arXiv:2211.02652v1 [cs.CR] 2 Nov 2022

AntFuzzer: A Grey-Box Fuzzing Framework for EOSIO Smart Contracts

Jianfei Zhou, Tianxing Jiang, Shuwei Song, Ting Chen*

Abstract—In the past few years, several attacks against the vulnerabilities of EOSIO smart contracts have caused severe financial losses to this prevalent blockchain platform. As a lightweight test-generation approach, grey-box fuzzing can open up the possibility of improving the security of EOSIO smart contracts. However, developing a practical grey-box fuzzer for EOSIO smart contracts from scratch is time-consuming and requires a deep understanding of EOSIO internals. In this work, we proposed AntFuzzer, the first highly extensible grey-box fuzzing framework for EOSIO smart contracts. AntFuzzer implements a novel approach that interfaces AFL to conduct AFL-style grey-box fuzzing on EOSIO smart contracts. Compared to black-box fuzzing tools, AntFuzzer can effectively trigger those hard-to-cover branches. It achieved an improvement in code coverage on 37.5% of smart contracts in our benchmark dataset. AntFuzzer provides unified interfaces for users to easily develop new detection plugins for continually emerging vulnerabilities. We have implemented 6 detection plugins on AntFuzzer to detect major vulnerabilities of EOSIO smart contracts. In our large-scale fuzzing experiments on 4,616 real-world smart contracts, AntFuzzer successfully detected 741 vulnerabilities. The results demonstrate the effectiveness and efficiency of AntFuzzer and our detection plugins.

Index Terms—Security and Privacy Protection, Distributed Systems, Software/Software Engineering

1 Introduction

EOSIOis a next-generation, open-source blockchain platform with industry-leading transaction speed and flexible utility [3]. It adopts DPoS (short for Delegated Proof-of-Stake) as its consensus protocol and supports smart contracts [8]. As of 2019, the total value of on-chain transactions of EOSIO has reached over $ 6 billion [26]. Unfortunately, smart contracts that hold a large number of cryptocurrencies have become profitable targets for attackers [15], [22]. For example, the rollback vulnerability within the gambling smart contracts, including FairDice, Betdice, EOSMax, Tobet, etc, has caused the loss of about $ 1.67 million [7]. And the fake transfer vulnerability and fake transfer notification vulnerability have led to the loss of 380K EOS tokens [11]. Calculating the price of $ 5 for each EOS, the cumulative financial damage from these vulnerabilities is about $ 3.2 million [10].

1.1 Existing approaches

Existing approaches for detecting vulnerability within smart contracts can be roughly divided into formal verification, symbolic execution, and fuzzing. The formal verification approach [28], [29], [30], [31] uses theorem provers or formal methods of mathematics to prove the specific properties in a smart contract [25]. However, most of them suffer from massive manual efforts to define the correct properties of smart contracts, which makes it challenging for using them to perform automated vulnerability detection [24]. The symbolic execution approach could effectively explore smart contracts to examine vulnerable patterns and flaws which would be expected in the run-time [11], [19], [25], [33]. Yet, the symbolic execution approach suffers from path explosion, which adds extra difficulty to applying them to analyze complex smart contracts [22]. The fuzzing approach provides random transactions as inputs for smart contracts. The smart contracts are then monitored for unexpected behaviors or vulnerability patterns. However, most of the existing fuzzing tools [10], [12], [20] use black-box fuzzing that suffers from low code coverage and could miss many vulnerabilities. To address this issue, some grey-box fuzzing techniques have been proposed [14], [21] and proved to be effective in improving code coverage. Yet, to the best of our knowledge, these existing studies focus on Ethereum smart contracts, and there are currently no studies on grey-box fuzzing of EOSIO smart contracts. Moreover, since the two platforms share little similarities in virtual machines, the structure of bytecode, and the types of vulnerabilities [26], applying them to EOSIO smart contracts is a very challenging task.

In addition, existing studies [9], [10], [11] only support detecting 3 types of vulnerabilities of EOSIO smart contracts. However, according to our investigation, there are 6 types of reported vulnerabilities in the wild, which far exceed their detection ability.

1.2 Our approach

In this work, we propose AntFuzzer, a novel grey-box fuzzing framework for EOSIO smart contracts. AntFuzzer distinguishes itself from existing security tools for EOSIO smart contracts through the features listed below:

First, as a lightweight test-generation approach, grey-box fuzzing has been proved to be capable of improving code coverage and fuzzing efficiency. However, it is very challenging and time-consuming to develop a practical grey-box

• *J. Zhou, T. Jiang, S. Song and T. Chen are with University of Electronic Science and Technology of China, Chengdu.*
E-mail: chriszhou.dev@gmail.com

Manuscript received April 19, 2005; revised August 26, 2015.

fuzzer from scratch for this new type of object program, EOSIO smart contracts. To address this challenge, we proposed a novel approach to interface AFL (short for American Fuzzy Lop) as the fuzzing input generation engine and conduct AFL-style grey-box fuzzing on EOSIO smart contracts. The insight behind our approach is that most existing grey-box fuzzers are based on AFL. And more importantly, in case more powerful AFL-based fuzzers are proposed in the future, our approach enables AntFuzzer to integrate them easily. In particular, our approach collects AFL-style code coverage information of smart contracts by instrumenting the EOSIO WebAssembly virtual machine (EOS VM) and feeding them to the AFL fuzzer as the feedback for fuzzing input generation via an interface program written in C.

Second, a practical smart contract fuzzing framework requires high extensibility to deal with the continuously emerging smart contract security vulnerabilities. Therefore, AntFuzzer is designed as a highly extensible fuzzing framework. To this end, it separates the information collection, fuzzing input generation, and vulnerability detection with a modular design, which empowers users to develop detection plugins for continuously emerging vulnerabilities easily. The information collection logic is implemented by instrumenting the blockchain software to obtain essential information for detecting vulnerabilities. The fuzzing input generation, as mentioned above, is also able to integrate more powerful AFL-style fuzzers as the fuzzing input generation engine easily. AntFuzzer provides unified interfaces for developing detection plugins. The detection plugins only need a small amount of code to describe attack scenarios and test oracles. To demonstrate the extendibility of AntFuzzer, we developed 6 plugins to detect major vulnerabilities of EOSIO smart contracts, including Fake EOS Transfer, Fake Transfer Notification, Block Information Dependency, Rollback, No Permission Check, and Receipt Hijacking. The average code amount for each detection plugin is only 98 lines of JAVA. In contrast, the code amount of the framework is more than 9,000 lines of JAVA.

Third, AntFuzzer is efficient. Due to the huge number and rapid growth of smart contracts, it requires numerous fuzzing iterations to test all the smart contracts in the wild. Therefore, AntFuzzer adopts several performance optimization designs. By using memory mapping and shared memory to exchange code coverage information and fuzzing inputs, AntFuzzer effectively reduces the communication overhead between its modules, which achieves a good balance between code coverage and fuzzing overhead. Moreover, unlike many other grey-box fuzzers that instrument smart contracts, AntFuzzer instruments the EOS VM for obtaining code coverage, which not only eliminates the overhead for instrumenting smart contracts one by one but also improves execution efficiency. Our large-scale fuzzing experiments show that the average fuzzing iteration per second is 15, which is fast enough for volume fuzzing experiments and only 16% less than EOSFuzzer [10] (i.e. 18 fuzzing iterations per second in our experiment). Due to the lack of a widely recognized benchmark dataset of EOSIO smart contracts, we collected 90 smart contracts with source code as our benchmark dataset. The experiment results on our benchmark dataset show that the vulnerability detection accuracy and F1-score of AntFuzzer reach 98.9% and 97.4%. Using AntFuzzer to perform large-scale fuzzing on 4,616 contracts in the wild, AntFuzzer detected 741 vulnerabilities.

We also conducted code coverage comparison experiments. The results show that AntFuzzer can effectively trigger those hard-to-cover branches. Compared to the black-box fuzzing tool EOSFuzzer [10], AntFuzzer improves code coverage on 37.5% of smart contracts within the benchmark dataset, and the code coverage increased by 22.7% on average.

1.3 Contributions

This work has three main contributions.

- We designed and implemented AntFuzzer, a novel grey-box fuzzing framework for EOSIO smart contracts. Compared with existing vulnerability detection tools, AntFuzzer has certain advantages in terms of supported vulnerability types, extensibility, code coverage, and detection accuracy.
- Based on AntFuzzer, we have implemented 6 plugins to detect 6 types of major vulnerabilities of EOSIO smart contracts, including Fake EOS Transfer, Fake Transfer Notification, Block Information Dependency, Rollback, Receipt Hijacking and No Permission Check. AntFuzzer and our plugins will be open-sourced once this paper is accepted.
- By conducting extensive fuzzing experiments on 4,616 smart contracts in the wild to evaluate AntFuzzer, we observe that AntFuzzer along with the detection plugins detected 741 vulnerabilities with high efficiency.

