contracts were also compared by the 3 users to the original contracts, and no differences were reported. As mentioned in Section II, the F* tool [8] does not support loop structures in Solidity. A significant fraction of contracts in our dataset contain loops and SIF is able to fully support loops. Self-destructs, throw statements and inline-assembly blocks are not supported by the Zues verification tool [10]. Our framework fully supports the use of these constructs in Solidity contracts.

2) *Correctness of AST query and instrumentation*: The seven tools discussed in Section III-C utilise SIF capabilities for AST query and instrumentation. We evaluate correctness of these capabilities in the context of the seven tools.

To validate output of tools `Function Listing` and `CG Generator`, we split the dataset into two halves. With contracts in the first half of the dataset, two users ran the tools to collect function lists and call graphs. The other two users validated the output produced by the tools by inspecting the original contract. We, then, swapped the roles of the users for contracts in the second half. This was done to reduce biases, if any, and the monotony in checking the contracts. The users did not find any anomalies in the tools over the given dataset.

Control flow graphs produced by `CFG Generator` were manually inspected by User-0 to check its correctness. `CFG generator` correctly generated CFGs for all contracts in the dataset. To evaluate `AST Diff`, we split the dataset into two halves again. User-0 and User-3 modified contracts in the first half by renaming identifiers, mutating operators and deleting code blocks for the first half. User-1 and User-2 ran the `AST Diff` on each set of original and modified contract (from first half), to check if all the modifications were reported correctly by the tool. The user roles were swapped for contracts in the second half of the dataset. `AST Diff` was found to work correctly by all users over all the contracts.

For the instrumentation capability, we used `SIF Rename` to locate and change names of identifiers within contracts in the dataset. We again, split the user responsibilities so that two users selected and changed identifier names and two users checked the changes were reflected in the generated Solidity code. Different identifier types were selected. The users found `SIF Rename` worked correctly for the given dataset and requested changes.

Each of the four users ran `Fault Seeder`, selecting different vulnerability types to be seeded in each Solidity contract, and inspected the mutated contracts to check if the vulnerabilities were seeded as per choice. The mutated contracts were also run using the Solidity compiler, `solc`,

to check for syntactic correctness. No issues were reported. The `Assertion Analyser` tool was run on all contracts, and the users individually inspected if the assertions were inserted correctly for the arithmetic operations. The instrumented contracts with assertions were checked through the `solc` compiler for syntactic correctness; no problems were found. The users were able to confirm the assertions were inserted by the tool correctly around arithmetic operations in all contracts.

User #	Existing 7 Tools		Using SIF	
	Extent of Automation	Ease of Use	Ease of Implementing New Tools Loop Count	Make Signed
1	4	4	4	4
2	5	4	5	4
3	4	5	4	3
Avg.	4.3	4.3	4.3	3.7

TABLE II: User experience of using SIF and tools

3) *Extent of automation and ease of use*: Table II reports how User-1, User-2, and User-3 rate the existing 7 tools for ease of use and extent of automation. The table also shows how the users rate feasibility and ease with which new tools, `Loop Count` and `Make Signed`, can be implemented using SIF. As the user interface for the existing 7 tools is similar, we report the average ratings across the 7 tools for ease of use and extent of automation for each user. User-1, User-2, User-3 rate the extent of automation for the 7 tools as 4, 5 and 4, respectively, giving a high average user experience of 4.3 (mostly fully automated). This implies that the users agree that the tools have a high level of automation with little user intervention. The users gave a score less than 5 as the users had to use the Solidity compiler to first produce the AST and then call the tool. We will look to automate this step in the future so the users do not have to explicitly call the Solidity compiler.

For ease of use of the 7 tools, all three users felt the interface for running the tools was easy to understand and use (ratings of 4, 4 and 5).

For writing new tools with SIF, all three users felt it was easy to implement the `Loop Count` tool (average implementation ease being 4.3) owing to implementation similarity to the existing `Function Listing` tool (replace functions by loops). The users found implementing the `Make Signed` tool was slightly more challenging, as it is different from implementations of the existing 7 tools. Implementing the `Make Signed` tool requires searching through variable declarations, function parameter lists to locate all unsigned integer qualifiers and replacing them with the signed qualifier. User-3 found implementation of the `visit` function was only slightly easy (rating of 3), while User-1 and User-2 found it easy (rating of 4). User-3 is not proficient in C++, and found the use of pointers in the `visit` function interface tricky. To address this issue, we have added comments in the interface showing how to use and cast to AST node pointers. Average ease of implementation across the users was 3.7 for the `Make Signed` tool.

Finally, we report the time taken in using the framework to generate Solidity code from unmodified ASTs. For contracts

with less than 1000 lines of code (90.0% of the dataset), the framework finished code generation in 4 seconds. One of the bigger contracts, *Expiry*⁴ with 6183 lines of code, takes 85.6 seconds for code generation from AST. To understand why code generation for this contract took significantly longer, we profiled the framework phases and found that 40.3 seconds was spent on parsing the JSON format AST. Our framework uses the text format AST as the main reference and only queries the JSON format AST when it encounters keywords missing in the text AST. However, each such query starts from the root of the JSON AST for each newly visited node, resulting in significant overhead. Other contracts with more than 5K lines of code finished code generation in less than 15 seconds as they did not have to refer to the JSON AST very frequently. In our future work, we will optimise and accelerate queries involving JSON data structures.

