Why Smart Contracts Reported as Vulnerable were not Exploited?

Tianyuan Hu, Jingyue Li, Bixin Li, André Storhaug

Abstract—Smart contract security is essential for blockchain applications. Studies show that few of the reported vulnerabilities are exploited. However, no follow-up study is performed to why the reported vulnerabilities are not exploited. We aim to understand the reasons for the low exploitation rate to help improve vulnerability detection practices. We first collect 136,969 unique real-world smart contracts and analyze them using seven vulnerability detectors. Then, we apply Strauss' grounded theory approach to understand if they are exploitable. In addition, we analyze the transaction logs of the exploitable vulnerabilities to understand their exploitations in history. Among the 4,364 smart contracts reported as vulnerable by the vulnerability detectors, 75.27% of them are unexploitable, and only 66 (0.015%) have been exploited. We uncover 11 reasons for making the detectors misidentify unexploitable vulnerabilities and six reasons for demotivating and preventing the attackers from exploiting the exploitable ones. We illustrate that: beyond treating the smart contracts as yet another Object Oriented (OO) application, it is essential to consider the Solidity programming language's design principle, smart contracts' application scenarios, and their execution environments. Our results can help differentiate exploitable smart contracts to help allocate efforts to exploitable ones to mitigate emergent risks.

Index Terms—Ethereum, smart contract, vulnerability detection, source code analysis

1 Introduction

Due to blockchains' monetary and anonymous nature, they are targets of adversaries. The security of smart contracts is critical because they may handle and store digital assets worth millions of dollars. The DAO hack [1] exploiting the reentrancy vulnerability in contract code resulted in a 60 million dollars loss. It is, therefore, imperative to prune out smart contracts' security problems before deploying them.

Many methods and corresponding detectors, e.g., *Oyente* [2], *Securify* [3], *sFuzz* [4], *ContractFuzzer* [5], *ContraMaster* [6], *DefectChecker* [7], *EXGEN* [8], *HONEYBADGER* [9], have been proposed to detect smart contract vulnerabilities. Ren et al. [10] point out that the detectors are evaluated by the tool authors using different datasets and metrics, which may result in biased conclusions. Several empirical studies were conducted by other researchers using manually annotated datasets or real-world smart contracts to evaluate these detectors fairly. *SolidiFI* [11] is used to evaluate six detectors [2], [3], [12], [13], [14], [15]. The results show that none of the detectors detect all the injected bugs correctly, and all the evaluated detectors report several false positives. Durieux et al. [16] evaluated nine vulnerability detectors [2], [3], [9], [12], [13], [14], [15], [17], [18], and found that 97% of the real-world contracts analyzed were labeled as vulnerable by the detectors. Perez and Livshits [19] analyzed 821,219 real-world contracts using six detectors, namely, *Oyente* [2], *ZEUS* [20], *MAIAN* [17], *Securify* [3], *TEETHER* [21], *Madmax* [22]. They classified the analyzed smart contracts as:

- **Vulnerable:** A contract is reported as vulnerable if the vulnerability detector flags it.
- **Exploitable:** A contract is exploitable if an attacker could exploit its vulnerability and cause security compromise.
- **Exploited:** A contract is exploited if a transaction on Ethereum's main network has triggered one of its vulnerabilities.

Results of the study by Perez and Livshits [19] show that, among the 73,62 contracts reported as vulnerable by at least two detectors, only 463 contracts were exploited. They hypothesized that most reported vulnerabilities are unexploitable. However, no follow-up study tried to confirm the hypothesis and understand the reasons for the low exploitation rate. "*Vulnerability deals with the theoretical, and exploitability deals with actuals. Understanding what is vulnerable and what remains exploitable can help companies prioritize and acknowledge where their security efforts can be improved.* [23]" Thus, we are motivated to answer two research questions (RQs):

- **RQ1:** Are vulnerable smart contracts reported by vulnerability detectors exploitable? If the exploitability is low, what are the possible reasons for that?
  If a smart contract is reported as vulnerable, we want to know if the adversary could execute the vulnerability and if the execution could lead to security compromises. In addition, we want to know what causes the vulnerability from being unexploitable.
- **RQ2:** Are exploitable smart contracts exploited? If not, what prevented attackers from exploiting them?
  If a smart contract is exploitable, we want to know if it has been executed at least once, although we

- *T. Hu and B. Li are with the School of Computer Science and Engineering, Southeast University, Nanjing, 211189, China.E-mail: tianyuan.hu@foxmail.com, bx.li@seu.edu.cn*
- *J. Li and A. Storhaug are with the Department of Computer science, Norwegian University of Science and Technology, Trondheim, Norway. E-mail: jingyue.li@ntnu.no, andre.storhaug@ntnu.no*

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

may not know the actual loss due to the execution. In addition, we want to identify commonalities between the exploitable smart contracts that have and have not been executed.

To answer RQ1, we collected 136,969 unique real-world smart contracts and used four efficient vulnerability detectors, namely, *Oyente* [2], *SmartCheck* [13], *Slither* [14], *SoliDetector* [24] to label their vulnerabilities. Then, we manually analyzed the source code of the vulnerable contracts to judge whether they are exploitable and adopted Strauss' grounded theory method [25] to identify the reasons for the low exploitability. The insights from the manual analysis are further verified by three vulnerability detectors, namely, *Mythril* [12], *ConFuzzius* [26], and *Smartian* [27]. To answer RQ2, we collected the transaction logs of the reported vulnerable contracts and replayed their transactions on a full Ethereum node. We designed transaction log analysis rules to identify the vulnerability exploitation and also used Strauss' grounded theory method to understand the contracts' exploitations.