2 Background

We first provide brief background knowledge and key concepts of the EOSIO platform, EOSIO smart contracts, and AFL in this section.

WebAssembly. WASM (short for WebAssembly) is a portable, fast, and compatible binary compilation target for most Web browsers [13], [16]. The EOSIO platform adopts WebAssembly virtual machine for executing its smart contracts [11].

EOSIO. EOSIO is an open-source public blockchain platform [10]. It adopts DPoS (short for Delegated Proof-of-Stake) as its consensus protocol, which is capable of achieving millions of TPS [26] (short for Transaction Per Second). The EOSIO platform currently supports writing smart contracts in C/C++ and compiling them to WASM bytecode [10], [26].

System Smart Contracts. EOSIO implements core features inside system smart contracts instead of the blockchain platform to improve its scalability [10]. For example, The native currency of EOSIO, EOS, is issued and managed by the eosio.token smart contract [9].

Computing Resources. EOSIO platform consumes computing resources to send and execute transactions [4], [6]. The required resources are charged to the account that signed this transaction with its private key [4]. To attract users, many smart contracts employ a resource payment model called "Receiver Pay". That is, the smart contract will use its own account to sign the transaction and pay for

the computing resources [17]. This resource payment model has the potential to be exploited by attackers to drain the resources of smart contracts. We will detail this attack in section 4.9.

Action & Transaction. Actions define atomic behaviors within an EOSIO smart contract [4]. Actions are always contained within transactions [4]. A transaction can contain one or more actions [4]. Typically, an EOSIO smart contract has a function named *apply()* to listen for incoming actions, and dispatch the actions to the corresponding handler function within the smart contract [11]. There are 3 parameters of the *apply()* function, named *receiver*, *code*, and *action* respectively [9]. The *receiver* represents the current receiver of the action (i.e. the account whose smart contract is currently executing). The parameter *code* is the original receiver of the action [10]. And the *action* carries the name of the incoming action [9], [10]. It is worth noting that there exists a difference between *receiver* and *code*. A smart contract can forward the action that it received to another account using the EOSIO built-in function *require_recipient()* [9], [10]. Before forwarding, the parameter *receiver* will be modified to the newly notified account, while the *code* stays the same [10].

Inline Action & Deferred Action. In the EOSIO communication protocol, actions can be classified into inline actions and deferred actions [4]. A group of inline actions will be packaged into the same transaction, and the inline actions within this transaction must all succeed, one by one, in a predefined order, or else all of the inline actions within this transaction will fail [4]. While a group of deferred actions will be packaged into several transactions respectively [4]. Therefore, if one of these deferred actions failed, the execution of other deferred actions is not affected [4]. Some attackers may use the characteristics of inline actions to carry out rollback attacks, which will be detailed in section 4.6.

Account & Permission. An account identifies a participant in the EOSIO blockchain [4]. Accounts also represent the smart contract actors that push and receive actions to and from other accounts in the blockchain [4]. Permissions associated with an account are used to authorize actions and transactions to other accounts [4]. EOSIO provides smart contract developers with built-in functions such as *require_auth()* and *has_auth()* to check that the actors specified in each action have the minimum permission required to execute it [4].

AFL. AFL (short for American Fuzzy Lop) is an instrumentation-guided grey-box fuzzer [2], [34]. It uses a modified form of edge coverage to effortlessly pick up subtle, local-scale changes to program control flow [2]. It loads user-supplied test cases and mutates the test cases repeatedly to explore new state transitions of the program being fuzzed [2].

3 Design and Implementation of AntFuzzer

3.1 Overview of AntFuzzer

AntFuzzer separates the information collection, fuzzing input generation, and vulnerability detection with modular design for the ease of developing detection plugins. In particular, As shown in Fig. 1, AntFuzzer is composed of three modules, namely *Fuzzing controller*, *AFL agent*, and *Data collector*.

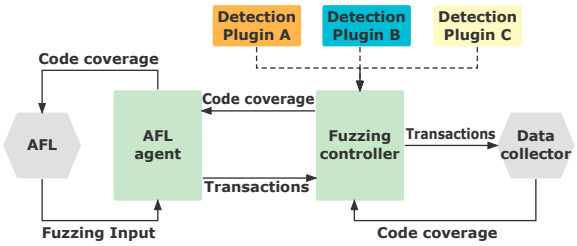

Fig. 1: Architecture of AntFuzzer

There are mainly 4 steps within the workflow of AntFuzzer. First, the *Fuzzing controller* scans and analyzes on the ABI (short for Application Binary Interface) and the bytecode of smart contracts to extract the essential corpus as the seed of fuzzing inputs and initializes the fuzzing environment (e.g. deploying contracts and creating necessary accounts, etc). Second, the *Fuzzing controller* loads and registers vulnerability detection plugins. Each detection plugin defines attack scenarios for triggering a specific type of vulnerability and test oracles for detecting it. Third, according to the attack scenarios, the *Fuzzing controller* starts an AFL agent which interfaces an AFL instance and converts fuzzing inputs generated by AFL to transactions. At last, the transactions are executed in the *Data collector*, an instrumented EOS VM that collects the execution traces and code coverage information of smart contracts. The *Fuzzing controller* used detection plugins to perform vulnerability detection based on the test oracles and recorded execution traces. Meanwhile, the code coverage information is fed to AFL fuzzer as the feedback for mutating fuzzing inputs at the end of each fuzzing iteration.

3.2 Technical Challenges

The key idea of this work, in short, is to interface AFL, the widely recognized and carefully researched grey-box fuzzer, as the fuzzing input generation engine and conduct AFL-style grey-box fuzzing on EOSIO smart contracts. However, as an instrumentation-guided genetic fuzzer for C/C++ programs, there are several major challenges to adopting AFL for generating fuzzing inputs for EOSIO smart contracts.

- **C1:** AFL is designed for fuzzing C/C++ programs, it neither can instrument WASM bytecodes, nor obtain the code coverage of EOSIO smart contracts directly to select fuzzing strategies.
- **C2:** AFL generates unstructured byte sequences as fuzzing inputs, while smart contract functions only take in structured transactions.
- **C3:** AFL can only detect straightforward misbehaviors such as crashes or failed assertions. However, the security vulnerabilities in smart contracts are closely related to their business logic, and can not be judged by these crashes or failures simply.

3.3 Fuzzing input generation.

Collecting code coverage information. To address **C1**, we implemented an AFL-complaint code coverage calculation

```
1 cur_location = <COMPILE_TIME_RANDOM>
2 shared_mem[cur_location ^ prev_location]++
3 prev_location = cur_location >> 1
```

Fig. 2: Code coverage calculation of AFL

```
1 cur_location = abs(rand(pc) % shared_mem.length)
2 shared_mem[cur_location ^ prev_location]++
3 prev_location = cur_location >> 1
```

Fig. 3: Code coverage calculation of AntFuzzer

algorithm in the *Data Collector*. Before we present our code coverage calculation algorithm, we first briefly introduce the algorithm of AFL, which is shown in Fig. 2. AFL instruments C/C++ programs to initialize compile-time random variables to identify code blocks. For instance, the variable *cur_location* indicates the currently executing code block (Line 1), while the *prev_location* indicates the previously executed code block (Line 2). AFL utilizes a slice of shared memory named *shared_mem* to preserve code coverage information. To be more specific, the value at the position $cur_location \oplus prev_location$ in *shared_mem* is incremented by 1, indicating that an execution path from *prev_location* to *cur_location* is hit once (Line 2). Then, *cur_location* is shifted right by one bit and assigned to *prev_location* to differentiate the direction of the execution paths.

Our code coverage calculation algorithm does not need to instrument WASM bytecode to identify code blocks, instead, it instruments the EOS VM to achieve the same purpose. As shown in Fig. 3, we use the *pc* (short for program counter) of the EOS VM to generate *cur_location* at run-time. To be specific, we instrumented the interpretation handlers of all control-flow-related instructions of the EOS VM, including *call*, *br*, and *br_if*, etc. Each time the EOS VM executes these control-flow-related instructions, we use the value of *pc* as the random seed to generate *cur_location* to identify the currently executing code block (Line 1). Then, same as AFL, we count the hit path, preserve code coverage information in a slice of shared memory, and shift right *cur_location* by one 1 bit as well (Line 2-3). The shared memory will be written in a bitmap memory mapping file and sent back to the AFL fuzzer. It is worth noting that, by instrumenting the EOS VM instead of WASM bytecode, AntFuzzer not only eliminates the costs of instrumenting smart contracts one by one but also improves the fuzzing efficiency since the instrumented code in the EOS VM is compiled into native code, which runs much faster than WASM bytecode that executed in the virtual machine.

Interfacing with AFL. In the *AFL agent*, we implemented a C program based on Kelinci [34] to interface with AFL. In fact, AFL fuzzer does not know that it is fuzzing an EOSIO smart contract but thinks that it is fuzzing this C program.