C. Feedback from the Global Community

In the 2 months since SIF was released on GitHub, we noted interest in its use and received queries from users around the world. In this Section, we present feedback and suggestions from global users.

a) *Bug report related to uninitialised variables:* This was reported by a user from Ohio State University. We found the bug was caused by un-initialised smart contract variable definitions. The framework looks for the missing initialisation statement in such variable definitions. We have now fixed this bug.

b) *Enquiry on gathering information from the AST:* 2 users from China did not know how to maintain and summarise information gathered from different AST nodes. We provided them with guidance and examples of using *before* and *after* functions to address this issue. We also updated our documentation to include this.

c) *Request for more Command Line options:* A user from Hong Kong University of Science and Technology made this request. After each SIF run, smart contract source code is generated by default. However, code generation is not always desired, especially for users who only want to query information from the AST. We have now provided this option to avoid code generation. We plan to add more command line options to SIF for additional user control.

d) *Request for detailed documentation:* A user from ShanghaiTech University requested more detailed documentation on each type of AST node and the associated fields and methods. We have provided this documentation for all types of AST nodes, along with examples using them.

V. CONCLUSION

SIF provides the capability to query, analyse and instrument the AST of Solidity contracts and generate Solidity code back. The framework uses a C++ intermediate representation, and provides helper methods to gather information on different AST nodes and manipulate them, as needed. The framework is generic and eases the implementation of custom query and

instrumentation capabilities. Users interact with the framework using a simple interface that is easy to understand and use. We built 7 tools for different types of Solidity code query and instrumentation to evaluate SIF's versatility and asked 3 users to rate the usability and automation of the tools along with ease of implementing new tools. We used 1838 unique contracts in our evaluation.

We found SIF was able to run all 1838 contracts, and the outputs of all 7 tools were confirmed to be correct by all the users. Additionally, the users found the 7 tools were easy to use and nearly fully automatic. The users also implemented 2 new tools using SIF and found the framework was well suited for implementing new tools, owing to the helper functions and simple interface. SIF's generality and ease of use makes it a promising framework for Solidity developers and testers to quickly and easily build, maintain new tools or leverage existing capabilities. Since the release of the framework on GitHub, we have received interest in its use for Solidity code analysis and instrumentation from several international users. We envision SIF becoming a large tooling infrastructure like LLVM/Clang, supporting diagnostics, easy development and maintenance of Solidity tools.

REFERENCES

- [1] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," *white paper*, 2014.
- [2] C. Dannen, *Introducing Ethereum and Solidity*. Springer, 2017.
- [3] P. Tsankov, A. Dan, D. Drachler-Cohen, A. Gervais, F. Buenzli, and M. Vechev, "Securify: Practical security analysis of smart contracts," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pp. 67–82, ACM, 2018.
- [4] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis and transformation," (San Jose, CA, USA), pp. 75–88, Mar 2004.
- [5] L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor, "Making smart contracts smarter," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 254–269, ACM, 2016.
- [6] MythX, "Mythril classic: Security analysis tool for ethereum smart contracts," 2019.
- [7] Remix, "Remix documentation," 2018.
- [8] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Gollamudi, G. Gonthier, N. Kobeissi, N. Kulatova, A. Rastogi, T. Sibut-Pinote, N. Swamy, *et al.*, "Formal verification of smart contracts: Short paper," in *Proceedings of the 2016 ACM Workshop on Programming Languages and Analysis for Security*, pp. 91–96, ACM, 2016.
- [9] S. Tikhomirov, E. Voskresenskaya, I. Ivanitskiy, R. Takhaviev, E. Marchenko, and Y. Alexandrov, "Smartcheck: Static analysis of ethereum smart contracts," 2018.
- [10] S. Kalra, S. Goel, M. Dhawan, and S. Sharma, "Zeus: Analyzing safety of smart contracts," NDSS, 2018.
- [11] "Clang libtooling," 2019. <https://clang.llvm.org/docs/LibTooling.html>.
- [12] "Soltar," 2016. <https://github.com/duaraghav8/soltar>.
- [13] "Porosity," 2017. <https://github.com/comaeio/porosity>.
- [14] P. Software, "Ethereum smart contract decompiler," 2019.
- [15] B. G. Ryder, "Constructing the call graph of a program," *IEEE Transactions on Software Engineering*, no. 3, pp. 216–226, 1979.
- [16] F. E. Allen, "Control flow analysis," in *ACM Sigplan Notices*, vol. 5, pp. 1–19, ACM, 1970.
- [17] M. Geimer, S. S. Shende, A. D. Malony, and F. Wolf, "A generic and configurable source-code instrumentation component," in *International Conference on Computational Science*, pp. 696–705, Springer, 2009.
- [18] M. M. Tikir and J. K. Hollingsworth, "Efficient instrumentation for code coverage testing," in *ACM SIGSOFT Software Engineering Notes*, vol. 27, pp. 86–96, ACM, 2002.
- [19] J. Ellson, E. Gansner, L. Koutsofios, S. C. North, and G. Woodhull, "Graphviz open source graph drawing tools," in *International Symposium on Graph Drawing*, pp. 483–484, Springer, 2001.

⁴The contract *Expiry* is available at <https://etherscan.io/address/0x0ecec224fbc24d40b446c6a94a142dc41fae76f2d>