

# SigRec: Automatic Recovery of Function Signatures in Smart Contracts

Ting Chen, Zihao Li, Xiapu Luo, Xiaofeng Wang, Ting Wang, Zheyuan He, Kezhao Fang, Yufei Zhang, Hang Zhu, Hongwei Li, Yan Cheng, and Xiaosong Zhang

**Abstract**—Millions of smart contracts have been deployed onto Ethereum for providing various services, whose functions can be invoked. For this purpose, the caller needs to know the *function signature* of a callee, which includes its function id and parameter types. Such signatures are *critical* to many applications focusing on smart contracts, e.g., reverse engineering, fuzzing, attack detection, and profiling. Unfortunately, it is challenging to recover the function signatures from contract bytecode, since neither debug information nor type information is present in the bytecode. To address this issue, prior approaches rely on source code, or a collection of known signatures from incomplete databases or incomplete heuristic rules, which, however, are far from adequate and cannot cope with the rapid growth of new contracts. In this paper, we propose a novel solution that leverages how functions are handled by Ethereum virtual machine (EVM) to automatically recover function signatures. In particular, we exploit how smart contracts determine the functions to be invoked to locate and extract function ids, and propose a new approach named *type-aware* symbolic execution (TASE) that utilizes the semantics of EVM operations on parameters to identify the number and the types of parameters. Moreover, we develop SigRec, a new tool for recovering function signatures from contract bytecode without the need of source code and function signature databases. The extensive experimental results show that SigRec outperforms all existing tools, achieving an unprecedented 98.7% accuracy within 0.074 seconds. We further demonstrate that the recovered function signatures are useful in attack detection, fuzzing and reverse engineering of EVM bytecode.

**Index Terms**—smart contract, function signature, Ethereum, automatic recovery, type-aware symbolic execution.



## 1 INTRODUCTION

More than 30 million smart contracts have been deployed on Ethereum, the largest contract-hosting blockchain. An Ethereum *smart contract* is a program (i.e., a collection of code and data) running on Ethereum virtual machine (EVM), which can be accessed through its address in Ethereum [1]. A smart contract is typically written in a high-level language (e.g., Solidity [2]) and then compiled into EVM bytecode. After the bytecode is deployed, its public and external functions can be invoked. To invoke a contract function, the caller needs to know its function signature which consists of a function id and the list of parameter types [3]<sup>1</sup>. A *function id* refers to the first 4 bytes in the Keccak-256 hash of a function name and the list of parameter types [3].

Function signatures play a critical role in many applications focusing on smart contracts because their functions (in terms of function signatures) need to be first identified before their behaviors can be evaluated. For example, some studies recognize wallet contracts [4], token contracts [5], [6],

[7], and the contracts for Ethereum name service [6] based on function signatures. §6 shows that function signatures can be used to enhance the results of reverse engineering results contract bytecode by adding types, meaningful variable names and simplifying the code for accessing parameters. Besides, knowing the parameter list of a target function, a fuzzer can strategically mutate the test cases [8], [9] for better vulnerability detection. Our experimental results (§6) show that function signatures enable ContractFuzzer [8] to find 23% more bugs than it could without these signatures. Function signatures are also important to the detection of attacks against smart contracts. A prominent case is the *short address* attack that exploits an EVM vulnerability in handling a malformed actual argument of the address type, which is shorter than a valid address [10]. We develop ParChecker, a new tool based on the recovered function signatures to detect malformed actual arguments including short address attacks (§6).

**Challenges in signature discovery.** It is challenging to recover function signatures from bytecode since there is no debug information in contract bytecode and the bytecode is *untyped*, i.e., parameters are kept in 256-bit words without type information [11]. All existing approaches fail to effectively recover function signatures (§5.6). More precisely, an intuitive method is to extract the function signatures from the source code of smart contracts, which, however, is only available for a small proportion of contracts (< 1% until 2017 [12]). Another method retrieves signatures from existing databases, like Ethereum Function Signature Database (EFSDB) [13], through function ids. For example, Gigahorse [14], Eveem [15], and Online Solidity Decompiler

- Ting Chen (brokendragon@uestc.edu.cn), Zheyuan He, Kezhao Fang, Yufei Zhang, Hang Zhu, Hongwei Li and Xiaosong Zhang are with University of Electronic Science and Technology of China, China.
- Zihao Li and Xiapu Luo (csxluo@comp.polyu.edu.hk) are with Hong Kong Polytechnic University, Hong Kong. (Corresponding Author: Xiapu Luo.)
- Xiaofeng Wang is with Indiana University, Bloomington, USA.
- Ting Wang is with Pennsylvania State University, USA.
- Yan Cheng is with Ant Group, China.  
Postal Code: 310013

1. A function signature defined by Ethereum includes a function name and the list of parameter types [3]. Since function id instead of function name is required to invoke a contract function, in this paper we just consider function id, parameter number and parameter types.

(OSD) [16] rely on EFSD, while EVM Bytecode Decompiler (EBD) [17] and JEB [18] maintain their own databases. However, these databases are incomplete, only covering 31.7% of the function signatures in the wild (Supplementary material A) and quickly become out of date without constantly updating.

Alternatively, one could reverse engineering a function id by enumerating all possible parameter type combinations to compute a hash, together with a possible function name. Given the huge space of possible combinations, this attempt is fragile in general. Abi Decompiler [19] adopts this approach, which only covers 12.3% of function signatures (Supplementary material A). Finally, one could derive function signatures using some heuristics: Gigahorse infers the number of parameters based upon the number of items that a function pops from stack [14]; Eveem regards a type as an array if its layout contains an *offset* field and a *num* field. Although this direction is promising, existing approaches turn out to be ineffective due to the incompleteness of their heuristics. For example, Eveem can only achieve an accuracy of 58.1% and 18.3% to infer the function signatures in closed-source smart contracts and synthesized contracts, respectively (§5.6), even it leverages an existing database, EFSD.

**Our work.** In this paper, we propose a new solution called `SigRec` to automatically recover function signatures from EVM bytecode compiled from two mainstream compilers (i.e., Solidity and Vyper) for smart contracts *without* the need of source code and other databases. Our key observation is that even in the absence of type information, the way EVM bytecode handles a function call and its inputs uniquely characterizes different parameter types. For example, a `uint32` argument will be extracted using a 32-bits mask applied to the input data (R11, §3.4). Based upon this observation, we first generalize the semantics of such type-related operations into rules (§3). Then, we design *type-aware* symbolic execution (TASE) to explore the EVM instructions that manipulate parameters, and use the rules for inferring the types of parameters (§4). Since the EVM instructions handling parameters are typically near a function’s entry point, TASE can handle them very effectively and efficiently to keep up with the ever-growing Ethereum smart contracts.

Studies on reverse engineering of variable types from binary executables by leveraging the semantics of instructions [20], [21], [22], [23], [24], [25], [26], [27], [28], [29] are related to our work. However, TASE has several differences with existing approaches. First, to recover complicated types that cannot be handled by existing techniques, TASE infers how variables are affected by parameters through symbolic execution (SE). For example, to identify a multidimensional array, TASE checks whether the read location is calculated by adding the value of the *offset* field to the base (§3.2). Second, semantic knowledge summarized from binary code cannot be applied to EVM bytecode, due to architecture differences (§8) and unique parameter types defined in smart contracts, e.g., 256-bit integers (no more than 64 bit in x64 binaries), the 20-byte address type, and reversed array notation (§2.3). New knowledge, therefore, needs to be discovered to support type inference through TASE. Third, by exploiting the fact that the code responsible for handling parameters is usually around a function’s entry point, TASE

runs extremely fast (0.074s to recover one function signature on average, §5.4).

We conduct extensive experiments to evaluate `SigRec` and compare it with all existing approaches. First, we collect all 119,404 unique (i.e., duplicates are eliminated) open-source smart contracts on Ethereum, which include 210,869 public/external functions with unique function signatures, and use them to evaluate the accuracy of `SigRec`. The experimental result shows that `SigRec` achieves an average accuracy of 98.7% (§5.2), and the accuracy never goes below 96% across all compilers (from V 0.1.1 to V 0.8.0) with or without optimization (§5.3). Second, we compare `SigRec` with all existing approaches, i.e., OSD [16], EBD [17], JEB [18], Gigahorse [14] and Eveem [15]. The experimental results show that `SigRec` correctly recovers much more signatures, outperforming them by at least 22.5 %, 40.1% and 80.5% in processing open-source, closed-source and synthesized smart contracts, respectively (§5.6). Manual investigation shows that for many function signatures, existing approaches report wrong types, produce *nonexistent* types, output an *unspecific* type, add *nonexistent* parameters, miss parameters or fail to generate function signatures (§5.6). Third, we apply `SigRec` to the bytecode of all 37,009,570 smart contracts deployed on Ethereum with 47,329,149 public/external functions, and found that it only takes 0.074 seconds on average to recover a function signature (§5.4).

We further demonstrate how the recovered function signatures can enable the detection of stealthy short address attacks, empower the state-of-the-art fuzzer to discover more vulnerabilities, and enhance the result of reverse engineering contract bytecode (§6).

**Contributions.** The major contributions of the paper are summarized as follows:

- We propose a novel solution that exploits the semantics of EVM instructions to correctly and efficiently recover function signatures from the bytecode of smart contracts.
- We develop `SigRec` based on our new solution, and conduct extensive experiments to evaluate it. The experimental results show that `SigRec` achieves nearly 100% accuracy in signature recovery, under different compiler versions, within 0.074 seconds on average, and it significantly outperforms existing methods. We have deployed `SigRec` as an online web service <http://bit.ly/SigRecWS> and will release its code after paper publication.
- We use three important applications to demonstrate the usage of recovering function signatures, including attack detection, fuzzing and reverse engineering of EVM bytecode.

## 2 ACCESSING PATTERNS OF PARAMETERS

This section first introduces how function invocation is performed in EVM and then elaborate more on the EVM instructions and variables involved in a function invocation. The concepts of account, bytecode of smart contracts, and public/external function can be found in Supplementary material B.

**Function invocation.** To invoke a public/external function, a message will be sent by an EOA account or a smart contract account [30]. The message contains the address of the smart contract whose function will be invoked and the *call data* field which indicates the function to be invoked



CALLDATALOAD is needed to read a  $\text{bytes}\langle M \rangle$  from the call data, and an AND is used for masking. After masking, the higher-order bytes of a  $\text{bytes}\langle M \rangle$  is retained. By contrast, for a  $\text{uint}\langle M \rangle$ , the masking retains the lower-order bytes since a uint is extended on the higher-order side. Since  $\text{bytes32}$  and  $\text{uint256}$  are not extended, we differentiate them by exploiting the fact that a BYTE instruction is used for accessing a single byte of a  $\text{bytes32}$ , while an AND masks a  $\text{uint256}$  for the same purpose.

**2. Array:** An array hosts array items of the same basic type. An array can be static, dynamic or nested. The size of a static array is known during compilation whereas the size of a dynamic array and a nested array is determined by the actual argument at runtime.

**(1) Static array.** We denote it as  $T[X_1]\dots[X_n]$ , where  $T$  is one of the basic types,  $X_1, \dots, X_n$  are constant numbers known in compilation, and  $n$  is the dimension. If  $n > 1$ ,  $T[X_1]\dots[X_n]$  is a multidimensional array. The size of a static array is known in compilation since the number of items in each dimension is fixed. As a basic type, each array item is extended through the process described above. Different from other languages (e.g., C), the notation of an array in EVM is reversed and the access is in the opposite direction of the declaration [32]. For example,  $\text{uint256}[3][2] x$  means an array  $x$  of two arrays of three  $\text{uint256}$ , and  $x[1][0]$  accesses the first item of the second  $\text{uint256}[3]$ . All items of a static array are stored consecutively in the call data. Fig. 5 shows the call data layout of a public/external function with one  $\text{uint256}[3][2]$  argument.




Fig. 5. The call data layout of a static array,  $\text{uint256}[3][2]$

The accessing patterns of a static array are different for public and external functions.

(a) *Public function.* Its static array will be read into the memory by one or more CALLDATACOPYs. Then, each array item is accessed by an MLOAD, which reads 32 bytes from the memory to the stack [30]. We find that a nested loop is used to read parameters from the call data, and the level of the nested loop is equal to the dimension of the static array minus one. In the most inner layer, a CALLDATACOPY is executed. A CALLDATACOPY can read a one-dimensional static array.

Listing 1 shows the pseudocode (we do not show the EVM instructions for the ease of presentation) to read a  $\text{uint256}[3][2]$  from the call data to the memory. The execution of each CALLDATACOPY copies the array of the lowest dimension (Line 3). The loop guard (Line 2) determines how many times the CALLDATACOPY should be executed. The loop guard is compiled into an LT instruction, which checks whether the current loop count is smaller than the item number of the highest dimension.

```

1 i = 0;
2 while (i < 2){ //how many uint256[3] needs to copy
3   calldatacopy(uint256[3]);
4   i++;
5 }

```

Listing 1. Pseudocode for reading a  $\text{uint256}[3][2]$  parameter to memory

(b) *External function.* Its static array will not be entirely copied. Instead, array items will be read from the call data to the stack using CALLDATALOAD on demand. We find that if the smart contract is compiled without optimization or the index of the array is a variable, the CALLDATALOAD is preceded by bound checks to prevent array overrun. For instance, before accessing the item  $x[i][j]$  of the array  $x[3][2]$ , two bound checks  $i < 2$  and  $j < 3$  are executed. Only after they are passed,  $x[i][j]$  can be accessed. Such runtime bound check is not needed if the array index is a constant and the smart contract is compiled with optimization due to the compile-time bound checks.

**(2) Dynamic array.** We denote it as  $T[X_1]\dots[X_{n-1}][]$ , where  $T$  is one of basic types,  $X_1, \dots, X_{n-1}$  are constant numbers known in compilation,  $n$  is the dimension. The size of a dynamic array is unknown in compilation because the item number of the highest dimension is unknown. Instead, the item number should be provided in the call data so that the invoked function can know it at runtime. Fig. 6 shows the call data layout of a dynamic array  $\text{uint256}[3][]$ , which is the first parameter, and the actual argument is  $\text{uint256}[3][2]$ . A 32-bytes value, i.e., the *offset* field, is located right after the function id, and this value is an offset relative to the first byte after the function id. The 32 bytes, i.e., the *num* field, pointed by the offset stores the number of items in the highest dimension. All array items are located after the *num* field.




Fig. 6. The call data layout of a dynamic array  $\text{uint256}[3][]$ , and the actual argument is  $\text{uint256}[3][2]$

The accessing patterns of a dynamic array are different for public and external functions.

(a) *Public function.* Its dynamic array will be read into the memory by one or more CALLDATACOPYs. More precisely, the smart contract uses a CALLDATALOAD to read the offset (i.e., the *offset* field), and another CALLDATALOAD to read the item number of the highest dimension (i.e., the *num* field). Then, the item number will be read from the stack to the memory by an MSTORE, which saves 32 bytes from the stack into the memory [30]. After that, the smart contract uses a nested loop to copy all array items into the memory right after the item number. Only one CALLDATACOPY is needed to read a one-dimensional dynamic array.

(b) *External function.* Two CALLDATALOADs are used to read the offset and the item number of the highest dimension. Different from the process for public functions, an array item will be read from the call data to the stack by a CALLDATALOAD on demand. The operand of the CALLDATALOAD, i.e., the location from where to copy, is computed from the value of the *offset* field at runtime, and the result contains the multiplication of 32, because each array item is extended to 32 bytes. For example, the location of the third array item is computed by  $\text{offset} + 4 + 32 + 2 \times 32$  (the function id is of 4 bytes, an item number is of 32 bytes, the first two array items

are of 64 bytes). The bound check of the highest dimension exists since its item number is unknown in compilation.

**(3) Nested array.** We denote a  $n$ -dimensional nested array as  $T[X_1] \dots [X_n]$ , where  $T$  is one of basic types. At least one of  $X_1$  to  $X_{n-1}$  should be empty indicating the corresponding dimension is dynamic. The key difference between a  $n$ -dimensional nested array and a  $n$ -dimensional dynamic array is that at least one dimension of the lower  $n - 1$  dimensions of the former can be dynamic whereas the lower  $n - 1$  dimensions of the latter must be static [3]. Hence, each dimension of a  $n$ -dimensional nested array associates an *offset* field and a *num* field to enable dynamic dimension size. Fig. 7 shows the call data layout of a nested array `uint[][]`, which is the first parameter of a function. Let's assume that the actual argument is `[[0x1, 0x2], [0x3]]` for the ease of explanation. *offset<sub>1</sub>* is the *offset* field of this nested array, and *num<sub>1</sub>* is the *num* field of its highest dimension. Two items (i.e., `[0x1, 0x2]` and `[0x3]`) in the highest dimension are placed after *num<sub>1</sub>*. Because both the two items are dynamic arrays, their *offset* fields will be placed immediately after *num<sub>1</sub>*. In this example, *offset<sub>2</sub>* and *num<sub>2</sub>* are the *offset* and *num* fields of the item `[0x1, 0x2]`, and *offset<sub>3</sub>* and *num<sub>3</sub>* are the *offset* and *num* fields of the item `[0x3]`. `0x1` and `0x2` which are two items of the array `[0x1, 0x2]` are placed immediately after *num<sub>2</sub>*. Similarly `0x3`, the only item of the array `[0x3]` is placed immediately after *num<sub>3</sub>*.

The accessing pattern of a nested array in the public mode is the same as that in the external mode, and an array item in nested array will be read from the call data to the stack by a `CALLDATALOAD` on demand. More specifically, two `CALLDATALOAD`s are used to read *offset<sub>1</sub>* and *num<sub>1</sub>*, the item number of the highest dimension. For example, to read the array item `0x3` as an example, three additional `CALLDATALOAD`s are required to read *offset<sub>3</sub>*, *num<sub>3</sub>* and the item `0x3`. To access an array item, there is a bound check for each dimension.




Fig. 7. The call data layout of a nested array `uint[][]`, and the actual argument is `[[1, 2], [3]]`

**3. bytes:** It is a dynamic byte sequence whose size is determined at runtime [3]. If a public/external function has a `bytes` type parameter, it will be extended on the lower-order (right) side with the minimum number of zero to make the length of the extended result be multiple of 32 bytes [3]. For example, if an un-extended `bytes` is `'abcd'`, the extended result is shown in Fig. 4. The layout of a `bytes` is similar to that of a one-dimensional dynamic array, as shown in Fig. 6, except that the *num* field stores the size (in byte) of the `bytes` before padding. Note that the *num* field of a one-dimensional dynamic array records the number of array items. The extended value is located right after the *num* field.

The accessing patterns of a `bytes` value are different for public and external functions.

(a) *Public function.* The accessing pattern of a bytes in a public function is the same as the accessing pattern of a one-dimensional dynamic array in a public function except the computation of read size, because every single byte of a bytes is not extended.

(b) *External function.* The accessing pattern of a bytes in the external mode is the same as the accessing pattern of a one-dimensional dynamic array in the external mode, except that accessing a single byte of a bytes does not need the multiplication of 32 because every single byte of a bytes is not extended.

**4. string:** It is a dynamic Unicode string, whose size is determined at runtime [3]. The call data layout of a string and reading a string from the call data are the same as that of a bytes. We distinguish them based on the observation that a bytes supports reading/writing its individual byte whereas a string does not support such operation.