Results show that 4,364 contracts are labeled as vulnerable by at least two of the effective vulnerability detectors. The vulnerabilities cover ten types. Among the 4,364 reported vulnerable contracts, 3,285 are unexploitable. After analyzing the 4,106,134 transaction logs of the 4,364 vulnerable contracts, we found that only 66 exploitable contracts had been exploited. Through open and axial coding, we identified 11 reasons causing the detectors to misidentify unexploitable vulnerabilities:

- missing path feasibility analysis
- overlooking preventive execution condition
- insufficient data flow analysis
- overlooking access control
- neglecting constraints caused by factory patterns
- neglecting constraints caused by contract inheritance
- assuming all fallback functions receive ether
- insufficient analysis of the values of the target contracts' addresses
- omitting the case that the ether transfer initiator is the ether's initial owner
- assuming critical operations after authorization
- assuming status inconsistency when the function call results are not checked.

We have also found six reasons that may have demotivated adversaries to attack the exploitable contracts:

- very little or no financial benefits for attacker
- insignificant impacts of the compromise
- attacker must develop attack contracts
- attacker must deposit ether as a prerequisite
- attacker must be lucky in random number competition
- attacker must be mining winner

Through selective coding, we concluded that the studied vulnerability detectors mainly adapted existing approaches to analyze OO applications. The detectors have not sufficiently considered Solidity's principle as a contract-oriented programming language for applications running on Ethereum Virtual Machine (EVM) and blockchain to securely transfer assets or ether between supplier and client and minimize the gas cost of execution.

TABLE 1: Smart Contract Vulnerability Detectors (*SC* represents source code; *BC* represents bytecode)

Year and ref.	Detector name	Vul. types covered	Inputs
Pattern Matching			
2018 [13]	SmartCheck	37 types	*SC*
2021 [24]	SoliDetector	20 types	*SC*
Symbolic Execution			
2016 [2]	Oyente	6 types	*SC*
2018 [20]	ZEUS	7 types	*SC*
2018 [18]	Osiris	Integer Vulnerability	*BC*
2018 [12]	Mythril	SWC Registry	*SC*
2018 [3]	Securify	37 types	*SC*/*BC*
2018 [21]	TEETHER	4 types	*BC*
2018 [17]	MAIAN	3 types	*SC*/*BC*
2019 [9]	HONEYBADGER	Honeypots	*BC*
2021 [7]	DefectChecker	8 types	*SC*
2022 [8]	EXGEN	4 types	*SC*
Data Flow Analysis			
2018 [22]	MadMax	3 types	*BC*
2019 [14]	Slither	71 types	*SC*
2020 [28]	Clairvoyance	Reentrancy	*SC*
2020 [29]	Ethainter	5 types	*BC*
Machine Learning			
2019 [30]	GNN-based	3 types	*SC*
2020 [31]	ContractWard	6 types	Opcode
2021 [32]	VSCL	6 types	*BC*
Fuzzing			
2018 [5]	ContractFuzzer	7 types	*BC*+ABI
2018 [33]	Reguard	Reentrancy	*SC*/*BC*
2020 [4]	sFuzz	9 types	*BC*
2020 [34]	Ethploit	3 types	*SC*
2021 [26]	ConFuzzius	10 types	SC
2021 [27]	Smartian	13 types	BC
2022 [6]	ContraMaster	5 types	SC

The contributions of this study are:

- We have derived novel insights and theories regarding the reasons for the low exploitation rate, which can help security practitioners differentiate vulnerable and exploitable smart contracts when analyzing and ranking vulnerabilities.
- We have created a benchmark dataset containing 4,364 real-world Solidity smart contracts, which are manually labeled with ten types of vulnerabilities. The dataset is around 20 times bigger than the similar state-of-the-art benchmark [10]. The dataset can help evaluate the vulnerability detectors' ability to detect vulnerabilities and their exploitability and is available at https://github.com/1052445594/SC_UEE.

The rest of the paper is organized as follows. Section 2 introduces related work, and Section 3 presents the research design. The answers to RQ1 and RQ2 are given in Sections 4 and 5, respectively. Section 6 discusses the results and Section 7 concludes.

2 Related Work

The approaches to detect smart contract vulnerabilities can be classified into pattern matching, symbolic execution, dependency analysis, machine learning (ML), and fuzzing, as shown in Table 1.

2.1 Detectors Using Pattern Matching Approaches

SmartCheck [13] translates Solidity source code into an XML-based intermediate representation and checks it against

XPath patterns. *SoliDetector* [24] is a static detection tool based on the knowledge graph of Solidity source code. For each smart contract to analyze, it constructs the knowledge graph containing the ontology and instance layers. Based on the knowledge graph, it uses the SPARQL [35] query to manipulate the knowledge graph and identify vulnerabilities.

2.2 Detectors Relying on Symbolic Execution

Oyente [2] is the first smart contract vulnerability detector based on symbolic execution. It builds control flow graphs from smart contract bytecode and symbolically executes the contract to identify vulnerabilities by analyzing execution traces. By analyzing dependency diagrams of smart contracts, *ZEUS* [20] combines abstract interpretation and symbolic execution to model smart contracts. As *ZEUS* analyzes artifacts at the low-level virtual machine (LLVM) intermediate level, it cannot locate vulnerabilities at the EVM bytecode level. *Osiris* [18] is a framework that combines symbolic execution and taint analysis to detect vulnerabilities related to arithmetic operations in Ethereum smart contracts. *Mythril* [12] uses symbolic execution and taint analysis to detect vulnerabilities. It performs decompilation and produces execution traces using a dynamic symbolic execution engine called Laser-EVM. However, *Mythril* is slow due to multiple symbolic executions. *Securify* [3] combines abstract interpretation and symbolic execution. The tool automatically classifies behaviors of a contract into three categories, compliance (matched by compliance properties), violation (matched by violation properties), and warning (no matches). Krupp et al. [21] first give a generic definition of vulnerable contracts and build *TEETHER*, a tool that employs symbolic execution to create an exploit automatically. However, *TEETHER* has difficulty solving hard constraints in execution paths and cannot simulate the blockchain behaviors very well, causing a loss of coverage. *MAIAN* [17] is a symbolic execution tool analyzing EVM bytecode. *MAIAN* classifies vulnerable contracts into three categories, namely, greedy, prodigal, and suicidal. *HONEYBADGER* [9] uses symbolic execution and pre-defined heuristics to expose honeypots. *DefectChecker* [7] symbolically executes the smart contract bytecode and generates their control flow graphs, stack events, and other features. Based on the generated information, it uses eight rules to detect different vulnerabilities. However, the public version of *DefectChecker* supports only Solidity 0.4.24. *EXGEN* [8] generates multiple transactions as exploits to vulnerable smart contracts and verifies the generated contracts' exploitability on a private chain with values crawled from the public chain. *EXGEN* is a cross-platform framework that supports Ethereum or EOS contracts.