The AFL has two requirements for instrumented applications [34]. First, the instrumented applications are required to connect to shared memory and write to the locations corresponding to branches of the application [34]. Second, the applications should run a fork server which is responsible for forking new processes on an input file provided by AFL [34]. The C program implements a fork server that is the same as the one in the program instrumented by the AFL compiler [34]. When the AFL fuzzer generates a fuzzing input and requests this interface program to fork a process, it hands over the fuzzing input to the *AFL agent*. The fuzzing input is then converted to a transaction and executed in the *Data collector*. After the execution, the C interface program receives the collected code coverage bitmap memory mapping file from the *Data collector* and writes them to shared memory. Therefore, from the perspective of AFL, the execution paths of this C program are exactly the same as the EOSIO smart contract. Note that the code coverage information is fed to AFL via memory mapping file and shared memory, AntFuzzer effectively reduces the communication overhead to archive an ideal fuzzing efficiency.

Generating fuzzing inputs with AFL. To address **C2**, in the first step of AntFuzzer's workflow, the *Fuzzing controller* scans and analyzes ABI of smart contracts to extract the parameter list of each ABI interface as well as the data types of each parameter. With the information of smart contracts' parameter lists, the *AFL agent* converts the byte sequence generated by AFL to transaction parameters for smart contracts by applying the Algorithm shown in Alg.1:

Algorithm 1 Conversion from bytes to parameters

Input: Byte sequence: $bytes$
Input: ABI: abi
Output: Parameters: $parameters$

1: $sorted_abi \leftarrow$ sort parameters in abi
2: $variable_length_parameters \leftarrow \varnothing$
3: **for** $p \leftarrow sorted_abi$ **do**
4: **if** length of type(p) is variable **then**
5: append p to $variable_length_parameters$
6: $n \leftarrow$ length of type(p)
7: $bits \leftarrow$ intercept n bits from $bytes$
8: $value \leftarrow$ convert $bits$ to type(p)
9: append $value$ to $parameters$
10: allocate and convert remaining bits in $bytes$ to $variable_length_parameters$
11: append $variable_length_parameters$ to $parameters$
12: **return** $parameters$

First, *AFL agent* will sort the parameters in the function ABI by the type's bit length (Line 1). Variable-length parameter types (e.g. string) will be sorted at the end of the parameter list. For each fixed-length parameter within a parameter list, the *AFL agent* intercepts the corresponding number of bits from the byte sequence according to the length of its type (Line 6-9). For parameters of variable length, after intercepting the bytes of fixed-length parameters, the *AFL agent* allocates the remaining bytes equally based on the number of variable-length parameters (Line 10-11).

As shown in Fig. 4, take the hexadecimal byte sequence 0x1623416e7446757a7a6572 as an example. The parameter list of a function in this example is **(string, uint8, uint8)**. After sorting the abi, the sorted parameter list is **(uint8, uint8, string)**. Then, the *AFL agent* intercepts the first 8bits (i.e. the first byte, 0x16) from the byte sequence and converts it into an uint8 type parameter (decimal representation is 22). Then, the *AFL agent* generates the second uint8 type parameter in the same way. At last, the remaining bytes,

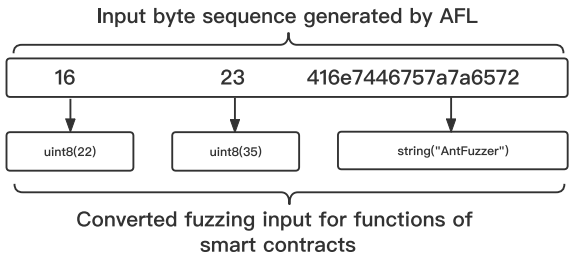

Fig. 4: The conversion from byte sequences to structured parameters

0x416e7446757a7a6572, are converted to a string type parameter (i.e. "AntFuzzer").

To ensure that all parameters can be allocated enough bits, we made a minimal modification to AFL (i.e. only 3 lines of C code), making it support fuzzing input generation with a specified minimum length.

Seed selection. AFL, as a mutation-based grey-box fuzzer, relies on a set of seed inputs (a corpus) to bootstrap the fuzzing process [35]. To select good seeds, AntFuzzer adopts a 2-phase method. First, the *Fuzzing controller* scans and analyzes the bytecode of smart contracts for constructing a corpus that consists of literal values (e.g. constant values and strings) that appear in the bytecode. Second, according to the collected corpus and ABI, AntFuzzer generates several transactions by randomly selecting values from the corpus and executing them for obtaining the code coverage information. Then, the transaction with the highest code coverage will be selected and converted to a byte sequence by using Alg. 1 reversely as the seed input of AFL.

3.4 Information collection.

The *Data collector* is responsible for executing fuzzing input transactions, collecting execution traces and code coverage information of smart contracts. The details of collecting code coverage information have been presented in section 3.3. In this section, we present the details of collecting execution traces, including executed WASM instructions as well as their results and invocation of EOSIO built-in functions. The recorded execution traces will be stored in separate trace files, which will be analyzed by detection plugins.

WASM Instructions. To support extensible analysis of various vulnerabilities, we added instrumentation code to the interpretation handlers of all the control-flow-related instructions and numerical instructions. The instrumentation code records the opcode, operands, and results of these instructions. For example, *call*, *call_indirect*, *br*, and *br_if* are control-flow-related instructions that jump to the address of a specific code block. Once these instructions are executed, the source address and destination address are recorded.

EOSIO built-in functions. EOSIO platform provides many built-in functions for smart contracts, which can be divided into 4 categories, including blockchain states query, smart contract interaction, on-chain data manipulation and permission check. These built-in functions are implemented as the host environment of the EOS VM and imported to the bytecode of smart contracts. We added instrumentation code to the built-in function handlers for recording the invocation parameters and the effects of these functions. The instrumented built-in functions are partly listed below:

1) **Blockchain states query.** The functions, *tapos_block_num()* and *tapos_block_prefix()* are used to query block information. We instrumented these functions to obtain blockchain states and detect blockchain state dependency vulnerabilities.
2) **smart contract interaction.** The functions, *send_inline()* and *send_deferred()* are used to send an inline action or a deferred action respectively. The sender and receiver in the parameters of the two functions are recorded to trace the interactions between smart contracts being fuzzed.
3) **On-chain data manipulation.** The functions, *db_store(), db_update(), and db_delete()* are responsible for manipulating the on-chain data that stored in the smart contract. AntFuzzer records the manipulated data and the resource payer in the parameters of these functions.
4) **Permission check.** The functions, *has_auth()* and *require_auth()* are responsible for checking whether actors specified in actions have the correct permission. We instrumented these functions for detecting permission-related vulnerabilities.

In addition to the above 2 types of execution traces, we instrumented the *transfer()* function of *eosio.token*, to record the EOS token transfer actions and monitor the balances of accounts.

3.5 Vulnerability Detection

To address **C3**, AntFuzzer only adopts AFL for input generation and utilizes the detection plugins to accomplish complex vulnerability detection tasks. Developing a detection plugin on AntFuzzer is easy because it has already provided the service for fuzzing input generation and data collection in its modular, *AFL agent* and *Data collector*. The detection plugins only need to implement attack scenarios for triggering specific vulnerabilities and test oracles for detecting them, without modifying the EOS VM or knowing any details of fuzzing input generation. In particular, AntFuzzer defines three high-level interfaces, namely *beforeFuzz()*, *fuzz()*, and *afterFuzz()*. Each detection plugin implements the three interfaces.

The interface *beforeFuzz()* is responsible for setting necessary blockchain state configurations that are required for triggering specific vulnerabilities. Meanwhile, some detection plugins will deploy attack agents which are smart contracts simulating attackers' behaviors to construct attack scenarios. For example, to construct the attack scenario for triggering the Fake Transfer Notification vulnerability, we need to create two accounts and deploy an attack agent on one of them to send fake transfer notifications. Therefore, the *beforeFuzz()* function of the Fake Transfer Notification vulnerability detection plugin will inform the *Fuzzing controller* to create 2 accounts, deploy an attack agent and issue some EOS tokens to them before the fuzzing starts.

The interface *fuzz()* is responsible for describing how to construct an attack scenario for the smart contract being fuzzed. Technically, an attack scenario specifies how to trigger the vulnerability, that is, generating fuzzing input

```
1 void apply(uint64_t receiver, uint64_t code,
      uint64_t action) {
2   if (code == receiver || action == name("
        transfer")) {
3     execute_action(action);
4   }
5 }
```

Fig. 5: An apply function with a Fake EOS Transfer vulnerability

for the functions of the smart contract or the deployed attack agent. According to the attack scenario, the *Fuzzing controller* requests fuzzing inputs from the *AFL agent*, and sends them to the *Data collector* to execute and collect transaction execution traces.

The interface *afterFuzz()* defines the test oracles for a specific vulnerability. After a fuzzing iteration is completed, the *Fuzzing controller* calls *afterFuzz()* of the detection plugin, to analyze the execution traces collected by the *Data collector* and check whether the execution traces conform to the test oracles.