**5. struct:** We denote a struct as  $(T_1, \dots, T_n)$ , where  $n$  is a constant number, indicating  $n$  struct items whose types can be any types in Solidity. The size of a struct can be dynamic or static. For a static struct whose size is known before compilation, the type of each struct item should be one of basic types, static array, or static struct. A dynamic struct whose size is unknown in compilation can host all possible Solidity types. The accessing pattern of a struct in the public mode is the same as that in the external mode. We find that the layout of a static struct is the same as that of all its items as if they were not inside the static struct. For example, Listing 2 is the definition of a struct `(uint256, uint256)` in Solidity, Listing 3 is the definition of a function with two `uint256` parameters in Solidity. Fig. 8 shows the same call data layout of them. Moreover, the bytecode to read an item of a static struct is the same as the bytecode to read the item as if it was not placed inside the struct. Therefore, there is no sufficient hint from the bytecode to distinguish these two different situations.

```
1 struct StructSample {
2     uint256 a;
3     uint256 b;
4 }
5 function Func(StructSample var1) {}
```

Listing 2. A function taking in a struct with two `uint256` items

```
1 function FuncSample(uint256 a, uint256 b) {}
```

Listing 3. A function taking in two `uint256` parameters




Fig. 8. The layout of a struct containing two `uint256` items is the same to that of two individual `uint256` parameters

The layout of a dynamic struct has an *offset* field, and all struct items are placed after the *offset* field. Fig. 9 shows the call data layout of a dynamic struct `(uint[], uint)`, which is the first parameter of a function, and we assume that the actual argument is `([0x1, 0x2], 0x3)`. In this example, *offset<sub>1</sub>* is the *offset* field of this struct. Two items (i.e., `[0x1, 0x2]` and `0x3`) of this struct are placed after the position which *offset<sub>1</sub>* indicated. Because the item `[0x1, 0x2]` is array, its *offset* field will be placed immediately after the position which *offset<sub>1</sub>* indicated, and the value of item `0x3` will be placed immediately after it. In this example, *offset<sub>2</sub>* and *num<sub>1</sub>* are the *offset* and *num* fields of item `[0x1, 0x2]`, and items of



[0x1, 0x2] are placed immediately after  $num_1$ . For example, to access the struct item 0x3, one CALLDATALOAD is used for reading  $offset_1$  and one additional CALLDATALOAD is required to read the item 0x3 directly.




Fig. 9. The call data layout of a dynamic struct (uint[],uint), and the actual argument is ([1, 2], 3)

### 2.3.2 Vyper

Vyper supports ten parameter types, including bool, int128, uint256, address, bytes32, decimal, fixed-size list, fixed-size byte array, fixed-size string, and struct [33]. The first five types are also supported by Solidity, and the layouts of these five types are the same with the layouts of Solidity types. Different with Solidity, the instructions to read these five types use comparison instructions (e.g., LT) rather than mask instructions (e.g., AND) to ensure that the values of these five types are in the valid ranges. For example, Listing 4 shows the accessing pattern for address parameter in Solidity and Listing 5 shows the accessing pattern for address parameter in Vyper. We also learn that Vyper generates the same bytecode for public and external functions. We describe the remaining five types below.

```
1 CALLDATALOAD //read the parameter
2 ...
3 PUSH20 0xffffffffffffffffffffffffffffffffffffffff
4 AND //mask by 0xff...ff
```

Listing 4. Accessing pattern for an address parameter in Solidity

```
1 PUSH21 0x0100000000000000000000000000000000000000000000000000000000000000
2 PUSH1 0x20
3 MSTORE
4
5 CALLDATALOAD //read the parameter
6 PUSH1 0x20
7 MLOAD //load 0x01...00
8 DUP2
9 LT //compare to 0x01...00
```

Listing 5. Accessing pattern for an address parameter in Vyper

**1. decimal.** A decimal is a fixed-point value with a precision of 10 decimal, whose value ranges from  $-2^{127}$  and  $2^{127} - 1$  [33]. We find that the layout and the accessing pattern of a decimal is similar to that of signed integer in Solidity, except that two comparisons rather than a SIGNEXTEND mask are used for ensuring that the value of the decimal is between  $-2^{127}$  and  $2^{127} - 1$ . If the decimal value is out of the range, the execution of smart contract will be aborted.

**2. Fixed-size list.** Fixed-size list records a constant number of items. [33]. We denote a fixed-size list as  $T[X_1] \dots [X_n]$ , where  $T$  is one of bool, int128, uint256, address, bytes32 and decimal.  $X_1, \dots, X_n$  are constant numbers known in compilation, and  $n$  is the dimension and  $n \geq 1$ . The layout and accessing pattern of a fixed-size list is the same as that of a static array in the external mode of Solidity. To access a list item, bound checks are used to prevent the overrun of the list if the index is a variable.

**3. Fixed-size byte array.** A fixed-size byte array is a byte sequence with a maximum length, which is given in the source code. Its real length which is provided at runtime, should

not be longer than the maximum size. We denote a fixed-size byte array as  $bytes[maxLen]$ ,  $maxLen$  is a constant value denoting the maximum length. The layout and accessing pattern for a fixed-size byte array are similar to a bytes in the public mode of Solidity, except that 32 (the size of the  $num$  field) +  $maxLen$  are read from the call data. That is, the extended bytes of a fixed-size byte array are not read.

**4. Fixed-size string.** We denote a fixed-size string as  $string[maxLen]$ , where  $maxLen$  is a constant value representing the maximum length of the string. The layout and the way to read a fixed-size string from the call data are the same as that of a fixed-size byte array. The difference between the two types lies in that a fixed-size byte array allows accessing its individual byte whereas a fixed-size byte array does not allow it.

**5. struct.** struct is a container type that can host several variables. We denote a struct as  $(T_1, \dots, T_n)$ , where  $n$  is a constant value indicating  $n$  struct items and  $T_i$  can be one of the following types, bool, int128, uint256, address, bytes32 and decimal [33]. The layout of a struct is the same as that of all its items as if they were not placed inside the struct. The bytecode to read a struct item is also the same as the bytecode to read the item as if it was not inside the struct. Therefore, there is no sufficient hint from the bytecode to distinguish these two different types. Listing 6 shows a function signature which takes in a struct containing two uint256 items, and Listing 7 presents a function taking in two uint256 parameters. The layouts of the struct and two individual uint256 parameters are the same, as shown in Fig. 10.

```
1 struct StructSample:
2   a: uint256
3   b: uint256
4 def Func(var1: StructSample)
```

Listing 6. A function taking in a struct with two uint256 items

```
1 def FuncSample(a: uint256, b: uint256)
```

Listing 7. A function taking in two uint256 parameters




Fig. 10. The layout of a struct containing two uint256 items is the same to that of two individual uint256 parameters

## 3 RULES FOR TYPE INFERENCE

This section first introduces the method to generate rules, and then describes the first four rules. We briefly introduce the remaining 27 rules, and elaborate more on them in Supplementary material C.

### 3.1 Rules Generation

We propose a systematic approach to generate rules. The basic idea is to first learn the accessing patterns automatically from the bytecode of self-generated smart contracts and then summarize the rules manually, which are used by SigRec to recover function signatures from the bytecode of other smart contracts. Our approach has five steps. Since the first four steps are automated, it can be easily extended to handle new accessing patterns due to new compilers/programming language/obfuscation techniques and new parameter types (discussed in §7).

**Step 1. Preparing smart contracts and their bytecode.** We develop a tool to automatically construct smart contracts and compile them into bytecode, from which the rules will be learned. We develop smart contracts in Solidity and Vyper since they are the most popular languages for developing Ethereum smart contracts. To ensure the completeness of the derived rules, we consider all parameter types supported in Solidity and Vyper, all major versions of compilers as well as various optimization levels. More specifically, each smart contract contains one public or external function. Since we find that smart contracts handle the different parameters of a function independently, each generated function just has one parameter to ease rule learning. To expose the accessing patterns of parameters, the body of each generated function contains statements to access the parameter, e.g., assignments, arithmetic operations, comparisons, reading one byte from bytes and bytes32, and reading one array item.

We consider all basic types with all possible widths, e.g., uint8, uint16, ..., uint256. Since the upper bound of the array dimension and the upper bound of the size of a static dimension are undocumented, we first identify the patterns with concrete number of dimensions and then generalize the pattern for all possible dimensions. Specifically, we set array dimensions from 1 to 5 and each dimension can be static or dynamic. For each static dimension, we set its size from 1 to 10. The items of an array are set to one of supported basic types with all possible widths (e.g., uint8, uint16, ..., uint256).

**Step 2. Collecting the accessing pattern of each parameter type.** We conduct data dependence analysis and control dependence analysis to locate the EVM instruction sequence used to access the parameter of each generated function, which are required in subsequent steps. We consider such instruction sequence as an accessing pattern. Specifically, the instructions which are data dependent with parameters are needed in step 4 for collecting the symbolic expressions of variables. Step 5 needs the control dependency relations to infer the structure of some complicated parameter types. For example, TASE infers the loop structure from the EVM instructions that are control dependent on the instruction for reading the call data to discover  $n$ -dimensional dynamic array.

**Step 3. Extracting common accessing patterns.** We observe that the access patterns for some parameters types (e.g., uint8 and uint16) are similar, and thus this step extracts common access patterns from the accessing patterns obtained in step 2. We regard an instruction sequence as the *common* accessing pattern of some accessing patterns if it appears in all these accessing patterns. For example, from step 2 we obtain the accessing patterns of uint8, uint16, ..., uint256 in Solidity, and then in this step we try to extract the common accessing pattern from these patterns to facilitate the derivation of the rule for uint with all possible widths. On the contrary, bool, int128, uint256, address, bytes32, decimal of Vyper are basic types, we skip the step for them because these types have fixed widths. Besides, we omit the description about the fixed-size list of Vyper because it is equivalent to the static array of Solidity.

*Basic types of Solidity.* We use unsigned integers as an example, and the other basic types are handled similarly. We extract the common accessing pattern from the accessing

patterns of uint8, uint16, ..., uint256, which is used in step 5 to generate a general rule for inferring the unsigned integers with all possible widths,  $\text{uint}\langle M \rangle$ ,  $8 \leq M \leq 256$ ,  $M\%8 == 0$ .

*One-dimensional static array of Solidity.* We extract the common accessing pattern from the accessing patterns of  $\text{uint8}[1]$ ,  $\text{uint8}[2]$ , ...,  $\text{uint8}[10]$ . By retaining the instructions in the common accessing pattern but not in the accessing pattern of uint8, we further obtain the accessing pattern of  $T[N]$ ,  $1 \leq N \leq 10$ ,  $T$  is one of the basic types.

*Multidimensional static array of Solidity.* We first extract the common accessing pattern from the accessing patterns of  $\text{uint8}[N_1][N_2]$ ,  $1 \leq N_1, N_2 \leq 10$ . Then we obtain the accessing pattern of  $T[N_1][N_2]$ ,  $1 \leq N_1, N_2 \leq 10$  by retaining the instructions in the common accessing pattern but not in the accessing pattern of uint8. Similarly, we obtain the accessing patterns of  $T[N_1][N_2][N_3]$ ,  $T[N_1][N_2][N_3][N_4]$ ,  $T[N_1][N_2][N_3][N_4][N_5]$ ,  $1 \leq N_1, N_2, N_3, N_4, N_5 \leq 10$ .

*One-dimensional dynamic array of Solidity.* We regard the EVM instructions that are in the accessing pattern of  $\text{uint8}[]$  but not in the accessing pattern of uint8 as the common accessing pattern of  $T[]$ .

*Multidimensional dynamic array of Solidity.* We extract the common accessing pattern from accessing patterns of  $\text{uint8}[N_1][]$ ,  $1 \leq N_1 \leq 10$ . By retaining the EVM instructions that appear in such common accessing pattern but not in the accessing pattern of uint8, we obtain the accessing pattern of  $T[N_1][]$ ,  $1 \leq N_1 \leq 10$ . We also obtain the accessing patterns of  $T[N_1][N_2][]$ ,  $T[N_1][N_2][N_3][]$ ,  $T[N_1][N_2][N_3][N_4][]$ ,  $1 \leq N_1, N_2, N_3, N_4 \leq 10$  through the same process.

*Nested array of Solidity.* We extract the common accessing pattern from the accessing patterns of  $\text{uint8}[][][N_1]$ ,  $1 \leq N_1 \leq 10$  or  $N_1$  is empty. By retaining the EVM instructions that appear in such common accessing pattern but not in the accessing pattern of uint8, we obtain the accessing pattern of  $T[][][N_1]$ ,  $1 \leq N_1 \leq 10$  or  $N_1$  is empty. We also obtain the accessing pattern of  $T[N_1][N_2][N_3][N_4][N_5]$ ,  $1 \leq N_1, N_2, N_3, N_4, N_5 \leq 10$  and at least one of  $N_1$  to  $N_4$  is empty through the same process.

*Fixed-size byte array of Vyper.* We obtain the accessing pattern of  $\text{bytes}[\text{maxLen}]$  by extracting the common accessing pattern from the accessing patterns of  $\text{bytes}[1]$ ,  $\text{bytes}[2]$ , ...,  $\text{bytes}[50]$ .

*Fixed-size string of Vyper.* We obtain the accessing pattern of  $\text{string}[\text{maxLen}]$  by extracting the common accessing pattern from the accessing patterns of  $\text{string}[1]$ ,  $\text{string}[2]$ , ...,  $\text{string}[50]$ .

*struct of Solidity.* we extract the accessing pattern of  $(\text{uint8}[])$ . By retaining the EVM instructions that appear in such accessing pattern but not in the accessing pattern of  $\text{uint8}[]$ , which is extracted above, we obtain the accessing pattern of the struct type of Solidity.

*struct of Vyper.* we extract the common accessing pattern from accessing patterns of  $(\text{uint256})$ . By retaining the EVM instructions that appear in such common accessing pattern but not in the accessing pattern of  $\text{uint256}$ , we obtain the accessing pattern of the struct type of Vyper.

**Step 4. Generating symbolic expressions from parameters.** To characterize how parameters are handled by EVM instructions, we conduct symbolic execution on the common accessing patterns obtained in step 3 by treating the call data as symbols, and collect the symbolic expressions for each

variable. Step 5 needs symbolic expressions to summarize rules.

**Step 5. Summarizing rules.** We summarize the rules according to the common accessing patterns obtained in step 3 and the symbolic expressions collected in step 4, and then organize them into a decision tree (Fig. 13) for fast rule checking. To determine the rules for arrays, we investigate how the accessing patterns change when its dimension increases from 1 to 5 and the number of items per-dimension increases from 1 to 10. Eventually, we generate 18 rules and divide them into three categories for CALLDATALOAD (§3.2), CALLDATACOPY (§3.3), and other instructions (§3.4), respectively.

### 3.2 Rules for CALLDATALOAD

A CALLDATALOAD can read 3 kinds of data from the call data (§2): (1) the value of a parameter; (2) the *offset* field of a parameter; (3) the *num* field of a parameter.

**R1:** R1 is used to infer a dynamic array/bytes/string. It holds if two CALLDATALOADs (termed by CALLDATALOAD<sub>1</sub> and CALLDATALOAD<sub>2</sub>) satisfy that  $x = \text{CALLDATALOAD}_1(\text{loc}) \wedge y = \text{CALLDATALOAD}_2(x + 0x4)$ , where  $x = \text{CALLDATALOAD}_1(\text{loc})$  means that a 32-bytes value is read from the *loc* byte of the call data into  $x$  by a CALLDATALOAD. Actually,  $x$  is the *offset* field of a dynamic array/bytes/string parameter. CALLDATALOAD<sub>2</sub> reads the *num* field of the parameter, because the location of the *num* field is the value of the *offset* field plus the length of the function id (i.e., 4 bytes).

**R2:** R2 is used to infer an  $n$ -dimensional ( $n > 1$ ) dynamic array in an external function. It holds if three requirements,  $v1$ ,  $v2$ ,  $v3$  are satisfied. Let  $\text{exp}(\text{loc})$  represent the symbolic expression of the location *loc*, which describes how *loc* is computed from symbols. Let  $\text{exp}(p) \diamond q$  denote that the symbolic expression of  $p$  (e.g.,  $x + y \times 5$ ) contains the symbolic expression  $q$  (e.g.,  $y \times 5$ ).  $v1$  is defined as  $\text{exp}(\text{loc}) \diamond (\text{offset} +)$ , meaning that the read location is computed by adding the value of the *offset* field, because array items are placed after the *num* field, which is pointed by the *offset* field. Hence, if a CALLDATALOAD reads an item from a dynamic array in the external model,  $v1$  is satisfied.

$v2$  is defined as  $\text{exp}(\text{loc}) \diamond (32 \times)$ , meaning that the symbolic expression of *loc* contains the multiplication of 32. If a CALLDATALOAD reads an item from a dynamic array parameter in an external function,  $v2$  is satisfied, because each array item is extended to 32 bytes so that the multiplication of 32 is needed to access an array item.

Before introducing  $v3$ , we let  $\text{num}_n = \text{CALLDATALOAD}(\text{num})$  indicate that a CALLDATALOAD reads the value of the *num* field and assigns it to  $\text{num}_n$ .  $\text{num}_{n-1}, \dots, \text{num}_1$  are constant numbers, and  $\text{LT}_j(i_j, \text{num}_j)$  means that an LT instruction compares a number  $i_j$  with a number  $\text{num}_j$ . Let  $\text{ins}_i \triangleleft \text{ins}_j$  indicate that the instruction  $\text{ins}_j$  is control-flow dependent on the instruction  $\text{ins}_i$ .  $v3$  is defined as  $\text{LT}_n(i_n, \text{num}_n) \triangleleft \text{LT}_{n-1}(i_{n-1}, \text{num}_{n-1}) \triangleleft \dots \triangleleft \text{LT}_1(i_1, \text{num}_1) \triangleleft \text{CALLDATALOAD}(\text{loc})$ .  $v3$  is satisfied if the CALLDATALOAD for reading the array item is within a nested loop. Hence, if a CALLDATALOAD(*loc*) reads an item from a dynamic array in an external function,  $v3$  is satisfied, because  $n$  bound checks (i.e.,  $n$  LTs) for preventing array overrun are located between the two CALLDATALOADs.

Hence, if  $v1$ ,  $v2$ , and  $v3$  are all fulfilled, the dynamic array is  $n$ -dimensional, and  $\text{num}_{n-1}, \dots, \text{num}_1$  are the sizes of the lower  $n - 1$  dimensions, respectively. We give an example in Supplementary material D to explain this rule.

**R3:** R3 is used to infer an  $n$ -dimensional ( $n > 1$ ) static array in an external function. It depends on two requirements  $v1$  and  $v2$ .  $v1$  is defined as  $\neg(\text{exp}(\text{loc}) \diamond (\text{offset}))$ , where “ $\neg$ ” means negation.  $v1$  means that the read location is not computed from the *offset* field. If a CALLDATALOAD reads an item from a static array in an external function,  $v1$  holds because the layout of a static array does not have an *offset* field (Fig. 5).  $v2$  is defined as  $\text{LT}_n(i_n, \text{num}_n) \triangleleft \text{LT}_{n-1}(i_{n-1}, \text{num}_{n-1}) \triangleleft \dots \triangleleft \text{LT}_1(i_1, \text{num}_1) \triangleleft \text{CALLDATALOAD}(\text{loc})$ .  $v2$  is the same as  $v3$  of R2. If a CALLDATALOAD(*loc*) reads an item from an  $n$ -dimensional static array parameter in an external function,  $v2$  holds because there are  $n$  bound checks for each dimension in a nested loop to prevent array overrun before reading the array item. Hence, if  $v1$  and  $v2$  are satisfied, the static array is  $n$ -dimensional, and the item numbers from the highest dimension to the lowest dimension are  $\text{num}_n, \dots, \text{num}_1$ , which are all constants since the item number of each dimension in a static array is known in compilation. Supplementary material D uses an example to explain this rule due to page limit.

**R4:**  $x$  is regarded as a uint256, if R1, R2 and R3 are not fulfilled. R4 means that without sufficient hints we just know that the length of  $x$  is 32 bytes and thus regard a 32-bytes parameter as a uint256. We will refine it to a specific type after using other rules to get more hints.

R19 is used for inferring a struct nested array parameter. R21 and R22 are used for inferring a struct parameter (R21) and a nested array parameter (R22). R24 and R25 are used for inferring a fixed-size list parameter (R24) and a uint256 parameter (R25) after R20 is satisfied.