2.3 Detectors Applying Data Flow Analysis

MadMax [22] is a gas-focused vulnerability detection tool consisting of a decompiler, which converts low-level EVM bytecode to code represented using an intermediate language. It then analyzes the code to detect out-of-gas vulnerabilities that require coordination across multiple transactions. *Slither* [14] is a highly scalable static analysis tool. It first converts Solidity smart contracts to an intermediate representation called SlithIR through control flow graph analysis. Then, it applies both data flow and taint analysis to detect vulnerabilities. *Clairvoyance* [28], [36] presents a static analysis tool that models cross-function and cross-contract behavior to detect the reentrancy vulnerability. Brent et al. [29] present *Ethainter* to detect composite vulnerabilities that escalate a weakness through multiple transactions. Based on the Datalog language [37] and the Soufflé Datalog engine [38], *Ethainter* constructs graphs containing data flow and control flow dependencies to identify vulnerabilities.

2.4 Detectors Using Machine Learning Technologies

Zhuang et al. [30] use a graph neural network (GNN) to classify vulnerable smart contracts. *ContractWard* [31] is a machine learning-based vulnerability detection tool targeting six vulnerabilities. It employs three supervised ensemble classification algorithms, namely, XGBoost, AdaBoost, and Random Forest (RF), and two classification algorithms, namely, Support Vector Machine (SVM) and *k*-Nearest Neighbor (KNN). Their evaluations show that XGBoost is the best-performing classifier algorithm. *VSCL* [32] is a smart contract vulnerability detection framework that constructs a control flow graph (CFG) to understand program run time behavior. Further, n-gram and Term Frequency–Inverse Document Frequency (TFIDF) techniques are used to generate numeric values (vectors) to present features of smart contracts. Finally, the generated feature matrix is used as input for the deep neural network (DNN) model.

2.5 Detectors Using Fuzz Testing

Fuzz testing [39] is an automated testing technique for analyzing computer programs. *ContractFuzzer* [5] is a fuzzing tool that generates random inputs to smart contracts according to the contracts' Application Binary Interface (ABI). *ContractFuzzer* defines a set of predefined test oracles that describes specific vulnerabilities. However, due to the randomness of the inputs, *ContractFuzzer*'s execution covers only limited system behavior. *ReGuard* [33] is a fuzzing tool to detect reentrancy vulnerabilities. It first converts the input to smart contracts into a C++ program via the Abstract Syntax Tree (AST) or CFG and generates random inputs to perform the fuzzing. *sFuzz* [4] employs an efficient, lightweight, adaptive strategy for selecting seeds to improve the fuzzing method based on random input generator [5]. *EthPloit* [34] adopts static taint analysis to generate exploit-targeted transaction sequences. It uses a dynamic seed strategy to pass hard constraints and an instrumented EVM to simulate blockchain behaviors. *ContraMaster* [6] is an oracle-supported dynamic exploit generation framework that can mutate transaction sequences. It uses data flow, control flow, and dynamic contract state to guide its mutations of the transaction sequences. *ConFuzzius* [26] combines evolutionary fuzzing, constraint solving, and dynamic data flow to generate test cases and detect vulnerabilities. *Smartian* [27] conducts both static and dynamic analysis for fuzz testing smart contracts. It analyzes the EVM bytecode to predict the transaction sequences and uses the dynamic data flow to guide the test case generation.

2.6 Empirical Evaluations of Vulnerability Detectors

Although studies proposing new vulnerability detectors always provide evaluation results, the evaluations can be biased. The studies may use different terms and definitions of the same vulnerability and use datasets that favor their detectors. Thus, other researchers performed empirical studies as shown in Table 2 to evaluate and compare the smart contract vulnerability detectors.

TABLE 2: Empirical Study of Vulnerability Detectors

Year and ref.	2020 [11]	2020 [16]	2021 [10]	2021 [19]
SmartCheck [13]	✓	✓		
Oyente [2]	✓	✓	✓	✓
ZEUS [20]				✓
Securify [3]	✓	✓		✓
Mythril [12]	✓	✓	✓	
Slither [14]	✓	✓		
Manticore [15]	✓	✓		
MAIAN [17]		✓		✓
Orisis [18]		✓	✓	
HONEYBADGER [9]		✓		
ContractFuzzer [5]			✓	
TEETHER [21]				✓
MadMax [22]				✓
sFuzz [4]			✓	