4 Detection plugins based on AntFuzzer

To demonstrate how to easily and quickly develop detection plugins on AntFuzzer, we developed 6 vulnerability detection plugins.

4.1 P1. Fake EOS Transfer

Vulnerability Description. Under normal circumstances, the *transfer()* function of an EOSIO smart contract will only be invoked when it receives EOS tokens transferred via *eosio.token* system contract [10]. The *apply()* function should check whether the parameter *code* is *eosio.token* when the input parameter *action* is *transfer* [10]. However, an EOSIO smart contract with a Fake EOS Transfer vulnerability may fail to do this check. As a consequence, an attacker can directly call the *transfer()* function of the vulnerable smart contract to mislead it that the attacker has transferred EOS tokens to it [10], [11].

Fig. 5 shows a simplified *apply()* function with a Fake EOS Transfer vulnerability. It only checks whether the parameter *code* is *eosio.token* or itself, but fails to check whether the action *transfer* is originally sent to *eosio.token* [10] (Line 2, i.e. check if the parameter *code* is *eosio.token* [10]).

Attack Scenario. To trigger the Fake EOS Transfer vulnerability, P1 deploys an attack agent contract named *fakeagent* [10]. As shown in Fig. 6, P1 uses the attack agent contract to send an inline action, trying to directly invoke the *transfer()* function within the smart contract being fuzzed [10] (Line 4 - Line 9), which behaves the same with an attacker.

Test Oracles. P1 needs the execution traces of WASM instructions. The test oracles for detecting the Fake EOS Transfer vulnerability under the attack scenario are as follows:

(i) The smart contract being fuzzed can receive the EOS token [10].

```
1  class fakeagent: public eosio::contract {
2  void attack(name target, asset quantity, string
       memo) {
3    eosio::action(
4      eosio::permission_level(_self, name("active"
           ));
5      name(target);
6      name("transfer");
7      std::make_tuple(self, target, quantity, memo
           );
8    ).send();
9  }
10 };
```

Fig. 6: The attack agent smart contract of P1

(ii) After the execution of the function, *attack()*, of *fakeagent* contract, the *transfer()* function of the smart contract is invoked [10].

P1 uses test oracle (i) to check if the smart contract can receive the EOS token [10]. Technically, P1 will send some EOS tokens to it. After that, P1 scans the execution traces of the smart contract being fuzzed and checks whether *call_indirect* instruction was executed twice [10] (i.e. the first one corresponds to the *transfer()* function of *eosio.token* and the second one corresponds to EOS transfer handler function of the smart contract being fuzzed [10]). Then, P1 uses test oracle (ii) to check if the *fakeagent* exploited the Fake EOS Transfer vulnerability successfully according to whether the execution trace of *transfer()* appears after the *attack()* function. Similarly, P1 scans the execution traces to check whether the *call_indirect* instruction appears twice [10] (i.e. the first for the *fakeagent*'s *attack()* function, the second for the *transfer()* function of the smart contract being fuzzed [10]). Note that 2 times appearances of *call_indirect* do not mean that the smart contract is vulnerable. In some smart contracts, the EOS transfer handler function may not be the *transfer()* function. Therefore, P1 will further check whether the destination address of the second *call_indirect* instruction is the EOS transfer handler.

4.2 P2. Fake Transfer Receipt

Vulnerability Description. In the EOS token transfer process, *eosio.token* system contract will forward the transfer action as a receipt to the receiver and sender through the *require_recipient()* function. However, as shown in Fig. 7, after receiving a transfer receipt, the smart contract with a Fake Transfer Receipt vulnerability may fail to check the real receiver of this transfer action (Line 4, i.e. check whether *data.to* equals to _self). As a result, an attacker could create two accounts, A1 and A2, and deploys an attacking contract within A2 as an attack agent. Then the attacker could use A1 to transfer EOS tokens to A2 through *eosio.token* system smart contract. When the transfer completes, A1 and A2 will receive the transfer receipt. And the attack agent deployed within A2 will invoke *require_recipient()* to forward the transfer receipt to the victim contract, making it mistakenly believe that it has received EOS tokens that transferred from A1 [10].

Attack Scenario. To trigger a Fake Transfer Receipt vulnerability, P2 uses 2 accounts, *sender* and *notifier*, and de-

```
void transfer(name sender, name receiver, asset
    amount) {
  auto data = unpack_action_data<st_transfer>();
  // no check for data.to
  if (data.from == _self) {
    return;
  }
  // do something
}
```

Fig. 7: A transfer function with a Fake Transfer Receipt vulnerability

```
1 class fakenotifier: public eosio::contract {
2 [[eosio::on_notify(eosio.token::transfer)]]
3 void transfer(name from, name to, asset quantity
    , string memo) {
4   require_recipient("victim");
5 }
6 };
```

Fig. 8: The notifier contract of P2

ploys an attack agent contract named *fakenotifier* within the account *notifier*. As shown in Fig. 8. The *notifier* listens to the notification sent by the *eosio.token* system smart contract's *transfer()* function (Line 2), and forward it to the victim smart contract. In the *fuzz()* function of P2, it uses the account *sender* to send EOS tokens to *notifier* via *eosio.token* system contract. Then *fakenotifier* forwards the receipt to the smart contract being fuzzed, which acts the same as an attacker.

Test Oracle. P2 needs the execution traces of WASM instructions. The test oracles for detecting the Fake Transfer Receipt vulnerability are as follows:

(i) The *transfer()* function of the smart contract being fuzzed was invoked [10].
(ii) The smart contract did not check the receiver of the notification [10].

P2 utilizes test oracle (i) to check whether the *transfer* function of the contract being fuzzed was invoked successfully. The judgment conditions of test oracle (i) are similar to the test oracle (i) of P1. P2 checks whether there are 3 execution traces of *call_indirect* instruction. The first one is for the *transfer()* function of *eosio.token* system contract [10]. And the second one is for the *transfer()* function of *notifier* [10]. The third one is for the *transfer()* function of the smart contract being fuzzed [10]. Since some smart contracts may check the receiver of transferred tokens within the *transfer()*, and exit when it checks out that it is not the receiver of the transfer action. P2 uses test oracle (ii) to detect whether the smart contract being fuzzed has executed a check statement. In particular, P2 scans the execution traces of the contract and checks whether there are comparison instructions, *eq* and *ne*, that compare *notifier* and the smart contract itself [10] (i.e. *data.to* and *_self* shown in Fig. 7).

4.3 P3. Block Information Dependency

Vulnerability Description. Most gambling or gaming smart contracts require random numbers to determine prize numbers or game-winners. However, due to the lack of randomness source on the blockchain, some smart contract developers mistakenly use EOSIO built-in functions to generate random numbers (e.g. *tapos_block_num()*, *tapos_blocks_prefix()* and *block_time_stamp()*) [10]. Since attackers can easily obtain the return value of these built-in functions, using them as the source of random numbers will make the prize numbers predictable.

Attack Scenario. Unlike P1 and P2, no attack agent contract is required to trigger the Block Information Dependency vulnerability. Since some smart contracts require users to participate in the game by directly transferring EOS tokens to it or require users to deposit some tokens before they invoke specific functions (e.g. *play()*, *bet()*, etc) to join the game. P3 randomly selects the functions of the smart contract to invoke them or sends EOS tokens to it.

Test Oracle. P3 needs the execution traces of EOSIO built-in functions. The test oracles for detecting the Block Information Dependency vulnerability are as follows:

(i) The smart contract being fuzzed invoked EOSIO built-in functions related to blockchain states query [10].
(ii) After the invocation of block-information-related functions, the *transfer()* function of the smart contract and *eosio.token* system contract was invoked [10].

P3 uses test oracle (i), scanning the execution traces of the smart contract being fuzzed to check whether the functions for blockchain states query were invoked (e.g. *tapos_block_num()*, *tapos_block_num()*, and *block_time_stamp()*). Then the test oracle (ii) is used to check whether the Block Information Dependency vulnerability was successfully exploited and the smart contract being fuzzed really transferred EOS tokens to our attack agent.

4.4 P4. No Permission Check

Vulnerability Description. Before a smart contract executes sensitive operations such as transferring EOS, sending inline actions or deferred actions to invoke another smart contract, and modifying users' data stored within the smart contract, the smart contract should check whether the actors specified in the action have the minimum permission required to execute it.

Attack Scenario. P4 selects functions within the smart contract being fuzzed that have parameters of data type *name* or *public_key*, and randomly invokes these functions.

Test Oracle. P4 needs the execution traces of EOSIO built-in functions. The test oracles for detecting the No Permission Check vulnerability are as follows:

(i) The smart contract being fuzzed performed sensitive operations.
(ii) No permission check functions were invoked before the sensitive operations.