### 3.3 Rules for CALLDATACOPY

A CALLDATACOPY reads the value of a parameter (§2). R5 – R10 and R23 are used for inferring a one-dimensional dynamic array/bytes/string parameter in a public function (R5), a one-dimensional static array in a public function (R6), a one-dimensional dynamic array in a public function (R7), a bytes/string in a public function (R8), an  $(n + 1)$ -dimensional ( $n > 0$ ) static array in a public function (R9), an  $(n + 1)$ -dimensional ( $n > 0$ ) dynamic array in a public function (R10) and a fixed-size byte array and string in Vyper (R23).

### 3.4 Rules for Other Instructions

R11 – R18 will be applied after R4 is satisfied (Fig. 13). By leveraging them, we can refine a uint256 parameter into a  $\text{uint}(256 - 8 \times x)$ ,  $0 < x < 32$  (R11), a  $\text{bytes}(32 - x)$ ,  $0 < x < 32$  (R12), an  $\text{int}((x + 1) \times 8)$ ,  $0 \leq x < 31$  (R13), a bool (R14), an int256 (R15), an address (R16), a bytes (R17), and a bytes32 (R18). R20 is used to distinguish Vyper bytecode from Solidity bytecode. R26 is used for inferring a fixed-size byte array parameter in Vyper after R23 is satisfied. R27 – R31 will be applied after R25 is satisfied (Fig. 13). By leveraging them, we can refine a uint256 parameter into an address (R27), an int128 (R28), a decimal (R29), a bool (R30), a bytes32 (R31) in Vyper.



Rules for CALLDATALOAD	
$x = \text{CALLDATALOAD}_1(\text{loc}) \quad y = \text{CALLDATALOAD}_2(x+4)$	(R1)
$\frac{\text{Solidity unknown} \triangleright \text{dynamic array/bytes/string} \quad \text{Solidity unknown} \triangleright \text{dynamic array/bytes/string}}{\text{Solidity unknown} \triangleright \text{dynamic array/bytes/string}}$	(R2)
$\frac{\text{Solidity unknown} \triangleright \text{dynamic array/bytes/string} \quad \text{Solidity unknown} \triangleright \text{dynamic array in external mode}}{\text{Solidity unknown} \triangleright \text{dynamic array in external mode}}$	(R3)
$\frac{\text{Solidity unknown} \triangleright \text{static array in external mode}}{\text{Solidity unknown} \triangleright \text{static array in external mode}}$	(R4)
$\frac{\text{Solidity unknown} \triangleright \text{struct/nested array}}{\text{Solidity unknown} \triangleright \text{struct/nested array}}$	(R19)
$\frac{\text{Solidity unknown} \triangleright \text{struct/nested array}}{\text{Solidity unknown} \triangleright \text{struct/nested array}}$	(R21)
$\frac{\text{Solidity unknown} \triangleright \text{struct/nested array}}{\text{Solidity unknown} \triangleright \text{struct/nested array}}$	(R22)
$\frac{\text{Solidity unknown} \triangleright \text{struct/nested array}}{\text{Solidity unknown} \triangleright \text{struct/nested array}}$	(R24)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size list}}{\text{Solidity unknown} \triangleright \text{fixed-size list}}$	(R25)
Rules for CALLDATACOPY	
$\frac{\text{Solidity unknown} \triangleright \text{one-dimensional dynamic array/bytes/string in public mode}}{\text{Solidity unknown} \triangleright \text{one-dimensional dynamic array/bytes/string in public mode}}$	(R5)
$\frac{\text{Solidity unknown} \triangleright \text{one-dimensional static array in the public mode}}{\text{Solidity unknown} \triangleright \text{one-dimensional static array in the public mode}}$	(R6)
$\frac{\text{Solidity unknown} \triangleright \text{one-dimensional dynamic array in public mode}}{\text{Solidity unknown} \triangleright \text{one-dimensional dynamic array in public mode}}$	(R7)
$\frac{\text{Solidity unknown} \triangleright \text{bytes/string in public mode}}{\text{Solidity unknown} \triangleright \text{bytes/string in public mode}}$	(R8)
$\frac{\text{Solidity unknown} \triangleright \text{(n+1)-dimensional static array in public mode}}{\text{Solidity unknown} \triangleright \text{(n+1)-dimensional static array in public mode}}$	(R9)
$\frac{\text{Solidity unknown} \triangleright \text{(n+1)-dimensional dynamic array in public mode}}{\text{Solidity unknown} \triangleright \text{(n+1)-dimensional dynamic array in public mode}}$	(R10)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R23)
Rules for Other Instructions	
$\frac{\text{Solidity unknown} \triangleright \text{uint}(M)}{\text{Solidity unknown} \triangleright \text{uint}(M)}$	(R11)
$\frac{\text{Solidity unknown} \triangleright \text{bytes}(M)}{\text{Solidity unknown} \triangleright \text{bytes}(M)}$	(R12)
$\frac{\text{Solidity unknown} \triangleright \text{bool}}{\text{Solidity unknown} \triangleright \text{bool}}$	(R14)
$\frac{\text{Solidity unknown} \triangleright \text{address}}{\text{Solidity unknown} \triangleright \text{address}}$	(R15)
$\frac{\text{Solidity unknown} \triangleright \text{address}}{\text{Solidity unknown} \triangleright \text{address}}$	(R16)
$\frac{\text{Solidity unknown} \triangleright \text{bytes32}}{\text{Solidity unknown} \triangleright \text{bytes32}}$	(R17)
$\frac{\text{Solidity unknown} \triangleright \text{bytes32}}{\text{Solidity unknown} \triangleright \text{bytes32}}$	(R18)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R20)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R26)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R27)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R28)
$\frac{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}{\text{Solidity unknown} \triangleright \text{fixed-size byte array/string}}$	(R29)
$\frac{\text{Solidity unknown} \triangleright \text{bool}}{\text{Solidity unknown} \triangleright \text{bool}}$	(R30)
$\frac{\text{Solidity unknown} \triangleright \text{bytes32}}{\text{Solidity unknown} \triangleright \text{bytes32}}$	(R31)

Fig. 11. Rules

## 4 DESIGN AND IMPLEMENTATION OF SIGREC

### 4.1 Overview

Fig. 12 shows the architecture of `SigRec` which takes in the *runtime* bytecode of a smart contract and outputs the function signatures of all public/external functions in it. `SigRec` first disassembles the bytecode using Geth disassembler [34] and recognizes basic blocks from them, then extracts function ids from the bytecode. Technical details are presented in Supplementary material E. After that, it uses TASE to infer the types of all parameters (§4.2). Finally, it outputs

the function ids as well as the list of parameter types. The reasons for using TASE rather than conventional SE and other methods are explained in Supplementary material F.

### 4.2 TASE: Type-aware Symbolic Execution

Being the core of `SigRec`, TASE has four steps. First, it conducts coarse-grained type inference to recognize Solidity and Vyper bytecode, and then it recognizes struct, arrays, bytes, strings and basic types of Solidity, and fixed-size lists, fixed-size byte arrays, fixed-size strings and basic types of

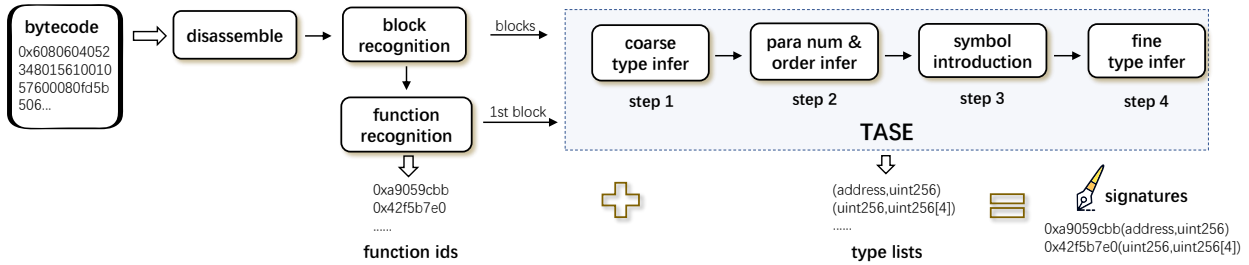


Fig. 12. Architecture of SigRec

Vyper. This step only determines whether a parameter is a basic type instead of deciding the specific basic type, and this step does not infer the type of array items, the type of struct items, and this step does not distinguish a bytes from a string. Second, it infers the number and the order of parameters. Third, by introducing parameter-related symbols when reading arguments from the call data, TASE determines whether an instruction operates on an argument and which argument is operated if that is the case. Finally, it conducts fine-grained type inference to distinguish different basic types for Solidity and Vyper, recognizes the item types of arrays, nested arrays, and each item type of struct for Solidity and the item types of fixed-size list for Vyper, and refines the types (e.g., from uint256 to bytes32) for Solidity and Vyper.

TASE treats the call data as symbols and maintains the symbolic expressions of all variables depending on the call data in order to apply the rules. TASE enhances conventional symbolic execution (CSE) with the ability to infer parameter types. It conducts symbolic execution (SE) to statically explore the paths of the smart contract and stops if the jump target is determined by inputs (e.g., function parameters), because inputs are unknown in static analysis. This restriction is not a big problem in practice since we only find 5 smart contracts deployed in Ethereum containing such kind of jump instructions, and TASE usually does not need to analyze the code deep inside smart contracts because parameters are usually handled near the entry point of each function. TASE treats each value read from the environment as a free symbol, because SigRec focuses on how a smart contract processes the parameters rather than its program logic. Experimental results show that these two restrictions do not affect SigRec’s accuracy (§5.2). After describing the details of TASE, we present the differences between TASE and CSE, as well as the rationale of using TASE instead of other approaches.

### Step 1. Coarse-grained type inference.

At this step, TASE recognizes struct, arrays, bytes, strings and basic types in Solidity, and fixed-size lists, fixed-size byte arrays, fixed-size strings and basic types in Vyper according to the rules R1 – R10 and R19 – R25. This step only determines whether a parameter is a basic type instead of deciding the specific basic type. This step does not infer the type of array items, the type of struct items, and this step does not distinguish a bytes from a string. Instead, these tasks are accomplished by step 4.

Fig. 13 shows the decision tree used by SigRec to determine the type of a parameter. A rectangle represents one or more types and an edge indicates applying one or more rules. The root rectangle represents the unknown

type. Two colors are used to differentiate between a public function and an external function. The  $\rightarrow$  indicates that the requirements of the corresponding rule are not satisfied. SigRec infers the type of a parameter if all the rules on a path from the root rectangle to a leaf node are satisfied.

We use the following example to explain the derivation process. SigRec regards a parameter as a bytes in a public function if R1, R5, R8, and R17 are fulfilled in order. More precisely, if R1 holds, the parameter should be a dynamic array/bytes/string, because the smart contract uses two consecutive CALLDATALOADs to read the *offset* field and the *num* field. Since R5 is also satisfied, the type should be a one-dimensional dynamic array or a bytes or a string in a public function, because exactly one CALLDATACOPY is used to copy the parameter. Then, R8 refines the type to a bytes or a string in a public function, because a bytes/string is extended to the length of a multiple of 32 bytes. Finally, R17 further refines the type to a bytes in the public mode, because the smart contract accesses a single byte of the parameter.

### Step 2. Determining the number and order of parameters.

TASE infers the number of parameters by counting the number of rules R1, R3, R4, R6, R9, R21, R22, R23, R24, R25 which are used in coarse-grained type inference. Specifically, by counting the applied number of R1, TASE knows the number of dynamic arrays/bytes/strings for Solidity (termed by  $n1$ ) since two consecutive CALLDATALOADs are used before reading a dynamic array/bytes/string. By counting the applied number of R3, R6 and R9, TASE knows the number of static arrays for Solidity (termed by  $n2$ ). Since R4 regards all basic types as uint256, TASE knows the number of basic types for Solidity (termed by  $n3$ ) by counting the applied number of R4. By counting the applied number of R21 and R22, TASE knows the number of nested array and struct for Solidity (termed by  $n4$ ). By counting the applied number of R23 and R24, TASE knows the number of fixed-size list, byte array and string for Vyper (termed by  $n5$ ). By counting the applied number of R25, TASE knows the number of basic types for Vyper (termed by  $n6$ ). Therefore, TASE knows the number of parameters, which is  $n1 + n2 + n3 + n4 + n5 + n6$ .

Then, TASE determines the order of parameters according to the locations of the corresponding arguments in the call data in ascending order. For example, if an argument  $x$  locates before an argument  $y$  in the call data,  $x$  is on the left-hand side of  $y$  in the parameter list. The location of a basic type is given by the operand of the CALLDATALOAD. The location of a dynamic array/bytes/string is indicated by the location of its *offset* field, which is the operand of the first CALLDATALOAD instruction in the two consecutive CALLDATALOAD instructions (R1, §3.2). The location of a static array in a public function is given in the operand of the first

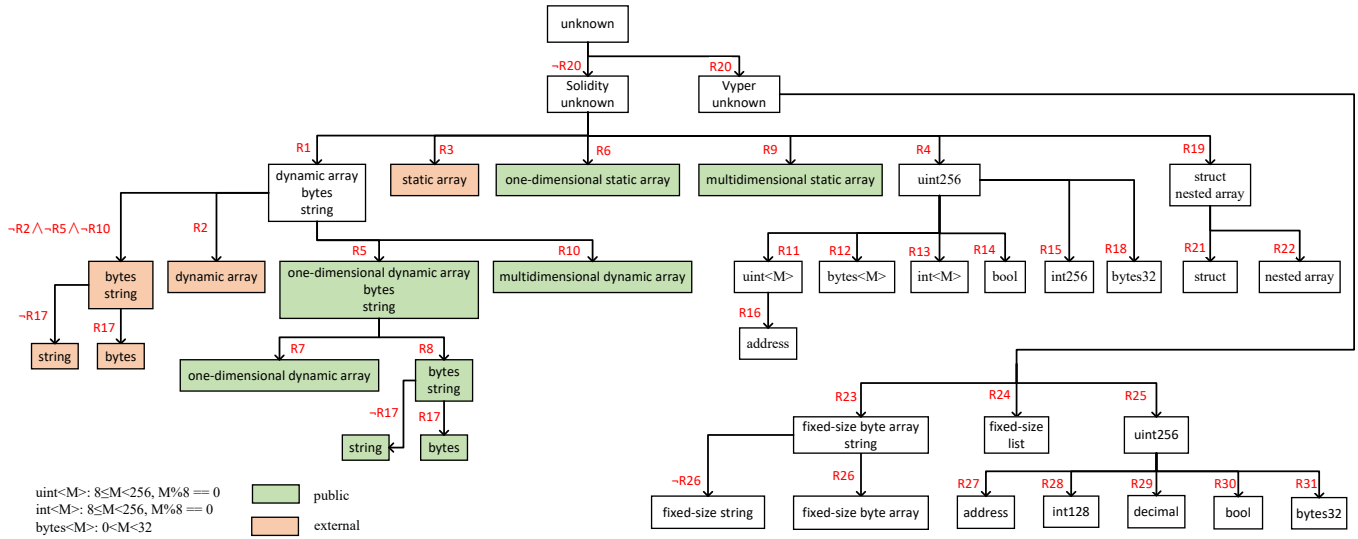


Fig. 13. Hierarchy of rules applied by SigRec

CALLDATACOPY because it copies the inner-most dimension. The start location of a static array in an external function cannot be directly obtained because it is not involved in reading array items. Instead, a CALLDATALOAD is used to read an array item whose location is a constant number determined during compilation. Hence, we get the start location indirectly, which is after the end of the parameter immediately before the static array. If the static array is the first parameter, the start location is  $0x4$  since the function id precedes it.

**Step 3. Introducing parameter-related symbols.** When a smart contract reads arguments, TASE introduces parameter-related symbols by marking all bytes of an argument with the same symbol. After executing a CALLDATALOAD and a CALLDATACOPY, we mark the corresponding stack item and the memory region with symbols. Besides, we will copy the symbol of the memory region to the top item of the stack after executing an MLOAD if it reads from the memory region that stores arguments. Fig. 14 shows an example of marking parameter-related symbols. The public function,  $fun()$ , takes in a  $\text{uint256}$  parameter  $x$  and a  $\text{uint256}[3]$  parameter  $y$ , and assigns a local variable  $z$  as  $y[0]$ . Before the assignment,  $x$  and  $y$  are read from the call data to the stack top and the memory, respectively. Since  $x$  and  $y$  are arguments, we mark the stack top and the corresponding memory region as symbols  $arg1$  and  $arg2$ , respectively. The assignment copies the first item of  $y$  from the memory to the stack top, and therefore we mark stack top with the same symbol of  $y$  (i.e.,  $arg2$ ).




Fig. 14. Marking parameter-related symbols

#### Step 4. Fine-grained type inference.

Using R11–R18 and R26–R31 according to Fig. 13, this step distinguishes basic types for Solidity and Vyper, infers the type of an array item and nested array item for Solidity,

infers each item type of struct for Solidity, and infers the type of a fixed-size list item for Vyper, and refines the type of a parameter to a specific one for Solidity and Vyper. Since TASE knows the structure of an array in step 1, it further knows the type of an array item in this step and then determines the type of the array. For example, given a public function with a  $\text{uint8}[]$  parameter, TASE knows that the parameter is a one-dimensional dynamic array in a public function by applying R1, R5 and R7 in order in step 1. Then, using R11 in this step, TASE learns that the type of an array item is  $\text{uint8}$ , and thus SigRec can recover the correct parameter type.

**An example to illustrate the four steps.** Listing 8 shows a public function with two parameters:  $values$  and  $to$ , whose types are  $\text{uint8}[]$  and  $address$ , respectively. Line 2 reads the first item of  $values$  and  $to$ . Listing 9 shows the corresponding EVM bytecode. For the ease of presentation, we just keep the EVM instructions that are needed by TASE to infer parameter types.

```

1 function test(uint8[] values, address to) public {
2   to.send(values[0]);
3 }

```

Listing 8. An example to explain the process of TASE

```

1 CALLDATALOAD // read offset
2 .....
3 CALLDATALOAD // read num
4 .....
5 CALLDATACOPY // read values
6 .....
7 CALLDATALOAD // read to
8 PUSH20 0xff ... ff // 20byte 0xFF
9 AND //mask
10 .....
11 MLOAD // read values[0]
12 PUSH1 0xff
13 AND //mask

```

Listing 9. The bytecode code of this example

**Step 1: Coarse-grained type inference.** In Listing 9, Line 1 reads 32 bytes, termed by  $x$ , from the 4th byte of the call data, and Line 3 reads 32 bytes, termed by  $y$ , from the  $(x + 4)$ -th byte of the call data. Since these two instructions satisfy R1, this parameter is a dynamic array/bytes/string. Since Line 5 reads  $32 \times y$  bytes from the  $(36 + x)$ -th byte of the call data to memory from the 160th byte without using a loop, R5 is fulfilled and thus this parameter is refined to one-dimensional dynamic array/bytes/string. SigRec further confirms that this parameter is a one-dimensional dynamic

array because R7 is satisfied. Specially, the read length is the multiplication of 32 bytes (i.e., the length of each array item) and the item number (i.e.,  $y$ ). As Line 7 reads 32 bytes, termed by  $z$ , from the 36th byte of the call data to the stack, R4 is fulfilled and thus  $z$  is a basic-type parameter.

*Step 2: Determining the number and order of parameters.* Line 3 reads a one-dimensional dynamic array, and its *offset* field is located at the 4th byte of the call data. Line 7 reads a basic-type parameter located at the 36th byte of the call data. Since no other parameters are read in this example, this function takes in two parameters. Moreover, the one-dimensional dynamic array is the first parameter and the basic-type parameter is the second parameter, because the *offset* field of the array precedes the basic-type parameter.

*Step 3: Introducing parameter-related symbols.* Since Line 5 reads the first parameter to the memory, `SigRec` marks the corresponding memory region with the symbol `arg1`. Similarly, since Line 7 reads the second parameter to the stack, `SigRec` marks the stack’s top item with the symbol `arg2`.