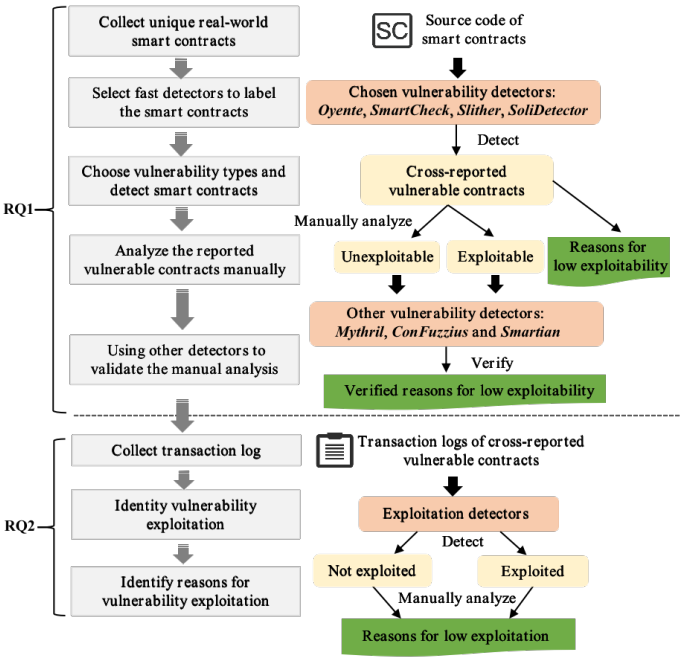

Fig. 1: Workflow of the study

Ghaleb et al. [11] proposed *SolidiFI* to evaluate six static vulnerability detectors [2], [3], [12], [13], [14], [15] using a dataset with injected vulnerabilities. Experiment results on a set of 50 contracts injected with 9,369 distinct vulnerabilities show that the evaluated detectors do not detect several instances of vulnerabilities despite their claims of being able to detect such vulnerabilities. Only one tool, i.e., *Slither* [14], detected all injected reentrancy and TxOrigin vulnerabilities. Ghaleb et al. [11] also found that all evaluated detectors have reported several false positives, ranging from 2 to 801 for different vulnerability types. However, they only manually analyze the vulnerabilities that are not reported by the majority of the detectors because manually inspecting all reported vulnerabilities involves a tremendous amount of effort and is therefore impractical. As a result, the number of false positives is underestimated.

Ferreira et al. [40] presented *SmartBugs*, an extensible and easy-to-use execution framework that simplifies the execution of detectors analyzing Solidity smart contracts. *SmartBugs* supports ten detectors [2], [3], [9], [12], [13], [14], [15], [17], [18], [41] and provides two datasets of Solidity smart contracts. One dataset contains 143 annotated vulnerable contracts with 208 tagged vulnerabilities, and another contains 47,518 unique contracts collected through Etherscan [42]. However, the 47,518 real-world contracts are not manually labeled. By using *SmartBugs*, Durieux et al. [16] evaluated nine detectors [2], [3], [9], [12], [13], [14], [15], [17], [18]. The evaluation was based on 69 annotated vulnerable smart contracts and all the real-world smart contracts in *SmartBugs*. The evaluation results showed that 97% of the real-world contracts were labeled as vulnerable. Durieux et al. [16] questioned that many reported vulnerabilities are false positives.

Ren et al. [10] evaluated six detectors [2], [4], [5], [12], [18], [43], and proposed a unified standard to eliminate the evaluation biases. They constructed a benchmark suite with three datasets, including unlabeled real-world contracts (UR), contracts with manually injected vulnerabilities (MI), and confirmed vulnerable contracts (CV). The experiment results on these datasets demonstrated that different choices of experimental settings could significantly affect tool performance and lead to misleading or even opposite conclusions. The experiment results in [14] show that *SmartCheck* has more false positives than *Slither*. However, Ren's study [10] gives opposite conclusions and shows that *SmartCheck* reports fewer false positives than *Slither* on UR and MI datasets.

Perez et al. [19] evaluated six detectors [2], [3], [17], [20], [21], [22] on real-world smart contracts and found many contradict results from different detectors [19]. Taking the reentrancy vulnerability as an example, *Oyente* and *Securify* agree on only 23% of the contracts reported as vulnerable to reentrancy, while *ZEUS* does not agree with any other detectors [19]. In addition, they analyzed more than 20 million Ethereum blockchain transactions and found that only 463 contracts related to six vulnerability types were exploited. Based on the evaluation results, they questioned whether the vulnerabilities reported by the evaluated detectors were either false positives or unexploitable.

Although the aforementioned empirical studies hypothesized that existing vulnerability detectors report many false positives or that the reported vulnerabilities are unexploitable, especially for real-world contracts, no follow-up study was performed to confirm the hypothesis and to understand the reasons for possible low exploitability. The insights could help security practitioners rank the reported vulnerabilities to allocate the security effort to the most urgent vulnerabilities to fix, i.e., those exploitable ones.

3 RESEARCH DESIGN

3.1 Research Design to Answer RQ1

To answer RQ1, we designed the research flow as shown in Figure 1.

3.1.1 Step 1. Collect unique real-world smart contracts

We crawled all available smart contracts with at least one transaction from Etherscan [42] on 1st April 2022. As there are duplicated smart contracts, we filtered contracts for uniqueness with a similarity threshold of 0.9, calculated using the Jacard index [44]. This means that if two contracts' code shares more than 90% of the tokens, one of the contracts will be discarded. The low uniqueness requirement is due to the often large amount of embedded library code. If the requirement is set to high, the actual contract code will be negligible compared to the library code. Most contracts will be discarded, and the resulting dataset will contain mostly unique library code.

3.1.2 Step 2. Select fast vulnerability detectors

For the vulnerability detector shown in Table 1, some, e.g., *SoliDetector* [24], run fast. *SoliDetector* took on an average of 1.33s to check 20 vulnerabilities on one real-world contract. A few others, i.e., *Mythril* [12] and fuzzing-based detectors, are a lot slower. The study [7] demonstrated that *Mythril* is a slow tool and the maximum time to analyze a smart contract using *Mythril* is 2480.26s. The fuzz-based detectors, e.g., *ConFuzzius* [26], *Smartian* [27], need to set a time-out, which is usually more than 10 minutes for each real-world contract to get decent detection accuracy. As we want to analyze a large amount of real-world smart contracts, we decide to first apply fast detectors to identify vulnerabilities. Then, we can use the slower vulnerability detectors to verify if the reported smart contracts by the fast detectors are vulnerable and exploitable.

We use the following criteria to select the fast vulnerability detectors.