Test oracle (i) is to check whether there are execution traces of functions that are relevant to sensitive operations (e.g. *send_inline()* is relevant to sending EOS and invoking another contract, *db_store()* is related to modifying users' data stored within the smart contract, etc). Then P4 checks whether there are execution traces of permission check functions such as *require_auth()* or *has_auth()* before each sensitive operation. If not, P4 reports the vulnerability.

```
1  class rbatk: public eosio::contract {
2  void makebet(name target, asset quantity, string
       memo) {
3    asset originBalance = getBalance(_self, "EOS")
       ;
4    eosio::action(
5      eosio::permission_level(_self, name("active"
         ));
6      name("eosdice");
7      name("transfer");
8      std::make_tuple(self, target, quantity, memo
         );
9    ).send();
10
11   currentBalance = getBalance(_self, "EOS");
12   assert(currentBalacne > originBalance, "
       rollback");
13 }
14 };
```

Fig. 9: The attack agent smart contract to exploit Rollback vulnerability

```
1  class biggame : public eosio::contract {
2  public:
3  void notify(account user, string reply) {
4    transaction tx;
5    const uint128_t sender_id = 0;
6    tx.actions.emplace_back(
7      action(permission_level(_self, name("active"
         )),
8        _self,
9        name("decereceipt"),
10       std::make_tuple(user, "received your bet!"
           ))
11   );
12   tx.deley_time = 0;
13   tx.send();
14 }
15
16 void decereceipt(account user, string reply) {
17     // send receipt to user
18     require_recipient(user);
19 }
20 };
```

Fig. 10: A smart contract with a Receipt Hijacking vulnerability

4.5 P5. Rollback

Vulnerability Description. Due to the characteristics of inline actions, if an error occurs during the execution of an inline action, all of the inline actions that are packed in the same transaction will fail. Some gambling or game smart contracts use an inline action to handle the bet sent from players and reveal the winner. However, such a design will lead to the following consequences, that is, an attacker can deliberately make his unwinnable bet action fail and roll back. As a result, if the attacker does not win, he can get his EOS tokens back.

The attacker first deploys an attack agent smart contract shown in Fig. 9. For simplicity of description, the code is simplified. In line 4, the attacker checks his balance before he sends EOS tokens to the victim smart contract *EOSDice*. Then, *EOSDice* reveals the winner in the *reveal()* function and sends EOS tokens as rewards to the winner. In line 12, the attacker checks his balance again and compares *currentBalance* and *originBalance* in order to determine whether he wins and receives rewards. If not, the attacker throws an assertion failure (Line 13). As a result, the action that sends EOS tokens to the victim smart contract will also fail and roll back, that is, the attacker gets his EOS tokens back.

Attack Scenario. In order to trigger the Rollback vulnerability within the smart contract being fuzzed, P5 deploys an attack agent smart contract that is similar to the contract shown in Fig. 9 and invokes the attack agent to perform a rollback attack on the smart contract being fuzzed.

Test Oracle. P4 needs 2 types of execution traces, EOSIO built-in functions and system smart contracts. The test oracles for detecting the Rollback vulnerability are as follows:

(i) The smart contract being fuzzed sent an inline action to reveal the winner.
(ii) The EOS token balance of the attack agent smart contract increased after fuzzing.

P5 utilizes test oracle (i) to determine whether the smart contract sent an inline action by checking the execution traces of *send_inline()*. Then, test oracle (ii) is used to confirm that P5 did exploit the rollback vulnerability within the smart contract being fuzzed.

4.6 P6. Receipt Hijacking

Vulnerability Description. Some EOSIO contracts that adopt the resource payment model named "Receiver pays" (section 2) has the potential to be exploited by attackers. Fig. 10 shows the gambling smart contract *biggame* with a Receipt Hijacking vulnerability. In line 6, *biggame* signs a transaction to call the function, *decerecept*, sending a receipt to notify the *user* that it has received the bet. In the function, *decereceipt* , *biggame* uses *require_recipient()* to make a copy of this action and forwards it to the *user*. Note that the transaction is signed by *biggame*, which means *biggame* will pledge tokens to get the computing resources that are needed to send and execute the transaction.

As shown in Fig. 11, an attacker could deploy an attack agent smart contract within his account, sending actions to the victim contract that adopts the resource payment model. Then, the attack agent smart contract monitors the transaction sent by the victim contract (Line 1). Then, the attacker utilizes the transaction to send other inline actions (Line 3-7, we name it receipt hijacking). Note that, the inline actions will be packed into the receipt transaction that *biggame* signed. As a result, *biggame* will pay for the computing resource needed to send and execute these inline actions. In this way, the attack agent smart contract can drain the computing resources of the victim smart contract quickly.

Attack Scenario. The attacker uses an attack agent smart contract to exploit the Receipt Hijacking vulnerability within a victim contract. However, P6 does not need to deploy an agent contract. This is because AntFuzzer has instrumented the *send_deferred()* and *require_receipt()*. In another word, we can tell from execution traces whether a smart contract signed a transaction to send a receipt without actually receiving it through an agent contract. To trigger

```
1 [[on_notify("biggame::decereceipt")]]
2 void hack(account user, string reply) {
3   action(permission_level(_self, name("active"))
      ,
4     _self,
5     name("mining"),
6     std::make_tuple()
7   ).send();
8 }
9 };
```

Fig. 11: An attack agent smart contract to exploit the Receipt Hijacking vulnerability

a Receipt Hijacking vulnerability within the smart contract being fuzzed. P6 randomly selects the functions of the smart contract to invoke or send some EOS tokens to it.

Test Oracle. P6 needs the execution traces of EOSIO built-in functions. The test oracles for detecting the Receipt Hijacking vulnerability are as follows:

(i) The smart contract being fuzzed signed a new transaction.
(ii) The smart contract being fuzzed forwarded the transaction to another account.

By checking the execution trace of *send_deferred()*, P6 can determine whether the smart contract being fuzzed signed a transaction using test oracle (i). Then, P6 uses test oracle (ii) to detect whether the transaction was forwarded to another account by checking the execution trace of *require_recipient()*.

5 Evaluation of AntFuzzer

5.1 Research Questions

We conducted extensive fuzzing experiments to evaluate AntFuzzer and our detection plugins and answer 4 research questions.

- *RQ1*. Can AntFuzzer facilitate the development of developing vulnerability detection plugins?
- *RQ2*. How accurate are AntFuzzer and the detection plugins in detecting vulnerabilities of EOSIO smart contracts?
- *RQ3*. How much code coverage can AntFuzzer improve compared to the existing black-box fuzzing framework?
- *RQ4*. What is the overhead of AntFuzzer?

5.2 RQ1: Amount of Code of Detection Plugins

We counted the lines of code of each detection plugin. As shown in Table. 1, the total code amount of the 6 detection plugins is 591 lines of Java code. The average code size of detection plugins is only 98 lines. Compared to the amount of code of the AntFuzzer framework (i.e. more than 9,000 lines of Java code), the code amount of detection plugins is much fewer.

TABLE 1: Amount of code of each detection plugin

Plugin	P1	P2	P3	P4	P5	P6
Lines of Java code	109	146	96	43	90	107

Answer to RQ1: AntFuzzer can facilitate the development of developing vulnerability detection plugins.

5.3 RQ2: Vulnerability Detection Accuracy of AntFuzzer

5.3.1 Experiment Setup

All the fuzzing experiments are performed within a Linux virtual machine. The host machine of the Linux virtual machine is equipped with an Intel i5-10400 CPU and 16GB of memory. The operating system of the Linux virtual machine is Ubuntu 20.04 LTS. We crafted AFL 2.53b to support fuzzing input generation with a specified length. We also instrumented nodeos v1.5.2 as the *Data collector* of AntFuzzer.

5.3.2 Creating Benchmark

First, we collect all the contracts with vulnerabilities that are disclosed by well-known smart contract auditing companies and blockchain security companies. We collected 31 unique smart contracts involved in these attacks in total. Of these attack reports, 5 attack reports only disclosed code snippets, but not the complete smart contract. We rebuilt the complete smart contract by adding code by hand. Then, we further verified the collected contracts to ensure that the attacks are indeed caused by the vulnerabilities of smart contracts and excluded 12 attacks that were caused by web attacks, compromised private keys, etc. Finally, we used 18 smart contracts related to publicly verified attacks as the ground truth. We also collected 71 smart contracts that have been open-sourced and deployed on EOSIO from GitHub and eosflare [1] by searching the keywords, EOSIO and smart contract.

We also extracted 5,500 smart contracts using the transaction data set provided by XBlocks-EOS [18]. Among them, 884 smart contracts have identical bytecode but are deployed on different accounts. We eliminated the duplicate ones and obtained 4,616 unique smart contracts. We utilized these smart contracts without source code to conduct large-scale fuzzing experiments.

5.3.3 Result on Smart Contracts with Source Code

The vulnerability detection result of AntFuzzer is shown in Table. 2. Plugin P1 and P2 reported 2 Fake EOS Transfer vulnerabilities and 4 Fake Transfer Notification vulnerabilities. After a manual code audit, we confirmed that the detection accuracy of P1 and P2 is 100.0%.