*Step 4: Fine-grained type inference.* Since Line 9 masks the stack top, which is marked with `arg2`, with 20-bytes `0xFF`, R11 is satisfied and thus the second parameter should be a `uint160`. Since the second parameter is not involved in any mathematics operations, R16 is held, and thus `SigRec` refines the type of the second parameter to `address`. Line 11 reads 32 bytes from the 160th byte of the memory to stack top. Since such memory region is marked with `arg1` in step 3, `SigRec` marks stack top with `arg1`. Since the first parameter is an array, the stack top stores an array item. As Line 13 masks the stack top with `0xFF`, R11 is satisfied and thus we learn that the type of the array item is `uint8`. Eventually, TASE infers that the type list is “`uint8[], address`”, the same as the source code.

## 5 EVALUATION

We implement `SigRec` in 8,327 lines of Python code and conduct extensive experiments to evaluate it by answering five research questions. **RQ1:** How is the accuracy of `SigRec` in recovering function signatures (§5.2)? **RQ2:** Will `SigRec` be affected by different compiler versions and optimizations (§5.3)? **RQ3:** How much time is required by `SigRec` to recover function signatures (§5.4)? **RQ4:** How frequently is each rule used for recovering function signatures (§5.5)? **RQ5:** Is `SigRec` superior to existing tools (§5.6)?

### 5.1 Data Collection

To evaluate the accuracy of `SigRec`, we download the source code of all open-source smart contracts which were deployed before Jan. 6, 2021 from Etherscan because the ground-truth (i.e., function signatures) can be obtained from the source code. We collect 119,404 unique open-source smart contracts including 119,126 Solidity contracts and 278 Vyper contracts. 210,869 unique public/external function signatures are found in Solidity open-source smart contracts and 1,076 unique function signatures are found in Vyper open-source smart contracts. To evaluate the efficiency of `SigRec` and the usefulness of heuristic rules, we collect the bytecode of all deployed smart contracts by instrumenting

an Ethereum full node as suggested by [35], [36]. Eventually, we download 11,600,000 blocks (the last block was mined on Jan. 6, 2021) and obtain the bytecode of 37,009,570 smart contracts. 368,679 out of them are unique. There are 47,329,149 public/external functions in all smart contracts with 383,522 unique function signatures. Note that these 37,009,570 smart contracts include all open-source smart contracts.

### 5.2 RQ1: How is the accuracy of `SigRec`?

We evaluate the accuracy of `SigRec` with all 210,869 and 1,076 unique function signatures in Solidity and Vyper open-source smart contracts with ground-truth, respectively. A recovered function signature is *correct*, if and only if the recovered function id, the number and the order of parameters, and the types of all parameters are the same as the ground-truth. The *accuracy* is the proportion of correctly recovered function signatures to the total number of function signatures. `SigRec` correctly recovers 208,218 and 1,052 function signatures for Solidity and Vyper, and hence its accuracy is 98.738%  $((208,218 + 1,052)/(210,869 + 1,076))$ . More specifically, `SigRec` correctly recovers 208,218 from all function signatures in Solidity smart contracts, and thus its accuracy for Solidity contracts is 98.743%  $(208,218/210,869)$ . Besides, `SigRec` correctly recovers 1,052 from all 1,076 function signatures in Vyper smart contracts, and thus its accuracy for Vyper contracts is 97.770%  $(1,052/1,076)$ .

By manually investigate 2,651 incorrect function signatures in Solidity and 24 incorrect function signatures in Vyper, we reveal five cases of inaccuracies as follows. The value in  $\langle \rangle$  is the number of inaccurately recovered function signatures in each case. Please note that one incorrect function signature may be affected by multiple cases. We discuss how to further increase the accuracy of `SigRec` in §7. **Case 1**  $\langle 498 \rangle$ : These smart contracts read the parameters that are not declared in function signatures by inline assembly. Listing 10 shows a practical case. The function `start()` has no parameters in declaration (Line 1), but it reads two parameters `foo` (Line 7) and `bar` (Line 8) using inline assembly. `SigRec` discovers these two parameters read by inline assembly, because it infers parameters by investigating how parameters are used.

```

1 function start() auth note{
2   stopped=false;
3 }
4 modifier note{
5   .....
6   assembly{
7     foo:=calldataload(4)
8     bar:=calldataload(36)
9   }
10  .....}

```

Listing 10. An inaccurately recovered function signature in case 1

**Case 2**  $\langle 387 \rangle$ : These smart contracts forcibly convert the type of a parameter before using it. The recovered types may be more useful than those declared in function signatures, and the smart contract uses the parameter as the converted type. Listing 11 shows a practical case. The function `setGen0Stat()` has one parameter whose type is declared as `uint256[6]` (Line 1). `SigRec` recovers the type of `_gen0Stat` as `uint8[6]`, because the high-order 31 bytes of `uint256` are discarded by type conversion (Lines 4 – 9).

```

1 function setGen0Stat(uint256[6] _gen0Stat) public onlyCOO
2 {

```






Fig. 15. Accuracies of `SigRec` for various Solidity compiler versions




Fig. 16. Accuracies of `SigRec` for Vyper compiler versions

```

3  gen0State=Gen0Stat({
4  retiredAge: uint8(_gen0Stat[0]),
5  maxSpeed: uint8(_gen0Stat[1]),
6  maxStamina: uint8(_gen0Stat[2]),
7  maxStart: uint8(_gen0Stat[3]),
8  maxBurst: uint8(_gen0Stat[4]),
9  maxTemperament: uint8(_gen0Stat[5])});
10 }

```

Listing 11. An inaccurately recovered function signature in case 2

**Case 3** < 314 >: These smart contracts have unused arrays/strings/bytes parameters in external functions. Since `SigRec` recognizes such types of parameters by checking how they are used, it cannot recognize the unused parameters. Missing such parameters may not be an important issue because they are not used. It is worth noting that `SigRec` can recognize all parameters in public functions and the basic-type parameters in the external functions, because such parameters will be read into the stack or the memory no matter whether or not they will be used.

**Case 4** < 602 >: Since `SigRec` recovers the type of each parameter with the `storage` modifier as `uint256`, its output may be incorrect if the storage variable referred by the parameter is not a `uint256` integer.

**Case 5** < 1, 123 >: Our rules cannot handle a few scenarios. First, `SigRec` cannot recognize static arrays in external functions, if the smart contract is compiled with optimization and the index of the array item being accessed is a constant number, because there are no bound checks which are necessary for inferring array structure (R2 and R3, §3.2). Second, `SigRec` cannot distinguish some types if it cannot obtain sufficient clues. For example, we cannot distinguish a bytes from a string if the smart contract does not access a single byte of the bytes value since R17 exploits the fact that a bytes allows accessing its individual byte but a string does not support such functionality. As another example, we cannot distinguish a struct whose items are all basic types from the same items which are not placed in a struct, due to the description about struct in §2.3.1 that the call data layout and the accessing pattern of a static struct are the same to that of the same items which are not placed in a struct.

**Answer:** The accuracy of `SigRec` is 98.7% for Solidity and 97.8% for Vyper.

### 5.3 RQ2: Will `SigRec` be affected by different compiler versions and optimizations?

We evaluate `SigRec` using the open-source smart contracts that are compiled by different versions of compilers w/o optimization. Since Etherscan lists the compiler version and whether or not a smart contract is optimized as well as the optimization level if any for each open-source smart contract [37], we eventually collect 119,126 and 278 unique




Fig. 17. Time consumption to recover function signatures

open-source smart contracts in Solidity and Vyper, respectively. These smart contracts were compiled by 155 versions of Solidity compilers and 17 versions of Vyper compilers, e.g., Solidity from V 0.1.1 to V 0.8.0 and Vyper from V 0.1.0b4 to V 0.2.8. Note that we consider a compiler version with optimization and that without optimization as two different versions. Then, we compute the accuracy of `SigRec` for each compiler version. Fig. 15 shows the accuracies in ascending order of Solidity versions, including the number of smart contracts compiled by each version, ranging from 1 to 11,430. Fig. 16 shows the accuracies in ascending order of Vyper versions, which also includes the number of smart contracts compiled by each version, ranging from 1 to 57. We can see that the accuracy of `SigRec` is never lower than 96% for all 155 Solidity compiler versions, and the accuracy of `SigRec` is more than 90% in 12 out of 15 Vyper compiler versions. By investigating the three Vyper compiler versions that make the accuracy of our tool be lower than 90%, we find that the relatively low accuracy is not caused by compiler features. Instead, the small number of Vyper contracts which are compiled by these three versions is the reason. Therefore, the accuracy does not show a downward trend with the evolving of compiler versions.

**Answer:** Compiler versions and optimizations bring minimum impact on the accuracy of `SigRec`. Its accuracy is never lower than 96% for all 155 Solidity compiler versions and higher than 90% in 80% of Vyper compiler versions.




Fig. 18. Time consumption to recover function signatures in different dimensions






Fig. 19. The number of times each rule is used

#### 5.4 RQ3: How much time is required by SigRec to recover function signatures?

We apply SigRec to all 47,329,149 public/external functions, and measure the time consumption for recovering each function. This experiment is conducted on a desktop equipped with an Intel Xeon E5-2609 CPU, 16GB main memory and 10TB hard disk. The results are shown in Fig. 17, where each cross  $(x, y)$  indicates that SigRec uses no more than  $x$  seconds to recover each of  $y$  function signatures. It shows that the time needed to recover a function signature ranges from  $5 \times 10^{-5}$  seconds to 23.5 seconds, and the average time is 0.074 seconds. For 99.7% (47,177,695/47,329,149) of function signatures, SigRec needs no more than 1 second to recover each of them.

We find three reasons for the function signatures which cost long analysis time. First, the analyzed function has many instructions. Second, the analyzed function signature contains any parameter types which will be confirmed after SigRec executes all the instructions of the function. Taking uint256 as an example, please recall that we initially assume every basic type as a uint256, and then we change uint256 to another type if the parameter is involved in a special instruction. In other words, SigRec recovers the type of parameter as uint256 after it confirms that the parameter is not involved in any special instructions by running all instructions of the analyzed function. On the contrary, SigRec can identify some other parameter types quickly because it needs not to run all instructions of the analyzed function. For example, SigRec identifies an  $\text{int}((x + 1) \times 8)$  when it encounters a SIGEXTEND instruction.

Third, we find that recovering a higher dimensional array costs more time than a lower dimensional array, because the accessing code for a higher dimensional array contains more bound checks and a larger nested loop for reading array items. To quantitatively study the effect of array dimension, we run SigRec to recover an array parameter whose dimension ranges from 1 to 20 and each array item is an uint256. Fig. 18 shows that the time consumption increases linearly along with the increasing of the array dimension. Hence, the array dimension is not a major reason for long analysis time because we find that the array dimension is no larger than 3 in practice.

**Answer:** SigRec is very efficient. For 99.7% of function signatures, it uses no more than 1 second to recover each of them.

#### 5.5 RQ4: How frequently is each rule used?

After recovering 47,329,149 public/external functions in all smart contracts, we count how frequently each rule is used.

TABLE 1  
Results of data set 1

	Gigahorse	Eveem	OSD	EBD	JEB	
1						
2	% same	58.9%	58.1%	18.2%	16.8%	21.3%
3	% fail to process	5.6%	0	0	0	0
4	% no signature	20.8%	20.6%	81.8%	83.2%	78.7%
5	% diff type	12.4%	18.1%	0	0	0
6	% diff num	2.3 %	3.2%	0	0	0

As shown in Fig. 19, all rules have been used. On average, each rule is used for 4,355,236 times. R4 is the most frequently-used rule because the number of basic types in Solidity is more than that of other types. R9 is the least frequently-used rule, because we find that multidimensional static arrays are infrequently used as parameters in public functions.

**Answer:** All rules have been used with different frequencies.

#### 5.6 RQ5: Is SigRec superior to existing tools?

We compare SigRec with five state-of-the-art decompilers, Gigahorse [38], Eveem [15], OSD [16], EBD [17], and JEB [18] in terms of the accuracy of recovering function signatures.

**Datasets.** We prepare three datasets for evaluation. Dataset 1 includes all unique closed-source smart contracts. Dataset 2 contains 1,000 synthesized functions. We construct the name of each function with 5 randomly-selected letters, and annotate it as a public function or an external function randomly. Each function takes in  $x$  parameters ( $1 \leq x \leq 5$ ). For each parameter, we create its name with 5 randomly-selected letters and randomly select a parameter type. Each array parameter has at most three dimensions. Besides, each array parameter has at most five items. The body of each function contains statements to access each parameter, including array items and individual byte of a bytes and a byte32. We construct 100 smart contracts in Solidity, each of which includes 10 synthesized functions, and compile them into bytecode by Solidity 0.5.5 with the probability of 50% to turn on optimization with the default optimization level. Dataset 3 includes all unique open-source contracts obtained from Etherscan.

**Results of dataset 1.** Table 1 lists the results of all closed-source smart contracts. Since there is no ground-truth for closed-source smart contracts, row 2 lists the ratio of function signatures that the five existing tools produce the same results as SigRec. Since tools may abort abnormally when recovering some function signatures, Row 3 presents the

TABLE 2  
Results of data set 2

1		Eveem	OSD	EBD	JEB	SigRec
2	% correct	18.3%	0%	0%	0%	<b>98.8%</b>
3	% fail to process	0	0	0	0	<b>0</b>
4	% no signature	1.1%	100%	100%	100%	<b>0</b>
5	% type error	65.9%	0	0	0	<b>0.5%</b>
6	% parameter num	14.7%	0	0	0	<b>0.7%</b>

TABLE 3  
Results of data set 3

1		Gigahorse	Eveem	OSD	EBD	JEB	SigRec
2	% correct	63.7%	76.2%	44.6%	44.9%	50.1%	<b>98.7%</b>
3	% fail to process	3.4%	0	0	0	0	<b>0</b>
4	% no signature	25.8%	15.7%	55.4%	55.1%	49.9%	<b>0</b>
5	% type error	6.2%	7.3%	0	0	0	<b>0.5%</b>
6	% parameter num	0.9%	0.8%	0	0	0	<b>0.8%</b>

ratio of such function signatures. Row 4 presents the ratio of function signatures that tools just find their function ids but fail to recover their parameter lists. Row 5 presents the ratio of function signatures that the five tools find the same number of parameters as SigRec, but at least one parameter type recovered by them is different from what is recovered by SigRec. Row 6 presents the ratio of function signatures that the five tools find different parameter numbers from SigRec.

We have two observations from the results. First, the accuracies of OSD, EBD and JEB are low because if their databases do not record a function signature, they cannot recover it. Please note that the databases of OSD, EBD and JEB just cover 31.7% of all function signatures in Ethereum public chain (Supplementary material A). Second, Gigahorse and Eveem outperform OSD, EBD and JEB because they try to infer parameter types if they cannot find function signatures from EFSD [14], [15]. Unfortunately, they still fail to correctly recover a large proportion of function signatures. Since Gigahorse and Eveem are not open-source, we investigate their errors by randomly selecting 50 recovered function signatures with different numbers of parameters and 50 recovered function signatures with different parameter types from SigRec for each of them from data set 1. After checking these 200 function signatures (100 ones are not recovered correctly by Gigahorse, and the other 100

TABLE 4  
Results of struct and nested array in Solidity

1		Gigahorse	Eveem	OSD	EBD	JEB	SigRec
2	% correct	10.1%	10.1%	0%	0%	0%	<b>61.3%</b>
3	% fail to process	9.2%	0	0	0	0	<b>0</b>
4	% no signature	72.8%	65.5%	100%	100%	100%	<b>0</b>
5	% type error	12.0%	20.4%	0	0	0	<b>26.6%</b>
6	% parameter num	6.4%	4.0%	0	0	0	<b>12.4%</b>

TABLE 5  
Results of function signatures in Vyper

1		Gigahorse	Eveem	OSD	EBD	JEB	SigRec
2	% correct	68.3%	0%	0%	0%	0%	<b>97.8%</b>
3	% fail to process	3.2%	0	0	0	0	<b>0</b>
4	% no signature	5.9%	100%	100%	100%	100%	<b>0</b>
5	% type error	15.8%	0%	0	0	0	<b>2.2%</b>
6	% parameter num	6.8%	0%	0	0	0	<b>0%</b>

ones are not recovered correctly by Eveem), we observe the following errors: (1) Gigahorse and Eveem report the wrong type. For example, Gigahorse regards the type of a `address/uint8/uint256[]/string/bytes` parameter as `uint256` for some function signatures. Moreover Gigahorse even reports the type of a `uint256` parameter as `uint2304`, which does not exist in smart contracts, in one function signature. (2) Gigahorse mistakenly regards several consecutive parameters as one parameter of a *nonexistent* type. For example, for one function with four parameters whose types are `address[]`, `uint256[5]`, `uint256[5]` and `uint256[2]`, respectively, Gigahorse outputs only one parameter of `uint3328`, a nonexistent type. (3) Gigahorse adds more parameters mistakenly. (4) Gigahorse and Eveem miss some parameters. Supplementary material G lists all errors made by Gigahorse and Eveem from the 200 function signatures.

**Results of dataset 2.** Table 2 lists the results of recovering 1,000 synthesized function signatures by five tools. Row 2 presents the ratio of the function signatures that are correctly recovered. Row 3 and row 4 have the same meaning as the 3rd and 4th rows in Table 1. Row 5 presents the ratio of function signatures that the tools can find the correct parameter numbers but fail to recover the types of at least one parameter. Row 6 presents the ratio of function signatures whose number of parameters cannot be correctly found by the tools. We do not evaluate Gigahorse with the synthesized function signatures because the web service of Gigahorse [38] takes in the addresses of deployed smart contracts rather than the bytecode of smart contracts while deploying 1,000 synthesized functions to Ethereum will cost much money.

The results show that SigRec’s accuracy is 98.8%. Manual investigation reveals that the 8 incorrectly recovered function signatures belong to case 5 (§5.2). OSD, EBD and JEB recover 0 function signatures, because none of the synthesized function signatures are recorded in their databases. Eveem correctly recovers 183 synthesized function signatures thanks to its heuristic rules. However, Eveem cannot produce function signatures for 11 functions, outputs incorrect parameter types for 659 function signatures, and outputs incorrect parameter numbers for 147 function signatures. We manually investigate the errors made by Eveem and the results are similar with the investigation of the errors made by Eveem in dataset 1. We present the detailed results in Supplementary material G.

**Results of dataset 3.** Table 3 lists the results. Row 2 to row 6 have the same meaning with the rows in Table 2. We have several observations. First, SigRec outperforms the other tools by at least 22.5% even the analyzed smart contracts are open-source. Second, Gigahorse is not stable because it aborts abnormally in processing 3.4% of function signatures and it fails to recover some function signatures even they are recorded in EFSD. Third, the accuracies of OSD, EBD and JEB are lower than 51%, even if these functions are implemented in open-source smart contracts. That is, more than 49% function signatures in open-source smart contracts are not recorded in existing function signature databases. Fourth, Eveem outperforms OSD, although they both query EFSD, because Eveem uses its simple rules to infer parameter types if it cannot find function signatures from EFSD. Finally, SigRec outperforms Eveem, because SigRec

has a complete set of rules than Eveem and the powerful TASE engine to infer parameter types. We also investigate the errors made by Gigahorse and Eveem by randomly selecting 200 incorrectly recovered function signatures. The results are similar to their errors in processing dataset 1. Supplementary material G contains the detailed results.