- *The detectors shall take smart contract Source Code as Input (SCI)*: As we want to confirm whether the vulnerable contract is exploitable and understand the reasons for the possible low exploitability, we need to access the smart contract source code. Thus, we exclude detectors that do not analyze Solidity smart contract source code.
- *The detectors shall provide Vulnerability Localization (VL)*: To analyze the reported vulnerabilities precisely, we only consider detectors that provide the location, i.e., code line number, of the vulnerabilities at the source code level. The detectors only label the smart contract as vulnerable without providing vulnerability location information are excluded.
- *The detectors shall support multiple Solidity Versions (SV)*: When we crawl real-world smart contracts, we get smart contracts developed using various Solidity versions. If the detectors support only a limited number of versions of Solidity, the smart contracts they can analyze are limited, meaning we cannot get sufficient vulnerable smart contracts to study. Thus, we require the detectors to support several versions of Solidity.
- *The detectors shall be available to us (Available)*: Not all papers make their detectors publicly available. We exclude detectors we cannot access.

3.1.3 Step 3. Choose vulnerability types to focus on and use the selected detectors to detect the chosen types of smart contracts

The vulnerability detection results from a particular detector can be biased by the detectors' design flaws or bugs. As we want to identify generic reasons for the low exploitability, we choose to detect only the vulnerability types supported by at least two detectors to reduce the possible biases introduced by a single detector.

After using the detectors to detect the chosen smart contracts on the chosen vulnerability types, we get detection results containing vulnerability type names and locations.

3.1.4 Step 4. Analyze the reported vulnerable contracts manually

The vulnerability detectors report different vulnerabilities and their locations for the same smart contracts. Again, to avoid the biases caused by a single detector, we analyze the smart contracts labeled as vulnerable to a particular vulnerability type by more than one detector. The chosen smart contracts are, hereafter, called cross-reported vulnerable contracts.

Our study aims to understand why reported vulnerabilities are not exploited. We believe that there must be unknown reasons for the low exploitation rate. Thus, we use Strauss' grounded theory approach [25], often used to identify generic and unknown theories from data. Strauss' grounded theory approach [25] is an iterative and recursive approach where the researchers must go back and forth until they achieve theoretical saturation. Our grounded theory analysis included several steps. First, we read the source code of each smart contract reported as vulnerable. We classified them into two categories, i.e., exploitable or unexploitable, in parallel with root cause analysis and open coding to categorize the reasons for the low exploitability. As a second step, these codes are grouped into conceptual categories through axial coding. We did a constant comparison and theoretical saturation to consolidate the reasons for the low exploitability across vulnerability types. The analysis ended when we could not derive more categories of reasons from the open codes. The axial coding resulted in 11 reasons for the low exploitability explained in Section 4.5. After that, we performed selective coding to connect reasons identified through axial coding to generate coherent explanatory schemes, i.e., the theories.

3.1.5 Step 5. Use other detectors to verify exploitability analysis results

As explained in Section 3.1.2, after we have identified the vulnerable contracts, we want to use other detectors, which are relatively slow but effective due to the use of dynamic approaches, to verify our manually-generated findings of smart contracts' exploitability. For the reported vulnerabilities that are found to be exploitable, we want to use these detectors to check if the identified vulnerable lines are reachable and triggerable. For the vulnerability found

to be unexploitable, we want to investigate if our identified reasons for the low exploitability are still valid with these detectors.

3.2 Research Design to Answer RQ2

As shown in Figure 1, the steps to answer RQ2 are as follows.

3.2.1 Step 1. Collect transaction logs

For all the cross-reported vulnerable contracts, we retrieve their transaction logs on Ethereum through the debug function of EVM, which supports replaying transactions and tracing transaction logs. The EVM's debug function is accessed through the Remote Procedure Call (RPC) provided by the Ethereum client.

3.2.2 Step 2. Analyze transaction logs to identify vulnerability exploitation

Step 4 to answer RQ1 identifies several exploitable contract. For the vulnerable contracts that we label as unexploitable, we analyze their transaction logs to check if they are exploited. The purpose is to verify that our low exploitability analysis is correct. We expect that there shall have no exploitation in the transaction logs of the unexploitable contracts.

Step 4 to answer RQ1 also identifies exploitable smart contracts. We analyze these smart contracts' transactions to determine if the exploitable vulnerabilities have been exploited. We developed different detectors for each vulnerability type to analyze the vulnerability exploitation.

3.2.3 Step 3. Identify reasons for vulnerability exploitation

Step 2 finds exploited contract on Ethereum's main network. Nevertheless, there are many exploitable smart contracts that are not exploited. We, again, use Strauss' grounded theory approach [25] to discover the possible reasons for this phenomenon. Besides the transaction logs, the extra data we analyze include the smart contracts' account types and balances. After the open coding and axial coding similar to what we did to answer RQ1, we derived several possible reasons, shown in Section 5.3. From the axial coding results, we performed selective coding to schemes, which will also be extensively recounted in Section 5.3.

4 Results of RQ1

4.1 Collected Unique Smart Contracts

We crawled 2,217,692 smart contracts from Etherscan. From these contracts, 2,080,723 duplications were found, giving a duplication percentage of 93.82%. After duplication filtering, we got 136,696 unique smart contracts with 318,026,937 transactions (*before 2022.6.1, UTC+2 08:23:22*). Figure 2 shows the transaction information of these contracts and indicates that 88.37% of contracts have more than one transaction. Figure 2 also shows that the contracts have broad coverage of different numbers of transactions. Thus, we believe the chosen smart contracts are representative.

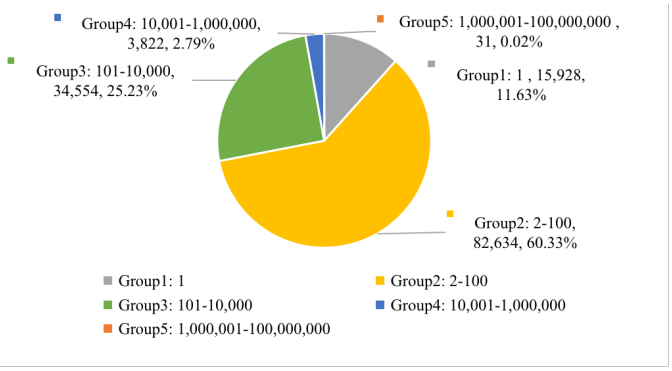

Fig. 2: Distribution of the number of transactions of the chosen smart contracts

4.2 Selected Fast Vulnerability Detectors

Based on the detectors selection criteria in Section 3.1.2, we choose to study four detectors, namely, *Oyente* [2], *Smartcheck* [13], *Slither* [14], *SoliDetector* [24]. The chosen detectors cover pattern matching, symbolic execution, and data flow analysis approaches. The other detectors are excluded because they do not satisfy one or multiple selection criteria.