Plugin P3 reported 2 Block Information Dependency vulnerabilities within the 90 smart contracts. After code auditing, we found that P3 of AntFuzzer missed a vulnerable smart contract. The smart contract, *biggame*, is a false-negative case, which is shown in Fig. 12. The *biggame* uses the *transfer()* to handle the bets of users, and requires that users who want to make a bet must have a recommender specified in the *memo* (Line 4), and the recommender must be an account that exists in the *recommenderTable* (Line 6). But in the ABI of *biggame*, there is no function to read or modify the *recommenderTable*. Therefore, it is hard for AntFuzzer to trigger such vulnerability within hundreds of fuzzing iterations. Combining the fuzzing approach with other approaches such as symbolic execution may help detect such vulnerable contracts, we leave it as future work.

P4, P5 and P6 identified 5 No Permission Check, 4 Rollback and 2 Receipt Hijacking vulnerabilities, respectively.

TABLE 2: Result of AntFuzzer on Smart Contracts with Source Code (P - Positive, N - Negative, TP – True-Positive, FP – False-Positive, TN – True-Negative, FN - False-Negative)

Vulnerability	Plugin	Samples(P/N)	Reported	TP	FP	TN	FN	Accuracy(TP+TN)/(P+N)	F1-score
Fake EOS Transfer	P1	90(2/88)	2	2	0	88	0	100.0%	100.0%
Fake Transfer Notification	P2	90(4/86)	4	4	0	86	0	100.0%	100.0%
Block Information Dependency	P3	90(3/87)	2	2	0	87	1	98.9%	80.0%
No Permission Check	P4	90(5/85)	5	5	0	85	0	100.0%	100.0%
Rollback	P5	90(4/86)	4	4	0	86	0	100.0%	100.0%
Receipt Hijacking	P6	90(2/88)	2	2	0	88	0	100.0%	100.0%
Total	P1-P6	90(19/71)	19	19	0	70	1	98.9%	97.4%

```
class biggame: public eosio::contract {
void transfer(name receiver, name code, string
    memo) {
  // get the recommender
  string recommender = memo.substr(
      recommenderPrefix);
  // check whether the recommander is valid
  if (recmanderTable.find(recommender) !=
      recmanderTable.end()) {
    // accept the bet
  }
}
};
```

Fig. 12: A False Negative case of P3

And after manual checking, we confirmed that the detection accuracy of these plugins of AntFuzzer reported is 100%.

Table. 2 shows that the vulnerability detection accuracy of AntFuzzer reaches 98.9% (i.e. 89/90) and the F1-score of all plugins is 97.4%.

5.3.4 Comparison with the State-of-the-Art detection tools

In this section, we further compared AntFuzzer with the EOSFuzzer [10] and EVulHunter [9]. EOSafe [26] is a more advanced detection tool, but it is not open-sourced. Therefore, we could not compare it with AntFuzzer. Since EOSFuzzer and EVulHunter only support 3 and 2 types of vulnerabilities respectively, we compared AntFuzzer, EOSFuzzer [10], and EVulHunter [9] in terms of detecting Fake EOS Transfer vulnerability, the Fake Transfer Notification vulnerability, and Block Info Dependency vulnerability only. The results are shown in Table. 3.

For the fairness of the comparison, we did not use our benchmark dataset for comparison experiments but used 82 smart contracts provided by EOSFuzzer as the benchmark dataset.

The EVulHunter successfully analyzed 74 smart contracts. It failed to analyze the other 8 contracts. The EVulHunter identified 12 Fake Transfer Notification vulnerabilities. But we found that 10 of them are false positives after a manual code audit. In addition, we confirmed that the EVulHunter missed 2 Fake Transfer Notification vulnerabilities. Hence, AntFuzzer is more accurate than EVulHunter when detecting Fake Transfer Notification vulnerabilities.

P1 of AntFuzzer successfully reported the only vulnerable contract with the Fake EOS Transfer vulnerability and did not report any false-positive cases. EOSFuzzer reported 2 vulnerabilities, but one of them is a false-positive case, which is shown in Fig. 13.

```
1  class vigor: public eosio::contract {
2  void apply(uint64_t receiver, uint64_t code,
       uint64_t action) {
3    if (code == name("eosio.token").value ||
         action == name("transfer").value) {
4      execute_action(name(receiver), name(code), &
           vigor::assetin);
5    }
6    if (code == receiver) {
7      switch (action) {
8        EOSIO_DISPATCH_HELPER(vigor, (transfer), (
             create), ...);
9      }
10   }
11 }
12 };
```

Fig. 13: A false positive case of EOSFuzzer

Since EOSFuzzer and P1 of AntFuzzer use similar attack agent contracts, to figure out why EOSFuzzer wrongly reported a false-positive case, we further analyzed test oracles of EOSFuzzer for detecting the Fake Transfer vulnerability. EOSFuzzer uses 2 test oracles, the first sub-oracle is the same as P1 (i.e. the smart contract can receive the EOS token). As shown in Fig. 13, the false-positive case, *vigor*, uses *assetin()* to handle EOS transfer actions (Line 4 - Line 5), therefore, the first sub-oracle is satisfied. The second sub-oracle of EOSFuzzer is that the attack agent can directly invoke the *transfer()* function of the smart contract [10]. However, such a test oracle implicitly assumes that all smart contracts use the *transfer()* as the handler function of EOS transfer actions. For the contract *vigor*, it does have a function named *transfer()* that can be invoked directly by the attack agent contract. Therefore, the second test oracle of EOSFuzzer is satisfied as well. However, the *transfer()* function of the vigor contract is not used to handle EOS transfer actions (the corresponding function is *assetin()*).

Since P1 uses a more reasonable test oracle, it will compare the 2 destination addresses, the address of the *transfer()* function that the attack agent directly invokes, and the address of the EOS transfer action handler of this smart contract, which can effectively eliminate such false-positive cases.

AntFuzzer and EOSFuzzer identified all 4 vulnerable smart contracts with the Fake Transfer Notification vulnerability and did not report any false-positive cases.

TABLE 3: Comparison of AntFuzzer, EOSFuzzer and EVulHunter (FP - False-Positive, FN - False-Negative)

Vulnerability	Samples(P/N)	AntFuzzer			EOSFuzzer			EVulHunter		
		Reported	FP	FN	Reported	FP	FN	Reported	FP	FN
Fake EOS Transfer	82(1/81)	1	0	0	2	1	0	9	8	0
Fake Transfer Notification	82(4/78)	4	0	0	4	0	0	12	10	2
Block Information Dependency	82(3/79)	2	0	1	2	0	1	-	-	-

```
1  void reveal(name from, name to, asset quantity,
       std::string msg) {
2    if (msg[0] == 'v') {
3      int luck_num = tapos_block_prefix() + msg.
           length();
4      if (msg[1] == 'l') {
5        if (tapos_block_num % 2) {
6          transfer(get_self(), name("eosio"),
               quantity, msg);
7        }
8      }
9    }
10 }
```

Fig. 14: Modified smart contract eosshark

For the Block Information Dependency vulnerability, both AntFuzzer and EOSFuzzer reported 2 vulnerabilities and missed one vulnerable smart contract. The false-negative case has been discussed in section 5.3.2.

To figure out the fuzzing capability limits of AntFuzzer and EOSFuzzer, we further modified a smart contract, *eosshark*, to inject a Block Information Dependency vulnerability and placed the vulnerable code under a three-layer conditional judgment statement to make it more difficult to be triggered.

The modified smart contract, *eosshark*, is shown in Fig. 14. The function *reveal()* of *eosshark* sends EOS tokens after invoking *tapos_block_prefix()* and *tapos_block_num()* (Line6 and Line9). Therefore, this function has a Block Information Dependency vulnerability. However, to trigger the vulnerability, the first character of argument *msg* must be "v" and the second character must be "l" (Line 5-7). We manually analyzed the bytecode of this smart contract and found that these two characters do not appear in the static data section of the bytecode file. Therefore, the strategy that EOSFuzzer adopted (i.e. extracting string parameters from the data section of the bytecode) cannot generate valid fuzzing inputs for this function that matches the condition to trigger the vulnerability. In other words, for such functions, EOSFuzzer can only generate random parameters using the black-box fuzzing strategy. In our experiment, EOSFuzzer failed to identify this vulnerability after more than 8,000 fuzzing iterations.

Since the fuzzing input generation of AntFuzzer is guided by code coverage, after about only 4,000 iterations of fuzzing, AntFuzzer successfully found a string that starts with "l" and caused the code coverage to rise, and after performing 16 mutations based on the string, the vulnerability was successfully triggered. Considering the detection efficiency of AntFuzzer, the time overhead for thousands of mutations is acceptable.