**Recovery of struct and nested array in Solidity.** We compare `SigRec` and existing tools in terms of recovering struct and nested array, which are new parameter types introduced from V 0.4.19 and are supported by the experimental version of Solidity, so-called `ABIEncoderV2` before V 0.8.0 [39]. 1,104 function signatures contain struct or nested array in dataset 3, and results are shown in Table 4. We find that the accuracies of existing tools are no higher than 11%, indicating that existing tools do not support these two new parameter types properly. Gigahorse and Eveem have the same accuracy in recovering struct and nested array because these 10.1% of function signatures are recorded in EFSD. In other words, the rules built in these two tools cannot handle struct and nested array. The accuracy of `SigRec` is 61.3%, which is much higher than existing tools. After investigating the function signatures recovered by `SigRec` incorrectly, we find that all of them belong to case 5 (§5.2). Although the accuracy of `SigRec` in recovering struct and nested array is not as high as the accuracy in recovering other types, we believe it is not a big limitation of `SigRec` because the function signatures taking in struct or nested array just account for 0.5% of the total function signatures. We will derive more rules for recovering struct and nested array as our future work.

**Recovery of function signatures in Vyper contracts.** We compare `SigRec` and existing tools in recovering the function signatures in Vyper contracts, since Vyper is an alternative to Solidity and aims to provide better security than Solidity [40]. 1,076 function signatures appear in all 278 open-source unique Vyper contracts in dataset 3. Results are shown in Table 5. The accuracy of `SigRec` is 97.8%, and we find that all the function signatures which are recovered by `SigRec` incorrectly belong to case 5 (§5.2). Eveem, OSD, EBD and JEB cannot handle Vyper contracts. The accuracy of Gigahorse is 68.3%, and we then investigate the reason for its high accuracy. We find that Gigahorse recovers all Vyper types as `uint256` and fortunately, 68.3% function signatures in Vyper contracts taking in `uint256` parameters only. Hence, existing tools are inadequate in recovering the function signatures in Vyper contracts.

**Importance of function signatures uncovered by existing databases.** 289,123, accounting for 75.4% of all unique function signatures are not included in EFSD, the most popular database for function signatures. To investigate whether these uncovered function signatures are important, we count the number of unique bytecode which contains at least one uncovered function signatures. The number of such unique bytecode is 148,268, which accounts for 40.2% of the total unique bytecode, and thus we believe these uncovered function signatures are important. In other words, EFSD misses important function signatures which can be complemented by `SigRec`.

**Answer:** *SigRec is much more accurate than the state-of-the-art tools in recovering function signatures without the need of function signature databases.*

## 6 APPLICATIONS OF THE RECOVERED FUNCTION SIGNATURES

We use three applications to demonstrate the usefulness of `SigRec`, including detecting short address attacks, fuzzing smart contracts, and reverse engineering the bytecode of smart contracts. Due to page limit, we detail the first application and briefly introduce the results of the last two applications here. Interested readers can refer to Supplementary material H, I for more details.

### 6.1 Detecting Short Address Attacks

We consider the actual arguments sent for function invocation as *invalid* if they are not encoded according to the specification [3]. Although Web3 APIs can generate valid actual arguments, many invalid actual arguments can still be found in practice for various reasons, such as malicious purposes [10], confusion of parameters padding [41], [42], compiler bugs [43], compatibility issues of Web3 [44], [45], changes in the new versions of compiler [46], etc. Invalid actual arguments may incur runtime exceptions, unexpected execution results, and even money stolen [10]. We design `ParChecker`, the first tool for automatically detecting invalid actual arguments sent to Solidity smart contracts by using the recovered function signatures. We plan to extend `ParChecker` to support Vyper smart contracts as our future work. Please note that without function signatures, it is difficult to detect them because different parameter types have different padding schemes (§2). For example, to detect short address attacks, we must focus on the functions with address type parameter(s).

`ParChecker` takes in the call data of a function invocation, and outputs (1) *false* if the actual arguments are invalid, (2) *true* otherwise. `ParChecker` first gets the function id from the call data, and looks for its function signature from the results of `SigRec`. Then, for each parameter in the function signature, `ParChecker` locates the actual argument in the call data and checks whether its padding is correct. Specifically, if the parameter is a basic type, `ParChecker` applies the rules in Table 6 to check the correctness of padding. The rules are derived from the padding scheme of each parameter type described in §2. If the parameter is a static array, `ParChecker` checks each item of the array according to the rules in Table 6. If the parameter is a dynamic struct, nested array, dynamic array, bytes, or string, `ParChecker` checks its structure and content. We use a bytes parameter as an example to illustrate the process. `ParChecker` first locates the *offset* field and the *num* field pointed by the *offset* field in the call data. If either field cannot be found, the structure of the actual argument is invalid. Otherwise, `ParChecker` reads the *num* field whose value is the length of the bytes before padding, termed by  $x$ . After that, `ParChecker` reads  $\lceil x/32 \rceil \times 32$  bytes, termed by  $v$ , after the *num* field, and checks whether the lower-order  $\lceil x/32 \rceil \times 32 - x$  bits of  $v$  are zeros, where “ $\lceil \cdot \rceil$ ” denotes rounding up to the next integer. If not, the content of the actual argument is invalid.

We use `ParChecker` to analyze all transactions in 556,361 blocks starting from the block number 6,625,132. It finds 91,257,261 transactions with non-empty call data and detects 1,024,974 (1% = 1,024,974/91,257,261) transactions with invalid actual arguments. These problematic transactions





cannot recover the type of a parameter with the *storage* modifier, because such parameter is a reference to a storage variable with any possible type and the address of the storage variable rather than its content is passed when passing such parameter [2]. The specialty makes the type of such parameter very challenging to be recognized, and we will extend our work to support such parameter in future work. Second, *SigRec* cannot recognize a function parameter if the clues from a function implementation are insufficient. For example, *SigRec* cannot distinguish a bytes from a string, if the smart contract does not access an individual byte of the bytes parameter. We will address this issue by exploiting the huge number of smart contracts, because one function signature may be found in many smart contracts with various function bodies that may provide sufficient clues. That is, one function signature may associate with many function bodies, and thus we may obtain sufficient clues from these function bodies. Third, the accuracy of *SigRec* to recover some parameter types (e.g., nested array) is not as high as the accuracy to recover other types. We will design more rules to infer these types with higher accuracy in future.

Moreover, if a new parameter type or different accessing pattern is introduced, we will determine the rules for recognizing them through the steps described in §3.1. Please recall that there are five steps to derive rules and the first four have been automated so that it will not cost much time to derive new rules. Similarly, malicious smart contracts may use obfuscation to hide their function signatures from being recognized by *SigRec*. A typical obfuscation technique is replacing the instruction sequence for accessing parameters which can be recognized by *SigRec* with a different instruction sequence with the same semantics which cannot be recognized by *SigRec*. We plan to propose general rules to resist obfuscation, where one rule can represent all possible instruction sequences with the same semantics.

## 8 RELATED WORK

The related studies about recovering function signatures in Smart Contracts have been presented in §1, so this section briefly introduces the related studies about reverse engineering of smart contracts and binaries.

**Reverse Engineering Smart Contracts.** Being a static analysis framework, Vandal decompiles EVM bytecode into IRs [51]. Porosity decompiles EVM bytecode into readable Solidity syntax contracts [52]. Erays transforms EVM bytecode into human readable expressions [50]. EthIR translates EVM bytecode into a rule-based representation for further analysis [53]. Mythril decompiles EVM bytecode into IRs and performs symbolic execution on IRs [54]. GasReducer disassembles the bytecode into assembly code and detects 24 anti-patterns in the bytecode [55]. [56] captures execution traces of smart contracts to evaluate the number of CFTs covered by traces that are not found by the selected tools. GasChecker disassembles the bytecode into assembly code and performs symbolic execution on the assembly code for further analysis [57]. However, none of them recover function signatures. As a necessary step to discover integer overflow bugs, Osiris applies heuristic rules that are similar to R11 and R13 to infer the types of integer variables [58].

However, Osiris can only recognize unsigned integers and signed integers.

**Reverse Engineering Binaries.** Reverse engineering techniques highly depend on the runtime architecture of the target programs. EVM bytecode is syntactically similar to Java bytecode. However, Java bytecode retains function signatures [59], and thus does not need to recover function signatures. Next, we briefly introduce the differences between EVM and x86/x64, which demand a new technique to recover function signatures of smart contracts.

First, EVM has 130+ instructions including many blockchain-specific instructions, which have different semantics with x86/x64 instructions [30]. Second, EVM is a stack-based virtual machine without registers [30]. Third, the following concepts are unique to EVM and x86/x64 does not have such design. (1) The call data stores parameters, and only two EVM instructions *CALLDATALOAD* and *CALLDATACOPY* can read parameters from the call data [30]. (2) EVM maintains a special memory space named *memory* to store some types of parameters, parameters to internal functions and execution results of smart contracts [30]. Moreover, every smart contract has a permanent storage space named *storage* to record <key, value> pairs [30]. (3) An intra-contract function invocation is compiled into a *JUMP* instruction, while an inter-contract function invocation is compiled into a *CALL/CALLCODE/DELEGATECALL/STATICCALL* [30].

Related studies on x86/x64 binaries include recovering parameter types [20], [21], recognizing parameters without recovering parameter types [60], [61], [62], [63], identifying function boundary [64], [65], [66], [67], and inferring variable types [22], [23], [24], [25], [26], [27], [28], [29]. Detailed description is in given Supplementary material K.

## 9 CONCLUSION

We propose and develop *SigRec*, a novel approach to automatically recover function signatures in both Solidity smart contracts and Vyper smart contracts without the need of function signature databases. The extensive experimental results show that *SigRec* has very high accuracy across different compilers and various compiler versions. Experiments also show that *SigRec* is efficient and much more accurate than the state-of-the-art tools. Moreover, we demonstrate that the recovered function signatures are very useful in attack detection, fuzzing and reverse engineering of EVM bytecode.

## ACKNOWLEDGEMENT

This work was partly supported by National Key R&D Program of China (2018YFB0804100), National Natural Science Foundation of China (61872057), HK RGC Project (No. 152193/19E), National Science Foundation under Grant No. 1951729, 1953813, and 1953893.

## REFERENCES

- [1] Ethereum, “Solidity documentation, — introduction to smart contracts,” <https://solidity.readthedocs.io/en/latest/introduction-to-smart-contracts.html>, 2019.
- [2] —, “Solidity documentation,” <https://solidity.readthedocs.io/en/latest/index.html>, 2019.



- [3] —, “Solidity documentation, — contract abi specification,” <http://solidity.readthedocs.io/en/latest/abi-spec.html>, 2019.
- [4] M. di Angelo and G. Salzer, “Wallet contracts on Ethereum,” <https://arxiv.org/pdf/2001.06909.pdf>, 2020.
- [5] M. Fröwis, A. Fuchs, and R. Böhme, “Detecting token systems on ethereum,” in *FC*, 2019.
- [6] M. di Angelo and G. Salzer, “Characterizing types of smart contracts in the Ethereum landscape,” in *WTSC*, 2020.
- [7] T. Chen, Y. Zhang, Z. Li, X. Luo, T. Wang, R. Cao, X. Xiao, and X. Zhang, “Tokenscope: Automatically detecting inconsistent behaviors of cryptocurrency tokens in ethereum,” in *CCS*, 2019.
- [8] B. Jiang, Y. Liu, and W. Chan, “Contractfuzzer: Fuzzing smart contracts for vulnerability detection,” in *ASE*, 2018.
- [9] J. He, M. Balunović, N. Ambroladze, P. Tsankov, and M. Vechev, “Learning to fuzz from symbolic execution with application to smart contracts,” in *CCS*, 2019.
- [10] P. Vessenes, “The erc20 short address attack explained,” <https://vessenes.com/the-erc20-short-address-attack-explained/>, 2017.
- [11] E. Albert, J. Correias, P. Gordillo, G. Román-Díez, and A. Rubio, “Safevm: a safety verifier for ethereum smart contracts,” in *ISSTA*, 2019.
- [12] M. Fröwis and R. Böhme, “In code we trust? measuring the control flow immutability of all smart contracts deployed on ethereum,” in *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, 2017, pp. 357–372.
- [13] 4byte, “Welcome to the ethereum function signature database,” <https://www.4byte.directory/>, 2019.
- [14] N. Grech, L. Brent, B. Scholz, and Y. Smaragdakis, “Gigahorse: thorough, declarative decompilation of smart contracts,” in *ICSE*, 2019.
- [15] Eveem, “Eveem,” <https://eveem.org/>, 2019.
- [16] ethervm, “Online solidity decompiler,” <https://ethervm.io/decompiler>, 2019.
- [17] MrLuit, “Evm bytecode decompiler,” <https://github.com/MrLuit/evm>, 2019.
- [18] JEB, “Ethereum smart contract decompiler,” <https://www.pnfssoftware.com/blog/ethereum-smart-contract-decompiler/>, 2019.
- [19] O. Ganiev, “Ethereum (evm) smart contracts reverse engineering helper utility,” <https://github.com/beched/abi-decompiler>, 2018.
- [20] J. Caballero, N. M. Johnson, S. McCamant, and D. Song, “Binary code extraction and interface identification for security applications,” in *NDSS*, 2009.
- [21] F. de Goër, R. Groz, and L. Mounier, “Lightweight heuristics to retrieve parameter associations from binaries,” in *PPREW*, 2015.
- [22] Z. Lin, X. Zhang, and D. Xu, “Automatic reverse engineering of data structures from binary execution,” in *NDSS*, 2010.
- [23] I. Guilfanov, “Simple type system for program reengineering,” in *WCRE*, 2001.
- [24] J. Lee, T. Avgerinos, and D. Brumley, “Tie: Principled reverse engineering of types in binary programs,” in *NDSS*, 2011.
- [25] K. Troshina, Y. Derevenets, and A. Chernov, “Reconstruction of composite types for decompilation,” in *SCAM*, 2010.
- [26] M. Noonan, A. Loginov, and D. Cok, “Polymorphic type inference for machine code,” in *PLDI*, 2016.
- [27] O. Katz, R. El-Yaniv, and E. Yahav, “Estimating types in binaries using predictive modeling,” *POPL*, 2016.
- [28] Z. Xu, C. Wen, and S. Qin, “Learning types for binaries,” in *ICFEM*, 2017.
- [29] A. Slowinska, T. Stancescu, and H. Bos, “Howard: A dynamic excavator for reverse engineering data structures,” in *NDSS*, 2011.
- [30] G. Wood, “Ethereum: A secure decentralised generalised transaction ledger byzantium version bdec180 – 2019-05-01,” <https://ethereum.github.io/yellowpaper/paper.pdf>, 2019.
- [31] A. Beregszaszi and P. Bylica, “Bitwise shifting instructions in evm,” <https://github.com/ethereum/EIPs/blob/master/EIPs/eip-145.md>, 2017.
- [32] Ethereum, “Solidity documentation, — types,” <https://solidity.readthedocs.io/en/latest/types.html>, 2019.
- [33] —, “Vyper documentation,” <https://vyper.readthedocs.io/en/latest/>, 2021.
- [34] —, “Go ethereum – official golang implementation of the ethereum protocol,” <https://github.com/ethereum/go-ethereum>, 2019.
- [35] T. Chen, Z. Li, Y. Zhu, J. Chen, X. Luo, J. C.-S. Lui, X. Lin, and X. Zhang, “Understanding ethereum via graph analysis,” in *INFOCOM*, 2018.
- [36] T. Chen, Z. Li, Y. Zhang, X. Luo, A. Chen, K. Yang, B. Hu, T. Zhu, S. Deng, T. Hu *et al.*, “Dataether: Data exploration framework for ethereum,” in *ICDCS*, 2019.
- [37] Ethereum, “Etherscan — ethereum (eth) blockchain explorer,” <https://etherscan.io/>, 2019.
- [38] DEDAUB, “Contract library,” <https://contract-library.com/>, 2019.
- [39] Ethereum, “Solidity documentation, — changelog,” <https://github.com/ethereum/solidity/blob/develop/Changelog.md>, 2021.
- [40] M. Kaleem, A. Mavridou, and A. Laszka, “Vyper: A security comparison with solidity based on common vulnerabilities,” in *BRAINS*, 2020.
- [41] Kristaps, “Solidity, problem with bytes decoding,” <https://ethereum.stackexchange.com/questions/69713/solidity-problem-with-bytes-decoding>, 2019.
- [42] E. Luis, “encode data input of the raw transaction to update a contract function?” <https://ethereum.stackexchange.com/questions/25144/encode-data-input-of-the-raw-transaction-to-update-a-contract-function>, 2017.
- [43] MeMyself, “Solidity “call” function with array as input,” <https://ethereum.stackexchange.com/questions/16144/solidity-call-function-with-array-as-input>, 2017.
- [44] filips123, “Compatibility with newer web3 versions,” <https://github.com/trufflesuite/truffle/issues/1690>, 2019.
- [45] taducquang, “Can’t install web3 version 1.2 on windows 10,” <https://github.com/ethereum/web3.js/issues/2971>, 2019.
- [46] Ethereum, “Solidity v0.5.0 breaking changes,” <https://solidity.readthedocs.io/en/v0.5.0/050-breaking-changes.html>, 2018.
- [47] T. D. Nguyen, L. H. Pham, J. Sun, Y. Lin, and Q. T. Minh, “sfuzz: An efficient adaptive fuzzer for solidity smart contracts,” in *ICSE*, 2020.
- [48] V. Wüstholtz and M. Christakis, “Targeted greybox fuzzing with static lookahead analysis,” in *Proc. ICSE*, 2020.
- [49] J. He, M. Balunović, N. Ambroladze, P. Tsankov, and M. Vechev, “Learning to Fuzz from Symbolic Execution with Application to Smart Contracts,” in *Proc. CCS*, 2019.
- [50] Y. Zhou, D. Kumar, S. Bakshi, J. Mason, A. Miller, and M. Bailey, “Erays: Reverse engineering ethereum’s opaque smart contracts,” in *USENIX Security Symposium*, 2018.
- [51] L. Brent, A. Jurisevic, M. Kong, E. Liu, F. Gauthier, V. Gramoli, R. Holz, and B. Scholz, “Vandal: A scalable security analysis framework for smart contracts,” <https://arxiv.org/pdf/1809.03981.pdf>, 2018.
- [52] M. Suiche, “Porosity: A decompiler for blockchain-based smart contracts bytecode,” in *DEFCON 25*, vol. 25, no. 11, 2017.
- [53] E. Albert, P. Gordillo, B. Livshits, A. Rubio, and I. Sergey, “Ethir: A framework for high-level analysis of ethereum bytecode,” in *ATVA*, 2018.
- [54] ConsenSys, “Mythril,” <https://github.com/ConsenSys/mythril>, 2019.
- [55] T. Chen, Z. Li, H. Zhou, J. Chen, X. Luo, X. Li, and X. Zhang, “Towards saving money in using smart contracts,” in *ICSE-NIER*, 2018.
- [56] T. Chen, Z. Li, Y. Zhang, X. Luo, T. Wang, T. Hu, X. Xiao, D. Wang, J. Huang, and X. Zhang, “A large-scale empirical study on control flow identification of smart contracts,” in *ESEM*, 2019.
- [57] T. Chen, Y. Feng, Z. Li, H. Zhou, X. Luo, X. Li, X. Xiao, J. Chen, and X. Zhang, “Gaschecker: Scalable analysis for discovering gas-inefficient smart contracts,” *IEEE Transactions on Emerging Topics in Computing*, 2020.
- [58] C. F. Torres, J. Schütte, and R. State, “Osiris: Hunting for integer bugs in Ethereum smart contracts,” in *ACSAC*, 2018.
- [59] J. Hamilton and S. Danicic, “An evaluation of current java bytecode decompilers,” in *SCAM*, 2009.
- [60] K. ElWazeer, K. Anand, A. Kotha, M. Smithson, and R. Barua, “Scalable variable and data type detection in a binary rewriter,” in *PLDI*, 2013.
- [61] K. Anand, M. Smithson, K. Elwazeer, A. Kotha, J. Gruen, N. Giles, and R. Barua, “A compiler-level intermediate representation based binary analysis and rewriting system,” in *EuroSys*, 2013.
- [62] J. Zhang, R. Zhao, and J. Pang, “Parameter and return-value analysis of binary executables,” in *COMPSAC*, 2007.
- [63] G. Balakrishnan and T. Reps, “Analyzing memory accesses in x86 executables,” in *CC*, 2004.
- [64] R. Qiao and R. Sekar, “Function interface analysis: A principled approach for function recognition in cots binaries,” in *DSN*, 2017.