- *Source Code Input (SCI)*: *TEETHER* [21], *HONEYBADGER* [9], *Osiris* [18], *MadMax* [22], *Ethainter* [29], *VSCL* [32] are excluded because they do not take source code as input. Even though the framework *SmartBug* [40] supports using *HONEYBADGER* [9] and *Osiris* [18] with the source code as input, we exclude them because they often report compilation errors, such as "Solc experienced a fatal error", when detecting real-world contracts and return null results. For example, when we run *HONEYBADGER* [9] and *Osiris* [18] with 100 real-world smart contracts, they reported 76 and 68 compilation errors respectively.
- *Vulnerability Localization (VL)*: *MAIAN* [17], GNN-based detector [30], *VSCL* [32], and *ETHPLOIT* [34] do not provide location information of the identified vulnerability and are, therefore, excluded.
- *Solidity Versions (SV)*: *Securify 2.0* [45] and *DefectChecker* [7] are excluded because they only support limited versions of Solidity. *DefectChecker* [7] aims at Solidity version 0.4.25, which is the most widely used version at the time of developing this tool. The updated tool *Securify 2.0* [45] only supports contracts written in Solidity after its version 0.5.8.
- *Availability*: We cannot get access to the source or executable code of three detectors [8], [20], [28], although we have contacted the paper authors and asked for the code.

A more structured summary of the reasons for excluding the detectors is shown in Table 8 in Appendix.

4.3 Chosen Vulnerability Types

To choose vulnerability types supported by at least two detectors, we did a mapping of the types between the detectors and decided to focus on ten types of vulnerability, as shown in Table 3. It is worth noting that the vulnerability

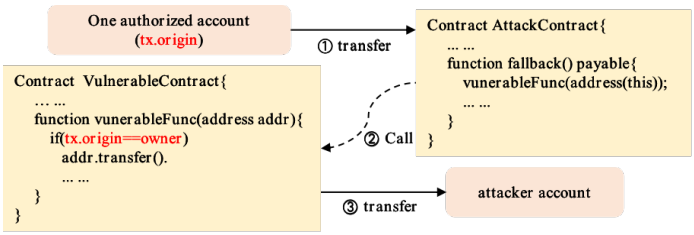

Fig. 3: An attack process exploiting TO

names in Table 3 are extracted from the detection results of the detectors, which may be different from the names in the papers presenting the detectors because the authors of the detectors did not make the names consistent.

According to state-of-the-art books and literature, e.g., [46] and [47], the definitions and characteristics of the ten chosen vulnerability types are as follows.

4.3.1 Unprotected Suicide (UpS)

The *selfdestruct(address)* function can remove all bytecode from the contract address and sends all ether stored in this contract to the *address*. If a contract is vulnerable to UpS, attackers can self-destruct the contract and transfer all contract balances to an attacker-specified *address*. According to [46], a contract vulnerable to the UpS attack has the following characteristics.

1) A function containing the *selfdestruct(address)* function.
2) No access control prevents attackers from calling the *selfdestruct(address)* function to destroy the contract.

4.3.2 TxOrigin (TO)

In Solidity, *tx.origin* returns the address of the originating Externally Owned Account (EOA) of a transaction [47]. Using *tx.origin* for authorization could make a contract vulnerable if an authorized user calls into a malicious contract. An example attack exploiting the TO vulnerability is shown in Figure 3. The attacker first lures the owner of *VulnerableContract* to transfer ether to the *AttackContract*. After that, the *tx.origin* of this transaction is the *owner* of the *VulnerableContract*. Once the *AttackContract* receives ether, the fallback function of the *AttackContract* will be triggered. A contract usually has one fallback function. The fallback function is executed on a call to the contract if none of the other functions match the given function signature or if no data was supplied and there is no receive ether function [48]. The fallback function can be declared using *function()*, *fallback()*, or *receive()* in different Solidity versions. Because the *tx.origin* of this transaction is the *owner* of the *VulnerableContract*, the call from the *fallback* function to the *VulnerableContract* can pass the authorization check *if(tx.origin==owner)* and execute the *addr.transfer()* function, which will cause unexpected ether transfer. According to [47], a contract with TO vulnerability has the following characteristics.

1) *tx.origin* is used for authorization in a function.
2) There are critical operations after successful *tx.origin* authorization.

4.3.3 Arithmetic Overflow and Underflow

An arithmetic overflow or underflow [18], [49], which is often also called Integer Overflow or Underflow (IOU), occurs when an arithmetic operation attempts to create a numeric variable value that is larger than the maximum value or smaller than the minimum value of the variable type. The popular IOU preventative technique is to use secure mathematical libraries, i.e., *SafeMath*, to replace the standard math operators, i.e., addition, subtraction, and multiplication. Thus, the arithmetic overflow or underflow may happen if a smart contract meets the following characteristic [46].

1) The arithmetic operation may pass a variable type's maximum or minimum value. However, the arithmetic operation is performed without using *SafeMath*.

4.3.4 DelegateCall (DC)

The function *address.delegatecall* allows a smart contract to dynamically load code from the target contract (*address*) at runtime. The code executed at the targeted address runs in the context of the calling contract. Calling into untrusted contracts can be dangerous. The code at the target address can change storage values of the calling contract, e.g., to change the caller's contract balance [50], [51], because state-preserving of *delegatecall* refers to the storage slots rather than the variable name. According to [50], [51], [52], a contract vulnerable to the DC attack usually has the following characteristics.

1) A function containing the *delegatecall* function.
2) No access control on the function prevents the attacker from specifying the calldata or changing the target contract address.