5.3.5 Results on real-world smart contracts without source code

We performed extensive fuzzing experiments on 4,616 EOSIO smart contracts without source code. As shown in Table. 4, AntFuzzer identified 275 No Permission Check vulnerabilities and 274 Fake Transfer Notification vulnerabilities. Such results imply that these two types of vulnerabilities may be very common in real-world smart contracts. AntFuzzer also identified 5 Block Information Dependency vulnerabilities, 9 Rollback vulnerabilities, and 23 Receipt Hijacking vulnerabilities. The percentage of vulnerable smart contracts reaches 16.05% in total. We also used EOSFuzzer [10] to perform large-scale fuzzing on the 4,616 smart contracts. The results of EOSFuzzer are shown in Table. 5 for comparison. Note that EOSFuzzer only supports 3 types of vulnerabilities, we can only compare the results of Fake EOS Transfer, Fake Transfer Notification, and Block Information Dependency. EOSFuzzer identified 308 Fake EOS Transfer vulnerabilities, which is twice as many as that of AntFuzzer. We manually check the execution traces of the 153 smart contracts that are reported vulnerable by EOSFuzzer but are reported non-vulnerable by AntFuzzer. We found that 147 of them are confirmed to be false positives of EOSFuzzer, which are similar to vigor (section 5.3.4). That is, smart contracts use other functions instead of *transfer()* to handle EOS transfer actions. Such results prove the test oracles of AntFuzzer for detecting the Fake EOS Transfer vulnerability are more accurate and reasonable than that of EOSFuzzer. In addition, AntFuzzer identified 24 more Fake Transfer Notification vulnerabilities and 1 more Block Information Dependency vulnerability than EOSFuzzer. And we confirmed none of them are false positives after manually mounting attacks on them and checking the execution traces. The results show that, with the grey-box fuzzing technique, AntFuzzer can trigger and identify more vulnerabilities with high accuracy.

Since we can not obtain the source code of the 4,616 real-world smart contracts, manually verifying all the smart contracts to validate the results is an impossible task. The results give us a lower bound of the number of vulnerable smart contracts in the wild. We recommend that smart contract developers perform code audits and security checks before releasing their smart contracts.

Answer to RQ2: On the benchmark dataset, the vulnerability detection accuracy of AntFuzzer reaches 98.9% (i.e. 89/90) and the F1-score of all plugins is 97.4%. In the extensive fuzzing experiments on 4,616 real-world smart contracts, AntFuzzer identified 741 vulnerabilities in total.

TABLE 4: Results on real-world smart contracts of AntFuzzer

Vulnerability	Samples	Reported	Percentage
Fake EOS Transfer	4,616	155	3.36%
Fake Transfer Notification	4,616	274	5.93%
Block Information Dependency	4,616	5	0.10%
No Permission Check	4,616	275	5.96%
Rollback	4,616	9	0.19%
Receipt Hijacking	4,616	23	0.50%

TABLE 5: Results on real-world smart contracts of EOSFuzzer

Vulnerability	Samples	Reported	Percentage
Fake EOS Transfer	4,616	308	6.67%
Fake Transfer Notification	4,616	250	5.42%
Block Information Dependency	4,616	4	0.08%

5.4 RQ3: Code coverage improvements of AntFuzzer

We compared AntFuzzer with the black-box fuzzing framework EOSFuzzer [10] to illustrate how much code coverage AntFuzzer can improve. For the fairness of the comparison, we conducted 2 comparison experiments on the 82 smart contracts offered by EOSFuzzer as the benchmark to compare the code coverage of the AntFuzzer and EOSFuzzer. In the first one, we compared the code coverage of AntFuzzer and EOSFuzzer under the same fuzzing iterations limit. In the second one, we compared them under the same time limit.

In the first comparison experiment, we modified the source code of AntFuzzer and EOSFuzzer, so that they both perform 2,000 rounds of fuzzing on each function of each smart contract. Then, we utilized an instrumented EOSIO local node to calculate code coverage branch counts of AntFuzzer and EOSFuzzer separately.

The results of the first experiment show that AntFuzzer has achieved an improvement in code coverage on 31 smart contracts, and the code coverage improvement rate was 22.7%. That is, on 37.5% of smart contracts in the dataset, the fuzzing inputs generated by AntFuzzer successfully cover the execution paths that EOSFuzzer failed to cover.

After manually checking the remaining 51 smart contracts with no code coverage improvement, we concluded that these smart contracts can be divided into two categories. The first category of smart contracts has only a few execution paths, and the fuzzing inputs generated by EOSFuzzer can already cover all of them. In the second category of smart contracts with no code coverage improvement, some functions require a specified account or public key to invoke, and these accounts and public key information do not have clues in the ABI and bytecode files. For these highly specialized data types, it is difficult to generate fuzzing inputs that can trigger judgment conditions even if taking code coverage as feedback. So neither AntFuzzer nor EOSFuzzer can successfully invoke these functions. More specialized fuzzing input generation strategies or larger-scale fuzzing campaigns against these smart contracts are required to improve code coverage of these smart contracts, which is left as future work.

Since it is hard to choose an appropriate time limit. In the comparison experiment under the same time limit, we first conducted black-box fuzzing on the benchmark and recorded the test time as the time limit when the coverage of these smart contracts peaked. Then, we used AntFuzzer to conduct grey-box fuzzing within the time limit. In the second experiment, we observed that AntFuzzer also achieved an improvement in code coverage on 14 smart contracts, and the average code coverage improvement rate is 17.5%.

Answer to RQ3: Compared to the black-box fuzzing framework EOSFuzzer, AntFuzzer has achieved significant improvement in code coverage. The results of the comparison experiments proved that AntFuzzer could effectively trigger those hard-cover branches.

5.5 RQ4: Overhead of AntFuzzer

In the large-scale fuzzing experiments on the smart contracts without source code, each plugin of AntFuzzer was configured to perform 2,000 fuzzing iterations on each smart contract. The fuzzing experiment took about 4 minutes on average on each contract. By analyzing the logs of EOSIO, we found that during our experiments, the average fuzzing iterations per second are 15. In a contrast, the fuzzing iterations per second of EOSFuzzer in our experiment are 18. That is, AntFuzzer is only 16% slower than this black-box fuzzing tool.

Since AntFuzzer utilizes code coverage as a feedback indicator to guide the mutation of fuzzing inputs, the workflow of AntFuzzer introduces additional overhead compared to black-box fuzzing, namely fuzzing input transformation and code coverage information transformation. Fuzzing input transformation is transforming fuzzing input from the *AFL agent* to the *Data collector*, and code coverage information transformation is transforming the code coverage bitmap from the *Data collector* to the *AFL agent*.

To evaluate the detection efficiency of AntFuzzer in detail, we conducted a quantitative analysis of the overhead of each step. We took a smart contract with four functions as the test subject and used AntFuzzer to perform 2,000 rounds of fuzzing on each function, and then carried out statistical analysis on the overhead of each step. As shown in Table. 6, In the workflow of AntFuzzer, the most time-consuming steps are smart contract execution and test case generation, which account for 57.9% and 38.3% of the total time overhead, respectively. Since these two steps, fuzzing input generation and smart contract execution, are also present in the workflow of black-box fuzzing tools, the extra overhead of AntFuzzer is mainly caused by code coverage information transformation and fuzzing input transformation. The results of our experiment showed that the overhead of these 2 steps only accounts for 0.4% and 2.4% of the total time overhead, respectively. It is worth noting that the execution time of each step of the AntFuzzer workflow fluctuates depending on factors such as the complexity of the function being fuzzed, and the length of the fuzzing inputs that need to be transformed, etc. However, the maximum time of the code coverage information transformation measured in our experiments is only 4 ms, which is far less than the time consumption of the smart contract execution and test case generation. And the maximum time of fuzzing input trans-

TABLE 6: Overhead of AntFuzzer

	Fuzzing Input Generation	Fuzzing Input Transformation	Smart Contract Execution	Code Coverage Information Transformation
Maximum time (ms)	16	2	112	4
Minimum time (ms)	4	0	3	0
Average time (ms)	13.98	0.13	21.14	0.62
Average Percentage	38.3%	0.4%	57.9%	2.4%

formation is only 2 ms. Hence, the extra communication overhead of AntFuzzer is almost negligible.

Answer to RQ4: By adopting several performance optimization designs, AntFuzzer achieves a good balance between detection capability and detection efficiency. The average fuzzing iterations per second of AntFuzzer are 15, which is fast enough for volume fuzzing experiments and only 16% less than that of the black-box fuzzing tool, EOSFuzzer [10].

6 Related Work

This work is related to vulnerability detection. The existing detection approaches for vulnerability detection of smart contracts can be divided into the following three categories:

Formal Verification The formal verification approach [28], [29], [30], [31] uses theorem provers or formal methods of mathematics to prove the specific properties in a smart contract [25] (e.g. functional correctness, run-time safety, etc). However, existing formal verification methods suffer from massive manual efforts to define the correct properties of smart contracts, which makes it challenging for using them to perform automated vulnerability detection [24].

Symbolic Execution The symbolic execution approach [5], [11], [19], [23], [26], [27], [33] could effectively explore smart contracts to examine vulnerable patterns and flaws which would be expected in the run-time [25]. Yet, the symbolic execution approach suffers from path explosion, which adds extra difficulty to applying them to analyze complex smart contracts [22].