- [65] T. Bao, J. Burket, M. Woo, R. Turner, and D. Brumley, “Byteweight: Learning to recognize functions in binary code,” in *USENIX Security Symposium*, 2014.
- [66] E. C. R. Shin, D. Song, and R. Moazzezi, “Recognizing functions in binaries with neural networks,” in *USENIX Security Symposium*, 2015.
- [67] S. Wang, P. Wang, and D. Wu, “Semantics-aware machine learning for function recognition in binary code,” in *ICSME*, 2017.
- [68] Ethereum, “Ethereum homestead documentation,” <http://www.ethdocs.org/en/latest/>, 2019.
- [69] T. Chen, Z. Li, Y. Zhang, X. Luo, A. Chen, K. Yang, B. Hu, T. Zhu, S. Deng, and T. Hu, “Dataether: Data exploration framework for ethereum,” in *ICDCS*, 2019.
- [70] T. Chen, X. Li, X. Luo, and X. Zhang, “Under-optimized smart contracts devour your money,” in *SANER*. IEEE, 2017.
- [71] G. Balakrishnan and T. Reps, “Divine: Discovering variables in executables,” in *VMCAI*, 2007.
- [72] B. P. Miller, L. Fredriksen, and B. So, “An empirical study of the reliability of unix utilities,” vol. 33, no. 12, pp. 32–44, 1990.
- [73] Ethereum, “Solidity documentation — contracts,” <https://solidity.readthedocs.io/en/v0.4.24/contracts.html>, 2018.

## SUPPLEMENTARY MATERIAL

### A. Function Signatures Recovered by Different Methods

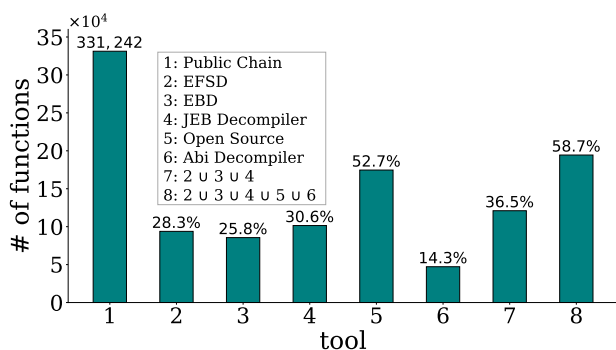


Fig. 21. Function signatures recovered by different methods.

In Fig. 21, the 1st bar gives the number (i.e., 383,522) of all unique function signatures in Ethereum public chain. The 2nd – 6th bars present the numbers of function signatures that can be found from EFSD, EBD, JEB, the open-source smart contracts collected from Etherscan, and Abi Decompiler which is based on brute-force search for function signature recovery, respectively. The 7th bar presents the number of function signatures that can be found in databases, including EFSD, EBD and JEB. The 8th bar gives the number of all function signatures that can be found from EFSD, EBD, JEB, the open-source smart contracts in Etherscan, and Abi Decompiler. The ratio of the function signatures that can be found by each tool is presented above each bar.

EFSD records 24.6% of the function signatures. EBD’s database contains 22.4% of function signatures. The database of JEB covers 26.5% of function signatures. Since function signatures can be easily extracted from open-source smart contracts, 55.0% of function signatures can be obtained through this approach. Abi-decompiler synthesizes function signatures from a list of words and a list of types, and it can recover 12.3% of the function signatures. Moreover, the union of them can recover 60.1% of the function signatures, and hence the accuracy of  $\text{SigRec}$  is 38.6% higher than the accuracy of the combination of all existing tools.

## B. Background

**Account.** There are two kinds of accounts in Ethereum, namely external owned accounts (EOAs) and smart contracts [?], [68]. The major difference between them is that only the latter contains smart contract bytecode. A smart contract can be invoked by an EOA or a smart contract.

**Bytecode of smart contracts.** The bytecode of a smart contract has two parts: initialization bytecode and runtime bytecode [2], [35], [69]. The former is only for contract deployment and thus cannot be invoked by others. Hence, we focus on the public/external functions in the latter. The *gas* mechanism adopted by Ethereum defines the gas cost of each instruction [?], [30], [55], [57], [70].

**Public/external function.** In a smart contract, only the functions declared as *public* or *external* can be invoked by others [7], [68]. The accessing patterns of some parameter types depend on whether the function is public or external (§2 in manuscript [?]). A smart contract can have a *fallback* function, which is unnamed and will be called if the transaction does not specify the invoked function or the specified function does not exist in the invoked smart contract [2].

## C. Rules 5 – 31

**R5:** R5 is used to infer a one-dimensional dynamic array/bytes/string parameter in a public function. R5 will be used after R1 is fulfilled (Fig. 13 in manuscript [?]) and therefore the CALLDATACOPY reads a dynamic array/bytes/string. Let  $\text{CALLDATACOPY}(\text{offset}_m, x + 36, \text{len})$  denote that a CALLDATACOPY reads the data right after the *num* field of a dynamic array/bytes/string, R5 adds one requirement:  $\nexists \text{LT} \triangleleft \text{CALLDATACOPY}(\text{offset}_m, x + 36, \text{len})$ , indicating that the CALLDATACOPY is not in a loop where LTs are loop guards. In this case, CALLDATACOPY reads a one-dimensional dynamic array/bytes/string in a public function.

**R6:** R6 is used to infer a one-dimensional static array in a public function. The  $\text{CALLDATACOPY}(\text{offset}_m, \text{offset}_c, \text{len})$  reads a static array in a public function, because  $\text{offset}_c$  and  $\text{len}$  are constants, indicating that the location and the length of the array are known in compilation. If  $\exists \text{LT} \triangleleft \text{CALLDATACOPY}(\text{offset}_m, \text{offset}_c, \text{len})$  holds, the array is one-dimensional, because the CALLDATACOPY is not inside a loop. Besides, the one-dimensional static array has  $\text{len}/32$  items because every array item is 32 bytes.

**R7:** R7 is used to infer a one-dimensional dynamic array in a public function. R7 will be used after R5 is fulfilled (Fig. 13 in manuscript [?]), and thus the CALLDATACOPY reads a dynamic array/bytes/string in a public function. R7 has a requirement:  $\text{len} == 32 \times y$ , where  $y$  is the value of the *num* field. This requirement indicates that the CALLDATACOPY reads an array since each array item is extended to 32 bytes.

**R8:** R8 is used to infer a bytes/string in a public function. It will be used after R5 is fulfilled (Fig. 13 in manuscript [?]), and thus the CALLDATACOPY reads a dynamic array/bytes/string in a public function. R8 has a requirement:  $\text{len} == 32 \times \lceil y/32 \rceil$ , where  $y$  is the value of the *num* field and “ $\lceil \cdot \rceil$ ” denotes rounding up to the next integer. If it holds, the CALLDATACOPY reads a bytes/string. Such requirement is accordant with the layout of a bytes/string, which is extended so that its length is multiple of 32 bytes.

**R9:** R9 is used to infer an  $(n + 1)$ -dimensional ( $n > 0$ ) static array in a public function. If its requirement  $LT_n(i_n, num_n) \triangleleft LT_{n-1}(i_{n-1}, num_{n-1}) \triangleleft \dots \triangleleft LT_1(i_1, num_1) \triangleleft CALLDATACOPY(offset_m, offset_c, len)$ , where  $num_n, \dots, num_1$  are constant numbers, holds, the  $CALLDATACOPY(offset_m, offset_c, len)$  reads from a multidimensional static array. The requirement indicates an  $n$ -layer nested loop to read an  $(n + 1)$ -dimensional array, and hence  $n$  LT instructions, which are  $n$  loop guards, execute before the  $CALLDATACOPY$  (an example is shown in Listing 1 of manuscript [?]).  $num_n, \dots, num_1$  are constant numbers because for a static array, the size of each dimension is known in compilation. Hence, the item numbers from the highest dimension to the lowest dimension are  $num_n, \dots, num_1, len/32$ , respectively.

**R10:** R10 is used to infer an  $(n + 1)$ -dimensional ( $n > 0$ ) dynamic array in a public function. R10 will be applied after R1 is fulfilled (Fig. 13 in manuscript [?]), and thus the  $CALLDATACOPY$  reads a dynamic array/bytes/string in a public function. If its requirement  $LT_n(i_n, y) \triangleleft LT_{n-1}(i_{n-1}, num_{n-1}) \triangleleft \dots \triangleleft LT_1(i_1, num_1) \triangleleft CALLDATACOPY(offset_m, offset_c, len)$ , where  $y$  is the value of the  $num$  field, holds, the  $CALLDATACOPY$  reads an  $(n + 1)$ -dimensional dynamic array in a public function, because  $n$  loop guards, belonging to a nested loop, execute before the  $CALLDATACOPY$  and the size of the highest dimension is given by the  $num$  field. Hence, only the size of the highest dimension is dynamic and the numbers of items in the lower  $n$  dimensions are  $num_{n-1}, \dots, num_1, len/32$ , which are constant values.

**R11:** For an  $AND(op_1, op_2)$ , if one operand is a uint256 parameter and the other one is a constant number whose higher-order side is  $x$  zero-bytes ( $0 < x < 32$ ), then the uint256 parameter is refined to a  $uint(256 - 8 \times x)$ . The rationale is that a  $uint(M)$ ,  $8 \leq M < 256$ ,  $M\%8 == 0$  will be extended on the higher-order side, and an AND is used for masking the value after extension before accessing its value.

**R12:** For an  $AND(op_1, op_2)$ , if one operand is a uint256 parameter and the other one is a constant number whose lower-order side is  $x$  zero-bytes ( $0 < x < 32$ ), then the uint256 parameter is adjusted to a  $bytes(32 - x)$ . The rationale is that a  $bytes(M)$ ,  $0 < M < 32$  will be extended on the lower-order side, and an AND is used for masking the extended value.

**R13:** For a  $SIGEXTEND(op, x)$ , if  $op$  is a uint256 parameter and  $x$  is a constant number  $0 \leq x < 31$ , then the type of  $op$  is refined to  $int((x + 1) \times 8)$ . The rationale is that  $SIGEXTEND$  is used for sign extension of a signed integer and  $x$  is the length (in bytes) of the signed integer minus 1 [30].

**R14:** If two ISZERO instructions satisfy  $x = ISZERO(op)$  and  $y = ISZERO(x)$  and  $op$  is a uint256 parameter, the type of  $op$  is refined to bool. The rationale is that a bool value is extended to 32 bytes and two consecutive ISZERO instructions are used for masking.

**R15:** If a uint256 parameter is an operand of an SDIV/SMOD/SLT/SGT, its type should be adjusted to int256, because such instructions take in signed integers [30].

**R16:** If a uint160 parameter is not involved in any mathematics instructions, its type is adjusted to address, because an address has 20 bytes and it should not be involved in mathematics computations.

**R17:** For a  $BYTE(op)/MSTORE8(op)$  which reads/writes 1 byte from/to a parameter  $op$  [30], if the type of  $op$  is a string or a bytes, we refine it into a bytes because a bytes supports byte reading and writing whereas a string does not support such operations [32].

**R18:** For a  $BYTE(op)$  which reads a byte from a parameter  $op$ , if the type of  $op$  is uint256, we refine it into bytes32. The reason is that a bytes32 supports byte reading by executing a  $BYTE$  [30], while an AND is used for extracting a byte from a uint256.

**R19:** R19 is used to distinguish the two types, struct and nested array from other Solidity types. R19 depends on one requirement  $v1$ , which is defined as  $offset_1 = CALLDATALOAD_1(loc_1) \triangleleft offset_2 = CALLDATALOAD_2(loc_2) \triangleleft \dots \triangleleft offset_n = CALLDATALOAD_n(loc_n)$ , where  $n > 2$ , and  $(loc_1$  is a constant number)  $\wedge (exp(loc_2) \diamond (offset_1)) \wedge \dots \wedge (exp(loc_n) \diamond (offset_{n-1}))$ .  $v1$  means that the read location of  $offset_1$  is a constant value, and the read location of  $offset_2$  is computed by  $offset_1$ , and so on. When  $x > 2$ , the value of  $offset_x$  is computed by nested  $CALLDATALOAD$ s for  $offset_1$ . For example, the value of  $offset_3$  is  $CALLDATALOAD(CALLDATALOAD(offset_1+))$ . That means, for an item in the parameter, which value is  $CALLDATALOAD(offset_1+)$ , it is the  $offset$  field for a value  $CALLDATALOAD(CALLDATALOAD(offset_1+))$ . So the explanation for  $v1$  is that there is a read location which is computed from the  $offset$  field of the item inside the parameter. So if  $v1$  is satisfied, the parameter contains an item with dynamic size, and please recall that only the nested array and struct can satisfy  $v1$ .

**R20:** R20 is used to distinguish Vyper bytecode from Solidity bytecode. R20 depends on one requirement  $v1$ , which is defined as  $x = CALLDATALOAD(0) \wedge MSTORE(offset_m, x)$ . If  $v1$  is satisfied, this bytecode is a Vyper bytecode. We observe that at the beginning of the Vyper bytecode, Vyper will apply  $CALLDATALOAD$  and  $MSTORE$  to store the first 4 bytes of call data indicating the function id to the memory. Differently, Solidity will store the 4 bytes function id to the stack from the call data by only a  $CALLDATALOAD$ .

**R21:** R21 is used to infer the struct type in Solidity, which is applied after R19 is satisfied. R21 depends on one requirement  $v1$ , which is defined as  $offset_1 = CALLDATALOAD(loc) \triangleleft \# LT(i, CALLDATALOAD(offset_1 + 0x4)) \triangleleft CALLDATALOAD(offset_1 + 0x4 + 0x32 \times i)$ , where  $loc$  is a constant value.  $v1$  means that the read location of  $offset_1$  is a constant value, there is a  $CALLDATALOAD$  that its read location is computed by  $i$  and  $offset_1$ , while the  $offset_1$  is the  $offset$  field of the parameter, and there is no comparison with  $i$  and  $CALLDATALOAD(offset_1 + 0x4)$ . The explanation for  $v1$  is that there is no  $num$  field for this parameter, and  $num$  field is used for bound check for accessing the item in parameter. So if  $v1$  is satisfied, the parameter is a struct type parameter, please recall that there is no  $num$  field for the struct type.

**R22:** R22 is used to infer the nested array type in Solidity, which is applied after R19 is satisfied. R22 depends on only one requirement  $v1$ .  $v1$  is defined as  $LT(i_n, num_n) \triangleleft LT(i_{n-1}, num_{n-1}) \triangleleft \dots \triangleleft LT(i_1, num_1) \triangleleft CALLDATALOAD_{n+1}(loc)$ , while  $(exp(loc) \diamond (offset_n))$  and  $num_x$  can be constant or variable number, and at least one of  $num_1 \dots num_{n-1}$  should be variable.  $v1$  means that there are  $n$  bound checks before read from a position computed from  $offset_n$ . So if  $v1$  is satisfied, this parameter is a  $n$ -dimension nested array, please recall

that for each dimension, there is a bound check. If  $num_x$  is a constant number, the corresponding dimension has  $num_x$  items. On the contrary, if  $num_x$  is a variable, the corresponding dimension is dynamic.

**R23:** R23 is used to infer the fixed-size bytes and string in Vyper. R23 will be applied after R20 is satisfied. R23 depends on one requirement  $v1$ , which is defined as  $x = CALLDAT-ALOAD(loc) \wedge CALLDATACOPY(offset_m, offset_c, len)$ , while  $len$  is a constant number and  $offset_c$  is computed by  $x$ . If  $v1$  is satisfied, this parameter is a fixed-size bytes or string. The explanation for  $v1$  is that Vyper uses a CALLDATACOPY to copy the item number field and the whole value from call data to memory for fixed-size bytes and string type, and the size of fixed-size bytes and string is known before compilation as  $maxLen$ . So the copy length  $len$  is a constant value and computed by  $32 + maxLen$  and the copy position in call data is computed by  $offset$  of this value as  $x$ .

**R24:** R24 is used to infer the fixed-size list type in Vyper. The requirements of R24 are the same as those of R3, but R24 should be applied after R20 is satisfied.

**R25:** R25 is used to infer the uint256 type in Vyper, which is applied after R20 is satisfied.  $x$  is regarded as a uint256 in Vyper, if R23 and R24 are not fulfilled. That means that without sufficient hints we just know that the length of  $x$  is 32 bytes and thus we currently consider a 32-bytes parameter as a uint256. We will refine it to a specific type after applying other rules to get more hints.

**R26:** R26 is used to distinguish fixed-size bytes from fixed-size string. If a parameter type is fixed-size bytes or fixed-size string and the parameter is an operand of a BYTE instruction or a MSTORE8 instruction, we adjust the parameter type to fixed-size bytes because a fixed-size bytes supports to access its individual byte whereas a fixed-size string does not support such operation. R26 is applied after R23 is satisfied.

**R27:** R27 is used to infer the address type in Vyper, which will be applied after R25 is satisfied. For a LT( $op_1, op_2$ ), if  $op_1$  is a uint256 parameter and the value of  $op_2$  is  $2^{160}$ , then the uint256 parameter is refined to an address type parameter, due to the existence of the bound check for the address type.

**R28:** R28 is used to infer the int128 type in Vyper. R28 should be applied after R25 is satisfied. R28 depends on two requirements,  $v1$  and  $v2$ . The parameter type is int128 when both of the two requirements are satisfied.  $v1$  requires a SLT( $op_1, op_2$ ) instruction where  $op_1$  is a uint256 parameter and  $op_2$  is  $2^{127} - 1$ .  $v2$  requires a SGT( $op_1, op_3$ ) instruction where  $op_1$  is a uint256 parameter and  $op_3$  is  $-2^{127}$ . The two requirements indicate the bound checks before accessing an int128 parameter.

**R29:** R29 is used to infer the decimal type in Vyper. R29 should be applied after R25 is satisfied. R29 depends on two requirements,  $v1$  and  $v2$ . The parameter type is decimal when both of the two requirements are satisfied.  $v1$  requires a SLT( $op_1, op_2$ ) instruction where  $op_1$  is a uint256 parameter and  $op_2$  is the 10 decimal value of  $2^{127} - 1$ .  $v2$  requires a SGT( $op_1, op_3$ ) instruction where  $op_1$  is a uint256 parameter and  $op_3$  is the 10 decimal value of  $-2^{127}$ . The two requirements indicate the bound checks before accessing a decimal parameter. It is worth noting that the precision of 10 decimal value of  $2^{127} - 1$  and  $-2^{127}$  is different with it in int128 type.

**R30:** R30 is used to infer the bool type in Vyper, which will be applied after R25 is satisfied. For a LT( $op_1, op_2$ ), if  $op_1$  is a

uint256 parameter and the value of  $op_2$  is 2, then the uint256 type is refined to a bool type, due to the existence of the bound check before accessing the bool parameter.

**R31:** R31 is used to infer the bytes32 type in Vyper. The requirements of R31 is the same as those of R18, but R31 should be applied after R25 is satisfied.