4.3.5 Unchecked Call (UcC)

If a smart contract does not check the return value of a message call and assumes that the call is always successful, the failing of the call may lead to inconsistency between the logic of the program and the system state [22], [46], [53]. The functions *address.call()* and *address.send()* are often used to transfer ether, and they return a Boolean value indicating whether the call succeeds. The transaction that executes these functions may return a false value but will not revert if the external call fails. So, a smart contract with the UcC vulnerability has the following characteristic [22], [46], [53].

1) The functions *address.call()* or *address.send()* is used without result checking.

4.3.6 Reentrancy (RE)

In Ethereum, insecure use of *call()* function can lead to reentrancy attacks. In the reentrancy attack, a malicious contract calls back into the vulnerable contract before the first invocation of the vulnerable function is finished. If the state variable change is after the *call()* function, the unexpected reentrancy into the vulnerable contract will result in program execution and state variable change inconsistency. Figure 4 shows an attack process exploiting the RE vulnerability. An attacker creates the *AttackContract* to call the *VulnerableContract* to transfer ether the attacker. In the *AttackContract*, there is a *fallback* function. Once *AttackContract*

TABLE 3: Mapping of the different vulnerabilities analyzed

Vul.	Oyente	SmartCheck	Slither	SoliDetector
UpS	-	-	Suicidal	Unprotected Suicide
TO	-	Tx_Origin	Tx-Origin	TxOrigin
IOU	Integer Overflow/Underflow	-	-	Integer Overflow and Underflow
DC	-	-	Controlled-Delegatecall	DelegateCall
UcC	Callstack Depth Attack	Unchecked_Call	Unchecked-Send	Unchecked Send
RE	Re-Entrancy	-	Reentrancy	Reentrancy
FE	-	Locked_Money	Locked-Ether	Frozen Ether
NC	-	Transfer_in_Loop	Calls-Loop, Costly-Loop	Nested Call
TD	Timestamp Dependency	-	Timestamp	Dependency of timestamp
TOD	Transaction-Ordering Dependency	-	-	Transaction Order Dependency

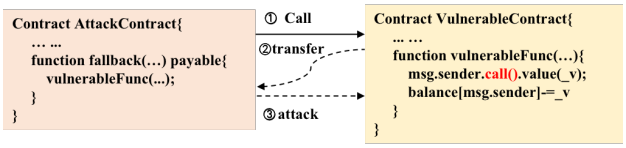

Fig. 4: An attack exploiting RE

receives the ether, the *fallback* function will be triggered to call back into the *VulnerableContract* to perform the attacks, e.g., transfer more ether to the attacker's account before changing the account's balance.

1) A function transfers ether to another contract using the *call()* function.
2) The state variable change is after the *call()* function.

4.3.7 Frozen Ether (FE)

A contract vulnerable to FE can receive ether but does not contain any functionalities to transfer ether. It relies on other contracts to transfer ether. However, if the contracts to be called to transfer ether are accidentally or intentionally terminated, the ether cannot be transferred from the contract and will be frozen. A contact with FE vulnerability has the following characteristic.

1) The contract can receive ether but cannot transfer ether by itself.

4.3.8 Nested Call (NC)

If a loop contains the gas-costly instruction but does not limit the loop iterations, the function containing the loop has a high risk of exceeding its gas limitation and causing an out-of-gas error [7]. An example of the gas-costly instruction is a non-zero value transfer as part of the CALL operation, which costs 900 gas [7]. According to [7], a contract vulnerable to NC attack has the following characteristics.

1) In the contract, dynamic data structures (e.g., array or mapping) or variables in the loop condition control the number of loop iterations.
2) The loop body contains gas-costly instructions, e.g., CALL operation.
3) No access control prevents an attacker from controlling the dynamic data structures or variables in the loop condition.

4.3.9 Timestamp Dependency (TD)

When mining a block, a miner has to set the timestamp for the block with the miner's local system time. The miner can vary this timestamp value by roughly 900 seconds while still having other miners accept the block [2]. Suppose the timestamp is used as a triggering condition to execute some critical operations [2], e.g., sending ether. In that case, miners can be incentivized to choose a timestamp that favors themselves. Thus, a contract vulnerable to TD has the following characteristic [2].

1) The contract uses the *timestamp* as the deciding factor for some critical operations, e.g., sending ether.

4.3.10 Transaction Order Dependency (TOD)

Miners decide the transaction order because transactions in the blockchain need to be packaged by miners before they are finally recorded on the chain. Malicious contract owners or attackers can exploit such order dependency. For example, if the contract is a game [19], which gives participants who submit a correct solution to a puzzle reward, a malicious contract owner could reduce the reward amount after the solution transaction is submitted. An attacker can watch the transaction pool and steal the correct answer. Then, he creates a transaction with the correct answer and gives a higher gas to get his answer packed in a block before the transaction of the answer provider is packed [47]. As there are many variations of TOD attacks, finding a precise characteristic of smart contracts vulnerable to TOD is challenging. According to [47], a high-level characteristic of a smart contract vulnerable to TOD is as follows.

1) The contract may send out ether differently according to different values of a global state variable or different balance values of the contract.

4.4 Vulnerability Reported by Chosen Detectors

The results of analyzing the 136,969 smart contracts focusing on the ten vulnerability types are shown in Table 4. The data in the *Overlap* column of Table 4 show the contracts flagged as vulnerable by multiple detectors. Results in Table 4 show that the detectors reported a large number of IOU and FE-type vulnerabilities. As we plan to use Strauss' grounded theory approach to manually analyze each reported vulnerability, analyzing all the reported IOU and FE-type vulnerabilities will take an enormous time. Hence, we randomly select 100 contracts for the IOU and FE-type vulnerabilities to analyze, as shown in the *Selected Contracts* column in Table 4. In total, we analyzed 4,364 contracts reported as vulnerable.