Fuzzing The fuzzing approach provides random transactions as inputs for smart contracts. The smart contracts are then monitored for unexpected behaviors or vulnerability patterns. However, most of the existing fuzzing tools [10], [12], [20], [32] use black-box fuzzing that suffers from low code coverage and could miss many vulnerabilities. To address this issue, some grey-box fuzzing techniques have been proposed [14], [21] and proved to be effective in improving code coverage. Yet, to the best of our knowledge, these existing studies focus on Ethereum smart contracts, and there are currently no studies on grey-box fuzzing of EOSIO smart contracts. Moreover, since the two platforms share little similarities in virtual machines, the structure of bytecode, and the types of vulnerabilities [26], applying them to EOSIO smart contracts is a very challenging task.

7 Conclusion

In this paper, we proposed AntFuzzer, a highly extensible grey-box fuzzing framework for detecting vulnerabilities of EOSIO smart contracts. Based on AntFuzzer, we implemented 6 vulnerability detection plugins. We fully validated the effectiveness and performance of AntFuzzer and our detection plugins on two datasets of smart contracts with source code and bytecode only. Experimental results show that AntFuzzer has certain advantages over the tools in vulnerability detection accuracy, code coverage, and the number of supported vulnerability types. In addition, AntFuzzer's extensibility and ease of use make this work effectively against the vulnerabilities of EOSIO smart contracts that may arise in the future.

References

[1] eosflare, "Eos blockchain explore," 2020. [Online]. Available: https://eosflare.io/

[2] Google, "american fuzzy lop," 2022. [Online]. Available: https://github.com/google/AFL

[3] Block.one, "Eosio overview," 2020. [Online]. Available: https://developers.eos.io/manuals/eos/latest/index

[4] Block.one, "Accounts and permissions," 2020. [Online]. Available: https://developers.eos.io/welcome/latest/overview/core_concepts

[5] ConsenSys, "Security analysis tool for evm bytecode." 2019. [Online]. Available: https://github.com/ConsenSys/mythril

[6] Block.one, "Resource payer," 2019. [Online]. Available: https://eos.io/news/eosio-2-2/

[7] SlowMist, "Rollback attack about blacklist in eos," 2019. [Online]. Available: https://slowmist.medium.com/roll-back-attack-about-blacklist-in-eos-adf53edd8d69

[8] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.

[9] L. Quan, L. Wu, and H. Wang, "Evulhunter: detecting fake transfer vulnerabilities for eosio's smart contracts at webassembly-level," *arXiv preprint arXiv:1906.10362*, 2019.

[10] Y. Huang, B. Jiang, and W. K. Chan, "Eosfuzzer: Fuzzing eosio smart contracts for vulnerability detection," in *12th Asia-Pacific Symposium on Internetware*, ser. Internetware'20. New York, NY, USA: Association for Computing Machinery, 2020, p. 99–109. [Online]. Available: https://doi.org/10.1145/3457913.3457920

[11] D. Wang, B. Jiang, and W. Chan, "Wana: Symbolic execution of wasm bytecode for cross-platform smart contract vulnerability detection," *arXiv preprint arXiv:2007.15510*, 2020.

[12] B. Jiang, Y. Liu, and W. Chan, "Contractfuzzer: Fuzzing smart contracts for vulnerability detection," in *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2018, pp. 259–269.

[13] B. McFadden, T. Lukasiewicz, J. Dileo, and J. Engler, "Security chasms of wasm," *NCC Group Whitepaper*, 2018.

[14] V. Wüstholz and M. Christakis, "Harvey: A greybox fuzzer for smart contracts," in *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2020, pp. 1398–1409.

[15] N. He, R. Zhang, L. Wu, H. Wang, X. Luo, Y. Guo, T. Yu, and X. Jiang, "Security analysis of eosio smart contracts," *arXiv preprint arXiv:2003.06568*, 2020.

[16] D. Lehmann, J. Kinder, and M. Pradel, "Everything old is new again: Binary security of webassembly," in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 217–234.

[17] S. Lee, D. Kim, D. Kim, S. Son, and Y. Kim, "Who spent my EOS? on the (in) security of resource management of eos. io," in *13th USENIX Workshop on Offensive Technologies (WOOT 19)*, 2019.

[18] W. Zheng, Z. Zheng, H.-N. Dai, X. Chen, and P. Zheng, "Xblock-eos: Extracting and exploring blockchain data from eosio," *Information Processing & Management*, vol. 58, no. 3, p. 102477, 2021.

[19] Z. Wang, B. Wen, Z. Luo, and S. Liu, "Mar: A dynamic symbol execution detection method for smart contract reentry vulnerability," in *International Conference on Blockchain and Trustworthy Systems*. Springer, 2021, pp. 418–429.

[20] M. Ding, P. Li, S. Li, and H. Zhang, "Hfcontractfuzzer: fuzzing hyperledger fabric smart contracts for vulnerability detection," in *Evaluation and Assessment in Software Engineering*, 2021, pp. 321–328.

[21] V. Wüstholz and M. Christakis, "Targeted greybox fuzzing with static lookahead analysis," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, ser. ICSE '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 789–800. [Online]. Available: https://doi.org/10.1145/3377811.3380388

[22] R. Cao, T. Chen, T. Li, X. Luo, G. Gu, Y. Zhang, Z. Liao, H. Zhu, G. Chen, Z. He, Y. Tang, X. Lin, and X. Zhang, "Soda: A generic online detection framework for smart contracts," in *Proceedings of the 27th Network and Distributed System Security Symposium (NDSS)*, Feb. 2020, pp. 1–17, network and Distributed System Security Symposium (NDSS) 2020 ; Conference date: 23-02-2020 Through 26-02-2020.

[23] J. Huang, J. Jiang, W. You, and B. Liang, "Precise dynamic symbolic execution for nonuniform data access in smart contracts," *IEEE Transactions on Computers*, vol. 71, no. 7, pp. 1551–1563, 2022.

[24] 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 *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, 2019, pp. 1503–1520.

[25] P. Praitheeshan, L. Pan, J. Yu, J. Liu, and R. Doss, "Security analysis methods on ethereum smart contract vulnerabilities: a survey," *arXiv preprint arXiv:1908.08605*, 2019.

[26] N. He, R. Zhang, H. Wang, L. Wu, X. Luo, Y. Guo, T. Yu, and X. Jiang, "EOSAFE: Security analysis of EOSIO smart contracts," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 1271–1288.

[27] S. So, S. Hong, and H. Oh, "SmarTest: Effectively hunting vulnerable transaction sequences in smart contracts through language Model-Guided symbolic execution," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 1361–1378.

[28] X. Li, C. Su, Y. Xiong, W. Huang, and W. Wang, "Formal verification of bnb smart contract," in *2019 5th International Conference on Big Data Computing and Communications (BIGCOM)*. IEEE, 2019, pp. 74–78.

[29] Z. Liu and J. Liu, "Formal verification of blockchain smart contract based on colored petri net models," in *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 2. IEEE, 2019, pp. 555–560.

[30] T. Sun and W. Yu, "A formal verification framework for security issues of blockchain smart contracts," *Electronics*, vol. 9, no. 2, p. 255, 2020.

[31] T. Abdellatif and K.-L. Brousmiche, "Formal verification of smart contracts based on users and blockchain behaviors models," in *2018 9th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, 2018, pp. 1–5.

[32] T. D. Nguyen, L. H. Pham, J. Sun, Y. Lin, and Q. Minh, "sfuzz: An efficient adaptive fuzzer for solidity smart contracts," in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2020, pp. 778–788. [Online]. Available: https://doi.ieeecomputersociety.org/

[33] L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor, "Making smart contracts smarter," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 254–269.

[34] R. Kersten, K. Luckow, and C. S. Păsăreanu, "Poster: Afl-based fuzzing for java with kelinci," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 2511–2513. [Online]. Available: https://doi.org/10.1145/3133956.3138820

[35] A. Herrera, H. Gunadi, S. Magrath, M. Norrish, M. Payer, and A. L. Hosking, "Seed selection for successful fuzzing," in *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 230–243. [Online]. Available: https://doi.org/10.1145/3460319.3464795

Jianfei Zhou received his B.S. degree from China University of Petroleum in 2020. Now he is currently pursuing his M.S. degree in UESTC. His current research interests include Blockchain, Smart Contracts.

Tianxing Jiang received his B.S. degree from Chongqing University of Posts and Telecommunications in 2020. Now he is currently pursuing his M.S. degree in UESTC.

Shuwei Song received his B.S. degree from University of Electronic Science and Technology of China (UESTC). Now he is currently working toward the PhD degree with UESTC. His current research interests include Blockchain, Smart Contracts.

Ting Chen received his PhD degree from University of Electronic Science and Technology of China (UESTC), China, 2013. He is an Professor in the School of Computer Science and Engineering in UESTC. His research interest focuses on blockchain, smart contract and software security. He has published tens of high quality papers in prestigious conferences and journals. His work received several best paper awards, including INFOCOM 2018 best paper award.