## D. Examples of Using Rules

### 1. An Example of Using R2:

```

1 offset = CALLDATALOAD(0x4)
2 num = CALLDATALOAD(offset + 0x4) //num = 2
3 LT //i2 < 2, or else abort
4 LT //i1 < 3, or else abort
5 loc = offset + 4 + 32 + 32 * 3
6 x = CALLDATALOAD(loc) //x[1][0]

```

Listing 12. The instructions to read an item from a dynamic array

Given an external function with a two-dimensional dynamic array parameter  $x[3][]$ , a transaction for invoking it contains the actual argument  $x[3][2]$ . Listing 12 shows the instructions to read  $x[1][0]$ . For the ease of presentation, we omit irrelevant instructions and merge some instructions into one statement. Line 1 and Line 2 read the values of the  $offset$  field and the  $num$  field, respectively. Since the size of the highest dimension is two, the value of the  $num$  field is two. Line 3 and Line 4 are two bound checks for two dimensions. The location of  $x[1][0]$  is computed in Line 5, and we explain how the location is computed in more detail. Since array items are located right after the  $num$  field, the smart contract first locates the  $num$  field, which is pointed by the  $offset$  field (i.e.,  $offset + 4$ , where 4 is the length of the function id).  $x[0][0]$  is at  $offset + 4 + 32$ , because the length of the  $num$  field is 32 bytes. Since  $x[1][0]$  is the first item of the second  $x[3]$ ,  $x[1][0]$  is located  $32 \times 3$  bytes after  $x[0][0]$ . From the instructions, we can see that the requirements of R2 are held, and therefore  $SigRec$  infers that the parameter is a dynamic array. More specifically, the symbolic expression of  $loc$  is computed by adding the value of the  $offset$  field with a number which contains the multiplication of 32 (Line 5). Besides, the two LTs are passed before reading the  $num$  field (Line 6). From the two LTs, we know that the array has two dimensions. From Line 4, we know that the size of its lowest dimension is three.

### 2. An Example of Using R3:

Given an external function with a two-dimensional static array parameter  $x[3][2]$ , Listing 13 lists the instructions to read  $x[1][0]$ . Lines 1 and 2 are bound checks. Line 3 computes the location to be read. More precisely,  $x[0][0]$  follows the function id and hence it is located at  $0x4$ . Since  $x[1][0]$  is the first item of the second  $x[3]$ ,  $x[1][0]$  is located  $32 \times 3$  bytes after  $x[0][0]$ . Line 4 reads  $x[1][0]$ . From the instructions, we learn that the requirements of R3 are held, and therefore the parameter is a static array. More precisely, the read location is not affected by the  $offset$  field (Line 3). Besides, two LTs are passed before reading the array item (Line 4). Moreover, the two LTs suggest that the array has two dimensions, and the sizes of its highest dimension and lowest dimension are two and three, respectively.

```

1 LT //i2 < 2, or else abort
2 LT //i1 < 3, or else abort
3 loc = 4 + 32 * 3
4 x = CALLDATALOAD(loc) //x[1][0]

```

Listing 13. The instructions to read an item from a static array



## E. Block & Function Recognition

To recognize blocks and public/external functions, `SigRec` partitions the disassembled code into blocks, and then extracts function ids and identifies the first block of every public/external function by parsing the dispatch routine which directs contract execution to the invoked function.

**Block Recognition.** A block starts with the first instruction of the bytecode or the instruction right after the last instruction of a block. Besides, every `JUMPDEST` instruction begins a block since every jump instruction will go to a `JUMPDEST` instruction [30], [56]. The last instruction of a block should be a `JUMP/JUMPI/STOP/RETURN/REVERT/INVALID`. `JUMP` and `JUMPI` are unconditional jump and conditional jump, respectively [30]. `STOP`, `RETURN`, `REVERT` and `INVALID` halt the execution of the smart contract [30].

**Function Recognition.** `SigRec` extracts function ids and identifies the first block of every public/external function by exploiting the fact that the invoked smart contract extracts the function id from the call data and jumps to the first block of the invoked function. Solidity and Vyper generate different bytecode for such function dispatch, so `SigRec` handles the two compilers differently.

*Solidity.* `SigRec` locates a function id and the first block if the first 4 bytes of the call data are compared with a 4-bytes constant number and a `JUMPI` instruction follows the comparison. Note that the 4-bytes constant number is the function id and the jump target of `JUMPI` is the first block. Listing 14 presents an example. Line 1 reads the first 32-bytes from the call data. Line 3 extracts the 4-bytes function id. Line 5 pushes a 4-bytes constant number `0x16d93ade` on the stack. Line 6 compares them. Line 7 pushes a constant number `0x55` on the stack. Line 8 jumps to `0x55` if the first 4-bytes of the call data is equal to `0x16d93ade`. Hence, the function id is `0x16d93ade` and the first block of the corresponding function locates at `0x55` from the beginning of the bytecode.

```

1 CALLDATALOAD //read the 1st 32 bytes from the call data.
2 .....
3 DIV //extract 4-byte function id
4 .....
5 PUSH4 0x16d93ade //push the function id on the stack.
6 EQ //whether they are equal?
7 PUSH1 0x55 //push the location of 1st block on the stack
8 JUMPI //jump to 1st block if function 0x16d93ade is
   called

```

Listing 14. A function with id `0x16d93ade`

*Vyper.* `SigRec` locates a function id and the first block if the first 4 bytes of the call data are compared with a 4-byte constant number and a `JUMPI` instruction follows the comparison. Note that the 4-bytes constant number is the function id and the first block follows the `JUMPI` instruction, which is different from Solidity. Listing 15 presents an example. Line 1 reads the first 32-bytes from the call data to stack. Line 3 store the first 4 bytes of the call data to memory. Line 5 pushes a 4-bytes constant number `0xd178231c` on the stack. Line 7 read the first 4 bytes of the call data from memory to stack. Line 8 compares them. Line 11 jumps to `0x043c` if the first 4-bytes of the call data is unequal to `0xd178231c`. Hence, the function id is `0xd178231c` and the first block of the corresponding function locates follows the `JUMPI` instruction in Line 11.

```

1 CALLDATALOAD//read the 1st 32 bytes from the call data
2 .....
3 MSTORE //store the 1st 4 bytes of the call data to memory
4 .....

```

```

5 PUSH4 0xd178231c //push the function id on the stack
6 PUSH1 0x00
7 MLOAD //read the 1st 4 bytes of the call data from memory
8 EQ //whether they are equal?
9 ISZERO
10 PUSH2 0x043c
11 JUMPI //jump to the next comparison if unequal

```

Listing 15. A function with id `0xd178231c`

## F. Reasons for using TASE instead of CSE and other methods

We propose TASE instead of using other methods (e.g., abstract interpretation [71], value set analysis [60], [61], [65], data dependence analysis [62], taint analysis [22]) to recover function signatures due to three reasons. First, symbolic execution is more suitable than other methods for our task, because it can identify how a variable is computed from the parameters through symbolic expressions and such information is required to apply our rules.

Second, conventional symbolic execution (CSE) is much slower than TASE and suffers from path explosion, and thus is not proper for recovering the ever-growing number of function signatures. More precisely, there are four differences between TASE and CSE. (1) TASE takes into account the rules when symbolically executing type-related instructions (e.g., `CALLDATALOAD`) to infer parameter types. These rules are proposed by this work and specific to EVM bytecode, and thus no CSE tools can achieve the same purpose. (2) TASE does not need to conduct the time-consuming checking of the branch feasibility. Instead, it explores both branches because it aims to recover parameter types and thus does not need to care about the program logic of smart contracts. By contrast, CSE usually spends lots of time in checking branch feasibility. (3) For each public/external function, TASE executes all its blocks only once starting from the first block. That is, TASE does not unfold loops, because executing each block once is enough for inferring parameter types. Therefore, there is no path explosion in TASE. In contrast, CSE usually attempts to explore program paths, thus suffering from severe path explosion. (4) TASE can cover all blocks except dead code since it explores all branches and does not get stuck into loops. In contrast, CSE may leave some reachable blocks unexplored due to the difficulty of checking branch feasibility and path explosion.

Third, we employ TASE instead of taint analysis because of two reasons. (1) some rules need to check whether a parameter is involved in a specific instruction rather than whether the variables are affected by any parameters. For instance, `R11` requires that an operand is a `uint256` parameter, and checks whether the parameter is involved in an `AND` instruction. Symbolic execution knows how an operand is computed from parameters (e.g., directly copy or computed by arithmetic instructions) whereas taint analysis just knows whether an operand is affected by parameters (i.e., tainted or not). (2) some rules need symbolic expressions. For example, `R2` checks whether the symbolic expression of the location contains the summation of the *offset* value and the multiplication of 32. However, taint analysis does not collect symbolic expressions.

## G. Errors Made by Gigahorse and Eveem

### 1. Dataset 1:



Table 7 presents the number of function signatures from dataset 1 (i.e., closed-source smart contracts) whose parameter types are incorrectly recovered by Gigahorse and Eveem. For each such function signature, this table also presents the correct type and the recovered type. If one recovered function signature has several incorrect parameter types, we count it repeatedly.

TABLE 7  
Errors of incorrect parameter types from dataset 1 recovered by Gigahorse and Eveem

Tool	Correct type	Recovered type	# of function signatures
Gigahorse	uint256	address	1
	uint8	uint256	8
	byte32	uint256	4
	byte16	uint256	1
	bool	uint256	5
	uint256	uint32	1
	uint256[]	uint256	10
	address[]	uint256	12
	uint32[]	uint256	1
	uint8[5]	uint256	1
	bool[2]	uint256	1
	uint256[3][]	uint256	1
	uint256	uint2304	1
	string	uint256	6
bytes	uint256	10	
Eveem	address	bool	1
	bool	uint256	4
	address	uint256	3
	byte32	uint256	6
	uint256[]	array	11
	address[]	array	11
	uint32[]	array	2
	uint256[3][]	array	1
	bytes	array	9
	string	array	10

Table 8 lists the number of function signatures from dataset 1 whose number of parameters are incorrectly determined by Gigahorse and Eveem. The correct parameter numbers and types are determined by manually investigating the EVM instructions for accessing parameters. This table also describes why the number of parameters is incorrect. If a recovered function signature misses or adds several parameters, we count it repeatedly.

## 2. Dataset 2:

Table 9 lists the errors made by Eveem in recovering synthesized function signatures. This table also presents the number of function signatures incorrectly recovered by Eveem for each kind of error. We count the number repeatedly, if one incorrectly recovered function signature contains several incorrect recovered parameter types, misses or adds several parameters.

## 3. Dataset 3:

Table 10 presents the number of function signatures from dataset 3 (i.e., open-source smart contracts) whose parameter types are incorrectly recovered by Gigahorse and Eveem. For each such function signature, this table also lists the correct type and the recovered type. If one recovered function signature has several incorrect parameter types, we count it repeatedly.

Table 11 lists the number of function signatures from dataset 3 whose number of parameters are incorrectly determined by Gigahorse and Eveem. This table also describes why the number of parameters is incorrect. If a recovered function signature misses or adds several parameters, we count it repeatedly.

## H. Boosting the Performance of Existing Smart Contracts Fuzzers

All existing fuzzing tools [8], [47], [48], [49] to discover vulnerabilities in smart contracts assume the availability of function signatures. Unfortunately, they are not available for most of the deployed smart contracts. Without function signatures, existing smart contract fuzzers may have to regard the list of parameters as a byte sequence and generate random byte sequences as input instead of applying specific mutation strategies for different parameter types.

To evaluate how function signatures recovered by SigRec can benefit smart contract fuzzers, we develop a tool, named ContractFuzzer<sup>-</sup> which is the same as ContractFuzzer [8], a state-of-the-art open-source smart contract fuzzer, except that ContractFuzzer<sup>-</sup> does not know function signatures. With function signatures, ContractFuzzer knows the type of each parameter and then adopts specific mutation strategy. In contrast, without knowing function signatures, ContractFuzzer<sup>-</sup> regards the list of parameters as a byte sequence and generates random byte sequences like a traditional black-box fuzzer [72]. To make a fair comparison, we ensure that the length of arguments generated by ContractFuzzer is equal to the length of arguments generated by ContractFuzzer<sup>-</sup>. More precisely, ContractFuzzer<sup>-</sup> first applies the same mutation strategy as ContractFuzzer to produce a list of arguments so that the length of arguments is obtained. Then, it replaces the list with a randomly-generated byte sequence of the same length.




Fig. 22. Vulnerabilities found by ContractFuzzer and ContractFuzzer<sup>-</sup>

We randomly select 1,000 smart contracts for this evaluation instead of conducting a larger scale experiment because ContractFuzzer has to maintain an Ethereum private chain and therefore it runs quite slow. We feed ContractFuzzer with the function signatures recovered by SigRec for all the public/external functions of the 1,000 smart contracts. ContractFuzzer discovers 53 vulnerabilities belonging to 5 kinds of vulnerabilities in 40 smart contracts. By contrast, ContractFuzzer<sup>-</sup> only finds 43 vulnerabilities in 32 smart contracts. Manual investigation shows that all vulnerable smart contracts and vulnerabilities found by ContractFuzzer<sup>-</sup> are also discovered by ContractFuzzer. Therefore, with the help of function signatures, ContractFuzzer discovers 25%  $((40 - 32)/32)$  more vulnerable

TABLE 8  
Errors of incorrect parameter numbers from dataset 1 found by Gigahorse and Eveem

Tool	Error Description	# of function signatures
Gigahorse	misses a parameter of address	6
	misses a parameter of bool	3
	misses a parameter of uint32	9
	misses a parameter of uint256[3]	6
	misses a parameter of uint256[6]	2
	misses a parameter of uint256[9]	2
	misses a parameter of uint256[10]	1
	misses a parameter of address[4]	1
	misses a parameter of byte32[4]	1
	misses a parameter of bool[9]	2
	misses a parameter of address[]	3
	misses a parameter of uint256[]	4
	misses a parameter of string	5
	misses a parameter of bytes	6
	considers 4 consecutive parameters of address[], uint256[5], uint256[5], uint256[2] as 1 parameter of nonexistent type uint3328	1
	considers 3 consecutive parameters of uint256[], uint256[5], uint256[5] as a 1 parameter of nonexistent uint2816	1
	considers 3 consecutive parameters of uint256, uint8[3], uint256[3] as a 1 parameter of nonexistent uint1792	1
considers 2 consecutive parameters of uint256, uint256[3][8] as a parameter of nonexistent uint6400	1	
Eveem	misses a parameter of uint256	8
	misses a parameter of address	10
	misses a parameter of uint8	1
	misses a parameter of uint64	1
	misses a parameter of bool	3
	misses a parameter of uint256[16]	6
	misses a parameter of uint256[7]	4
	misses a parameter of uint8[5]	2
	misses a parameter of bool[2]	2
	misses a parameter of address[4]	2
	misses a parameter of uint256[11]	2
	misses a parameter of uint256[3][8]	1
	misses a parameter of address[]	3
	misses a parameter of uint256[]	6
	misses a parameter of string	7
misses a parameter of bytes	4	

smart contracts and 23% ((53 – 43)/43) more vulnerabilities, as shown in Fig. 22. By further investigating why ContractFuzzer<sup>-</sup> misses the 10 vulnerabilities, we find that it is due to the lack of function signatures and detail the observations as follows.

**Invalid parameters.** Some parameter types have structural layouts. For example, a bytes has an *offset* field which must point to the *num* field. Unfortunately, since ContractFuzzer<sup>-</sup> is not aware of such structure, it just randomly generates a byte sequence as arguments, thus making it almost impossible to generate a valid structural input. Note that a smart contract halts immediately when reading an invalid input. Two missed vulnerabilities are due to this reason.

Listing 16 presents one such missed vulnerability. It is a reentrancy bug allowing an attacker to steal money from the smart contract. More precisely, *\_to.call.value()* will call the fallback function of *\_to* since no function is specified. If the fallback function of *\_to* calls the function *checkForward()* (Line 1) again, the money of this vulnerable smart contract will be sent to *\_to* repeatedly by executing Line 6. *checkForward()* is a public function because a function is public by

default before Solidity 0.5.0 [73] though it can be specified as being public, external, internal, or private. ContractFuzzer<sup>-</sup> cannot find the vulnerability since the contract will halt the execution at Line 1 due to reading the invalid input, *\_data*.

```

1 function checkForward(bytes _data) constant returns (bool,
2   bool) {
3   return _forward(allowedForwards[sha3(_data[0], _data
4     [1], _data[2], _data[3])], _data);
5 }
6 function _forward(address _to, bytes _data) internal
7   returns (bool, bool) {
8   ...
9   if (!_to.call.value(msg.value)(_data)) { // vulnerability
10    _returnFee(msg.sender, fee);
11    return (false, _safeFalse());
12  }
13  return (true, _applyRefund(startGas + additionalGas));
14 }

```

Listing 16. A missed vulnerability due to an invalid parameter

**Improper unsigned integer generation.** ContractFuzzer adopts a special mutation strategy for unsigned integers to generate small unsigned integers. In particular, it divides the input space of an unsigned integer into regions, and randomly selects integers from each region. For instance, for a uint256, ContractFuzzer divides the input spaces into  $[0, 2^8 - 1]$ ,  $[0, 2^{16} - 1]$ , ...,  $[0, 2^{256} - 1]$ . However, it is

TABLE 9  
Errors made by Eveem when processing dataset 2

	Correct type	Recovered type	# of function signatures
Incorrect type	int<M>	uint<N>	870
	bytes<M>	uint<N>	799
	uint<M>	uint<N>, M≠N	621
	uint<M>	address	123
	bytes<M>	address	107
	bytes	array	324
	string	array	328
	bytes<M>[X <sub>1</sub> ]...[]	array	4
	uint<M>[X <sub>1</sub> ]...[]	array	11
	int<M>[X <sub>1</sub> ]...[]	array	7
Incorrect number	Error description		# of function signatures
	misses a parameter of bytes<M>		83
	misses a parameter of address		3
	misses a parameter of int<M>		77
	misses a parameter of uint<M>		91
	misses a parameter of bool		7
	misses a parameter of string		76
	misses a parameter of bytes		97
	misses a parameter of bytes<M>[X <sub>1</sub> ]...[X <sub>n</sub> ]		4
	misses a parameter of bytes<M>[X <sub>1</sub> ]...[]		4
	misses a parameter of int<M>[X <sub>1</sub> ]...[X <sub>n</sub> ]		6
	misses a parameter of int<M>[X <sub>1</sub> ]...[]		2
	misses a parameter of uint<M>[X <sub>1</sub> ]...[X <sub>n</sub> ]		3
	misses a parameter of uint<M>[X <sub>1</sub> ]...[]		1

TABLE 10  
Errors of incorrect parameter types from dataset 3 recovered by Gigahorse and Eveem

Tool	Correct type	Recovered type	# of function signatures
Gigahorse	bool	uint256	4
	string	uint256	16
	bytes32	uint256	9
	bytes20	uint256	1
	address	uint256	1
	bytes	uint256	5
	uint8	uint256	10
	uint80	uint256	1
	bool	uint256	1
	string	uint256	1
	uint256[]	uint256	5
	address[]	uint256	7
	Eveem	string	array
bytes		array	4
bytes32		uint256	17
address[]		array	7
uint256[]		array	5
bool		uint256	1
int256[]		array	1
bytes32		uint64	1
uint16	uint256	1	

difficult for ContractFuzzer<sup>-</sup> to generate small unsigned integers since it treats an unsigned integer as a 32-bytes byte sequence and it sets each bit to 0 or 1 randomly. For example, if ContractFuzzer<sup>-</sup> needs to produce an unsigned integer smaller than  $2^8$ , it must set each bit in the 31 higher-order bytes to zeros, and the probability is  $(1/2)^{31 \times 8}$ . Conse-

quently, ContractFuzzer<sup>-</sup> cannot find the vulnerability that can only be triggered by small unsigned integers. Seven missed vulnerabilities are due to this reason.

Listing 17 presents one such vulnerability. *startBlock* is set to the current block number (Line 6), and therefore the comparison at Line 11 depends on the block number. ContractFuzzer detects a vulnerability because the following requirements are satisfied [8]. First, the miner can control the comparison result at Line 11 because it can control the block number [68]. Second, Line 12, which is control dependent on Line 11, modifies the global variable, *endBlock*. Line 12 is control dependent on Line 11 because the execution of Line 11 determines whether Line 12 can be executed. For example, if the condition at Line 11 does not hold, the smart contract will halt immediately, and thus Line 12 will not be executed. Unfortunately, ContractFuzzer<sup>-</sup> fails to discover the vulnerability because it is difficult for it to generate a small unsigned integer *\_block* that can pass the check at Line 10. Note that uint is shorthand of uint256 [32].

```

1 uint public startBlock;
2 uint public endBlock;
3 .....
4 function start(uint _block) external onlyOwner() {
5     require(_block < 54000);
6     startBlock = block.number;
7     endBlock = safeAdd(startBlock, _block);
8 }
9 function adjustDuration(uint _block) external onlyOwner()
10 {
11     require(_block <= 72000);
12     require(_block > safeSub(block.number, startBlock));
13     endBlock = safeAdd(startBlock, _block); //vulnerability
14 }

```

Listing 17. A missed vulnerability due to improper uint generation

**Improper bool value generation.** ContractFuzzer randomly sets a bool value to true or false. Please recall that a bool argument is extended to 32 bytes, and hence ContractFuzzer<sup>-</sup> considers a bool argument as a 32-bytes byte sequence. Consequently, the probability for ContractFuzzer<sup>-</sup> to set a bool value to false is just  $(1/2)^{32 \times 8}$  (i.e., all bits of the byte sequence are set to 0). Therefore, ContractFuzzer<sup>-</sup> cannot find the vulnerability which can only be triggered by a bool parameter set to false. One undiscovered vulnerability is due to this reason as shown in Listing 18.

*now* means the current timestamp, and therefore the comparison at Line 7 depends on the current timestamp. ContractFuzzer detects the vulnerability because the following requirements are satisfied. First, the miner can control the comparison result at Line 7 because it can control the timestamp [68]. Second, the miner can affect the global variable, *rateSale* at Line 8 which is control dependent on the comparison at Line 7. However, ContractFuzzer<sup>-</sup> always sets *isGlobalPause* to true (Line 4) because it always sets the bool parameter *\_state* (Line 3) to true. Therefore, the last two expressions *now > startSale* and *now < finishSale* will not be evaluated (Line 7) since *isGlobalPause* is already true. That is why ContractFuzzer<sup>-</sup> misses the bug.

```

1 bool public isGlobalPause=false;
2 uint public rateSale = 400*10**18;
3 function globalPause(bool _state) public onlyOwner {
4     isGlobalPause = _state;
5 }
6 function changeRateSale(uint _tokenAmount) public
7     onlyOwner {
8     require(isGlobalPause || (now > startSale && now <
9         finishSale));
9     rateSale = _tokenAmount;
10 }

```

Listing 18. A missed vulnerability due to improper bool generation

TABLE 11  
Errors of incorrect parameter numbers from dataset 3 found by Gigahorse and Eveem

Tool	Error Description	# of function signatures
Gigahorse	misses a parameter of bool	4
	misses a parameter of address	8
	misses a parameter of uint256	11
	misses a parameter of uint8	1
	misses a parameter of bytes4	1
	misses a parameter of bytes32	1
	misses a parameter of bytes	4
	misses a parameter of uint16[10]	1
	misses a parameter of uint256[8]	1
	misses a parameter of uint256[3]	1
	misses a parameter of address[4]	1
	misses a parameter of address[3]	4
	misses a parameter of address[2]	1
	misses a parameter of address[16]	2
	misses a parameter of address[]	1
	misses a parameter of bytes32[2]	1
	recognizes a parameter of bytes32[5] as a nonexist type uint35184372087264	1
	recognizes a parameter of bytes as a nonexist type uint17592186044416	1
	adds a parameter of string	1
	adds a parameter of uint32	7
adds a parameter of uint256	7	
adds a parameter of bytes	2	
adds a parameter of address	3	
adds a parameter of a nonexist type bytes1	1	
Eveem	misses a parameter of bool	2
	misses a parameter of bytes6	1
	misses a parameter of string	2
	misses a parameter of bytes	1
	misses a parameter of uint16[10]	1
	misses a parameter of address[4]	1
	misses a parameter of uint256[5]	3
	misses a parameter of uint256[2]	1
	misses a parameter of address[3]	4
	considers a parameter of address[3] as 3 consecutive parameters of address	2
	considers a parameter of uint256[4] as 4 consecutive parameters of uint256	1
	misses a parameter of uint256	10
	misses a parameter of bytes32	5
	misses a parameter of uint8	2
	misses a parameter of uint64	1
	misses a parameter of address	3
	misses a parameter of address[]	9
	misses a parameter of uint256[]	7
misses a parameter of int256[]	1	
misses a parameter of byte32[]	4	

**Summary:** *SigRec empowers existing smart contract fuzzers to analyze all deployed smart contracts by providing function signatures. The function signatures recovered by SigRec help fuzzers adopt proper fuzzing strategies and find much more vulnerabilities.*

### I. Improving the Result of Reverse Engineering the Bytecode of Smart Contracts

We demonstrate the usefulness of SigRec to reverse engineering the bytecode of smart contracts by enhancing Erays [50]. We believe that other reverse engineering tools will also benefit from SigRec. Erays takes in the bytecode, and outputs a register-based instructions which are more readable than EVM bytecode [50]. Unfortunately, Erays recover neither function signatures nor variable types, and its results

contain lots of code produced by the compiler for accessing parameters, making it difficult to understand the program.

We develop a tool named Erays<sup>+</sup>, which first uses Erays to reverse engineering the bytecode and then converts the outputs of Erays into a more readable form. Erays<sup>+</sup> adds 1,456 lines of Python to Erays. The enhancement consists of five parts. First, Erays<sup>+</sup> adds the function signature for each public/external function.

Second, Erays<sup>+</sup> replaces meaningless variable names with meaningful parameter names if these variables are copied from parameters (e.g., Erays<sup>+</sup> replaces a variable  $x$  with  $arg1$  indicating that  $x$  the 1st parameter; Erays<sup>+</sup> also replaces a variable  $y$  with  $num(arg1)$  if  $y$  is copied from the  $num$  field of the 1st parameter)

Third, for a assignment statement, if a parameter is as-

signed to a variable, Erays<sup>+</sup> adds the type of the parameter to the variable. Fourth, if a variable refers to the *num* field of a dynamic array, or a bytes, or a string, Erays<sup>+</sup> replaces its meaningless variable name with a meaningful name, *num(argN)* indicating the *num* field of the *Nst* parameter. Finally, Erays<sup>+</sup> replaces the bulk of code generated by the compiler for accessing parameters with simple assignment statements. To recognize such code, we first synthesize a set of contracts containing synthesized functions which take in various parameter types in the same way described in §5.6 of manuscript [?] and then compile these synthesized contracts into EVM bytecode using all major versions of Solidity. After reverse engineering of these EVM bytecode by Erays, we learn the code patterns. Hence, Erays<sup>+</sup> can recognize the code for accessing parameters by scanning the results of Erays according to the code patterns. Please note that Erays<sup>+</sup> is used for demonstrating the usefulness of SigRec to reverse engineering the bytecode of smart contracts, rather than a full-fledged decompiler.

```

1 contract EraysABI{
2   string s;
3   function f1(uint[2][3][4][1] data, string name) public {
4     uint item = data[3][2][0][1];
5     s = name;
6   }
7 }

```

Listing 19. The source code of an example contract

Listing 19 shows a contract which defines a string *s* (Line 2) and a public function *f1* (Line 3). *f1* takes in a multidimensional dynamic array *data*, and a string *name*. *f1* gets access to one array item (Line 4) and the string *name* (Line 5). We compile this contract using Solidity v 0.4.22 and use Erays to reverse engineer the bytecode. The result, as shown in Listing 20, is difficult to interpret because it consists of 112 lines of instructions and complicated program constructs (e.g., nested *while*, *if...else*, *break*) that cannot be found in the source code. Line 1 guarantees that no Ether can be sent to *f1*. Lines 2 – 58 read the 1st parameter (i.e., the multidimensional dynamic array, *data*) into the memory by a nested loop with 3 layers (Lines 13 – 58). Lines 59 – 65 read the 2nd parameter (i.e., the string, *name*) into the memory by a `CALLDATACOPY`. Line 70 reads the array item from the memory to the stack, which corresponds to Line 4 in the source code. Before accessing the array item, several bound checks (Lines 66 – 69) are performed to prevent array overrun. In particular, Line 66 reads the size of the highest dimension which can only be obtained at runtime, and checks whether the index (i.e., 3) is smaller than the size. Lines 71 to 111 copies the string (i.e., *name*) from the memory to the string *s*. The last line halts the execution. We can see that most code is produced by the compiler for accessing parameters.

```

1 assert(0 == msg.value)
2 $s5 = 0x4 + c[0x4]
3 $t = c[$s5]
4 $s7 = $t
5 $s8 = $m
6 $m = $m + (0x20 + (0x20 * $t))
7 $t = 0x20 + $s5
8 $s5 = $s8
9 $s6 = $t
10 m[$s8] = $s7
11 $s9 = 0x0
12 $s10 = 0x20 + $s8
13 while (0x1) {
14   if ($s9 >= $s7)
15     break
16   $s13 = $m
17   $m = 0x80 + $m
18   $s12 = (0x300 * $s9) + $s6

```

```

19 $s14 = 0x0
20 $s15 = $s13
21 while (0x1) {
22   if ($s14 >= 0x4)
23     break
24   $s18 = $m
25   $m = 0x60 + $m
26   $s17 = (0xc0 * $s14) + $s12
27   $s19 = 0x0
28   $s20 = $s18
29   while (0x1) {
30     if ($s19 >= 0x3)
31       break
32     $s23 = $m
33     $m = 0x40 + $m
34     calldatacopy($s23, (0x40 * $s19) + $s17, 0x40)
35     m[$s20] = $s23
36     $t = $s19
37     $s19 = 0x20 + $s20
38     $s20 = 0x1 + $t
39     $t = $s19
40     $s19 = $s20
41     $s20 = $t
42   }
43   m[$s15] = $s18
44   $t = $s14
45   $s14 = 0x20 + $s15
46   $s15 = 0x1 + $t
47   $t = $s14
48   $s14 = $s15
49   $s15 = $t
50 }
51 m[$s10] = $s13
52 $t = $s9
53 $s9 = 0x20 + $s10
54 $s10 = 0x1 + $t
55 $t = $s9
56 $s9 = $s10
57 $s10 = $t
58 }
59 $s6 = 0x4 + c[0x24]
60 $t = c[$s6]
61 $s9 = $m
62 $m = $m + (0x20 + (0x20 * ((0x1f + $t) / 0x20)))
63 m[$s9] = $t
64 calldatacopy(0x20 + $s9, 0x20 + $s6, $t)
65 $s3 = $s9
66 assert(0x3 < m[$s5])
67 assert(0x1)
68 assert(0x1)
69 assert(0x1)
70 $s4 = m[0x20 + m[m[0x40 + m[0x60 + (0x20 + $s5)]]]]
71 $s8 = m[$s9]
72 $s10 = s[0x0]
73 m[0x0] = 0x0
74 $s9 = sha3(0x0, 0x20)
75 $t = 0x20 + $s3
76 $s7 = $s9 + (((0x1f + (((0x100 * (0 == (0x1 & $s10))) - 0x1) & $s10) / 0x20)
77 $s10 = $t
78 if (0x1f >= $s8){
79   s[0x0] = ($s8 + $s8) | (0
80     xffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffffff00
81     [ $t])
82 } else {
83   s[0x0] = 0x1 + ($s8 + $s8)
84   if ($s8){
85     $t = $s8
86     $s8 = $s10
87     $s10 = $s10 + $t
88     while (0x1) {
89       if ($s10 <= $s8)
90         break
91       s[$s9] = m[$s8]
92       $t = $s8
93       $s8 = $s10
94       $s10 = 0x20 + $t
95       $t = $s8
96       $s8 = $s10
97       $s10 = $t
98       $t = $s9
99       $s9 = $s10
100      $s10 = 0x1 + $t
101      $t = $s9
102      $s9 = $s10
103      $s10 = $t
104    }
105  }
106 }
107 $s10 = $s9
108 while (0x1) {
109   if ($s7 <= $s10)
110     break
111   s[$s10] = 0x0
112   $s10 = 0x1 + $s10
113 }
114 stop()

```

Listing 20. Reverse engineering result outputted by Erays

```

1 0xd780c374(uint256[2][3][4][1] arg1, string arg2) {
2   assert(0 == msg.value)
3   assert(0x3 < num(arg1))
4   assert(0x1)

```



```


5  assert(0x1)
6  assert(0x1)
7  uint256 $s4 = arg1[3][2][0][1]
8  string s = arg2
9  stop()
10 }

```

Listing 21. Reverse engineering result outputted by Erays<sup>+</sup>

Listing 21 shows the result of Erays<sup>+</sup> from the same bytecode. Erays<sup>+</sup> adds the function signature provided by SigRec (Line 1). Since the recovered function signature does not include function name and parameter names, Erays<sup>+</sup> uses the function id and uses *argN* to represent the *N*th parameter. Lines 7 and 8 in Listing 21 correspond to Lines 4 and 5 in the source code. Besides, Erays<sup>+</sup> retains the code for checking Ether transfer (Line 2), array overrun (Lines 3 – 6) and terminating execution (Line 10) which are produced by the compiler. Please note that Erays simplifies the bound checks for the lower dimensions to true (Lines 4 – 6) because the sizes of the lower dimensions are known and thus the outcomes of these bound checks can be determined without execution. We then explain how Erays<sup>+</sup> simplifies the result of Erays. First, Erays<sup>+</sup> replaces a meaningless variable name *m*[\$s5] with *num(arg1)* denoting the size of the highest dimension of *arg1* which is a multidimensional array. Second, Erays<sup>+</sup> replaces 58 lines of code in Listing 20 for reading one item of *arg1* with a single line of code in Listing 21 (Line 7). In the meanwhile, Erays<sup>+</sup> adds a type to *\$s4* because Erays<sup>+</sup> knows that *\$s4* is an array item of the first parameter. Then, Erays<sup>+</sup> replaces 48 lines of code in Listing 20 for reading *arg2* with a single line of code in Listing 21 (Line 8). Erays<sup>+</sup> also adds a type to *s* because it is directly copied from the second parameter. Therefore, the results of Erays<sup>+</sup> are much easier to interpret than the original outputs of Erays.

We use Erays<sup>+</sup> to process 53,166 unique open-source contracts and obtain the results from 39,935 out of them (Erays crashes in processing the other 13,231 contracts). We find that Erays<sup>+</sup> improves the readability of the outputs of all 39,935 contracts. The numbers of smart contracts whose reverse engineering results can be improved by Erays<sup>+</sup> in terms of adding types, adding parameter names, adding *num* names, and removing compiler-generated code are 37,249, 39,740, 17,381 and 17,983, respectively. Fig. 23 depicts the cumulative distribution function (CDF) plots for each case. Each green point (*x*, *y*) in this figure denotes that there are *y* smart contracts whose reverse engineering results are improved by Erays<sup>+</sup> in terms of adding types, for each, the number of added types is no more than *x*. The other colors of points can be interpreted in the same way. We find that the numbers of added types, added parameter names, added *num* names, removed code lines range from 1 to 165, 1 to 272, 1 to 200, and 1 to 303, respectively.

Fig. 23. The numbers of smart contracts whose reverse-engineering results are improved by Erays<sup>+</sup>

**Summary:** The information provided by SigRec including function signatures can obviously improve the readability of the results of reverse engineering tools for smart contracts

## J. Short Address Attacks Detected by ParCheck

Table 12 presents the addresses of the victim smart contracts and the number of attacking transactions.

TABLE 12  
Short address attacks detected by ParChecker

Address	# of transactions
0x42a952ac23d020610355cf425d0dfa58295287be	20
0x183630c3afa08957e588eaa26b748cc5e2d42dc6	11
0xfbe7f429dfa08edc28cb0085c49901d6ffdfaa33	6
0xc9c4d9ec2b44b241361707679d3db0876ac10ca6	6
0x1b6c5864375b34af3ff5bd2e5f40bc425b4a8d79	5
0xda3cd7eed7dc8a0bc76968a9ae67d318d1634b8	3
0x348c796f9902dfdecf572f1ddb262c2bc24ffa5b	2
0x04f33d30c9feebe0c5beb0fe20931958c5fd52b6	2
0x51b5e3291976d4fc7cf7bc76efce6b85aace9f6	2
0xf8520c235406c3d599dff5bb29147ef9ed2e46c2	1
0xd5c8e92d18116abf02e6d07eca55f5c897d49151	1
0x1268e0636680c90b8aae70f90e7f3d94b546ee80	1
0x922f8672d9daf7d5fe784d741eab3f585b033658	1
0xdbb409e7c15f58186775959850df18967f7deade	1
0xec491c1088eae992b7a214efb0a266ad0927a72a	1
0xe5f6035bafd4b1ad3c0010a3ca4086c0891b4a68	1
0x0aef025e7955919a1a2ad98d2e58bc7b3db99436	1
0xe75e6b147fd831a6b79c05dcf2401bc187303b79	1
0x52d19f0ba6f33f4222f55c030085f5cd24bf984d	1
0x3243f38c7f251526b1948b09f0aaa1d2752ad9ae	1
0x056fd409e1d7a124bd7017459dfea2f387b6d5cd	1
0x695dd499ec104dcba3f90d1e77f77aff15e9c666	1
0x21a6b7ed9560604fbc2e459abfee3b433c7d6bf8	1
0xdac17f958d2ee523a2206206994597c13d831ec7	1
0x5bffc45d740c213e19b68b40e9ed89705f495e44	1

## K. Related Work on Reverse Engineering Binaries

Caballero et al. propose to recover function interfaces by dynamic analysis [20]. Their approach applies several architecture-specific rules, e.g., it ignores any access to ESP to exclude unnecessary input locations [20]. Different from our work, their recovered function interfaces provide the information of the parameter location (i.e., on the stack, in a register or in the memory), parameter length, whether the parameter is a pointer or not, etc., rather than the type list declared in the source code. Goër et al. propose a set of architecture-specific heuristic rules (e.g., floats are passed through %xmm0 to %xmm7) to infer parameter number and parameter types [21]. However, the recovered type is either an integer, a float, or an address [21] rather than the accurate type.

SecondWrite proposes a modified value set analysis to recognize parameters [60], [61]. For a call instruction *i*, if a location *l* is defined at somewhere before *i* and is used after *i*, and then Zhang et al.'s work considers the location *l* as a parameter of *i* [62]. Balakrishnan and Reps consider the stack items with positive offsets to the location pointed by ESP as parameters [63]. These studies [61], [62], [63] just recognize parameters, but do not recover parameter types. Qiao and Sekar apply static analysis to discover function boundaries [64]. BYTEWEIGHT learns signatures for function starts using a weighted prefix tree, recognizes function

starts by matching binary fragments with the signatures, and applies value set analysis to recognize function bodies and discover function boundaries [65]. Shin et al. train a recurrent neural network to take bytes of the binary as input, and predict, for each location, whether a function boundary is present at that location [66]. FID applies symbolic execution to generate semantic information, learns the function recognition model, and recognizes function boundaries by machine learning [67]. However, these studies [64], [65], [66], [67] just discover function boundaries, but do not recover function signatures.

REWARDS applies architecture-specific rules which take advantage of system calls, standard library calls and type-revealing instructions (e.g., `MOVS` moves a string) to infer variable types [22]. Guilfanov infers variable types from standard library function prototypes [23]. Besides standard library functions, TIE extracts rules from instructions and applies these rules to infer variable types [24]. Troshina et al. extract architecture-specific rules (e.g., the address computed by  $m + j \times 4$  is used for accessing the array item  $m[j]$ ) from the accessing patterns of structs and arrays, and apply the rules to discover structs and arrays from binaries [25]. Retypd is a powerful type system to recover variable types, and it is more accurate than TIE, REWARDS, and SecondWrite [26]. Considering the code which accesses a variable as a feature to identify the variable type, Katz et al. uses statistical language models to predict variable types [27] and Xu et al. use machine learning to recover variable types [28]. Howard recovers data structures (e.g., array) by applying rules (e.g. if  $A$  is the address of a structure,  $*(A+4)$  presumably accesses a field in this structure) to execution traces [29]. However, these rules for other architectures cannot be used to recover the function signatures of smart contracts due to the significant differences between those architectures and EVM.