TABLE 4: Detection Results of Fast Vulnerability Detectors

Vul.	Oyente	SmartCheck	Slither	SoliDetector	Overlap	Selected Contracts
UpS	-	-	218	1,046	137	137
TO	-	2,292	1,807	45	45	45
IOU	65,829	-	-	80,121	28,457	**100**
DC	-	-	1,186	24,227	924	924
UcC	940	2,0361	1,683	1,316	219	219
RE	314	-	31,287	2,031	97	97
FE	-	27,882	10,240	17,901	2,934	**100**
NC	-	807	15,702	33,393	473	473
TOD	4,298	-	-	8,814	913	913
TD	4,298	-	29,941	28,365	1,356	1,356

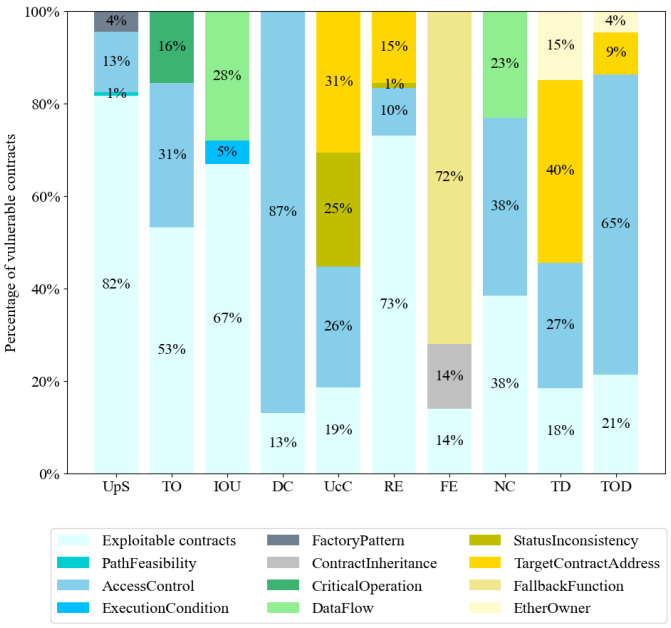

Fig. 5: The reasons for the low exploitability of reported smart contract vulnerabilities

4.5 Results of Analyzing the Reported Vulnerabilities

After analyzing the 4,364 vulnerable contracts, we identify 11 reasons for the low exploitability and show them in Figure 5. The reasons can be further abstracted into three schemes.

Scheme 1: Weaknesses of the detector in adapting approaches to analyze OO-based applications result in reporting vulnerabilities that are not reachable and triggerable.

Reasons related to Scheme 1 are explained in detail in Sections 4.5.1 to 4.5.3.

4.5.1 *Missing path feasibility analysis (PathFeasibility)*

SoliDetector [24] and *Slither* [14] overlook the path feasibility analysis and assume all code paths are reachable. One contract reported as vulnerable to UpS is unexploitable because the vulnerable function *selfdestruct* locates in an infeasible path and will never be executed. The condition to execute the vulnerable *selfdestruct* is *require(cancel == 1)*. However, the initial value of *cancel* is 0, and no arithmetic operation changes this state variable value to be one.

4.5.2 *Overlooking preventive execution condition (ExecutionCondition)*

Vulnerability detectors reported the vulnerability that could not be triggered due to the preventive execution condition. For instance, in Listing 1, *balances[_from]* is greater than *_value* is checked before the arithmetic operation *balances[_from] -= _value* to avoid underflow changes of the storage value of *balances[_from]*.

```
function transferFrom(address _from, address _to, uint256
    _value) {
    if (balances[_from] >= _value ) {
        balances[_from] -= _value;
        Transfer(_from, _to, _value);
    }
}
```

Listing 1: A code with IOU vulnerability but is unexploitable

4.5.3 *Insufficient data flow analysis (DataFlow)*

Data flow analysis is essential to detect variable-related vulnerabilities, such as IOU and NC. By analyzing the 100 smarts reported as vulnerable to IOU, we found cases where variables involved in the arithmetic operation are fixed values or have a fixed range. The arithmetic operations on these variables will not trigger overflow or underflow due to data flow control within a function or across functions.

Listing 2 shows an example of a vulnerable code with IOU that is unexploitable caused by the data flow control within one function. In Listing 2, the arithmetic operation in line 2 is labeled as vulnerable to IOU. However, line 1 shows that the variable *allCards* has at least one element. Therefore, the length of *allCards* is always bigger than 1, which will not cause arithmetic underflow.

```
1 allCards.push(Card(ids[i],0,CardStatus.Tradable,upIndex));
2 idToCardIndex[ids[i]] = allCards.length - 1;
3 cardToOwer[ids[i]] = _address;
4 ownerCardCount[_address] = ownerCardCount[_address].add(1);
```

Listing 2: A code with IOU vulnerability but is unexploitable

An example of data flow control preventing the reported NC vulnerability from being exploitable is shown in Listing 3. In Listing 3, the *session* is a struct which has multiple properties, including *investor*, *investorCount*, and *amountInvest*. An attacker can add a new element into the *session* by calling the function *invest*. However, the value of the property *investorCount* cannot be greater than the value of the global variable *MaxInvestor*, which is pre-defined as 20 in line 1. Therefore, the total gas cost of the function *closeSession* will not be more than the given gas. Thus, the reported NC vulnerabilities will not lead to the permanent failure of the function *closeSession*.

```
1  uint public constant MaxInvestor = 20;
2  function closeSession (uint _priceClose) public onlyEscrow{
3      for (uint i = 0; i < session.investorCount; i++) {...}
4      session.investorCount = 0;
5  }
6  function invest (bool _choose) public payable{
7      require(msg.value >= minimunEth && session.investOpen);
8      require(session.investorCount < MaxInvestor);
9      session.investor[session.investorCount]=msg.sender;
10     session.amountInvest[session.investorCount]=msg.value;
11     session.investorCount+= 1;
12 }
```

Listing 3: A code with NC vulnerability but is unexploitable

Scheme 2: Overlooking the characteristics of the Solidity programming language results in reporting unexploitable vulnerabilities.

Gavin Wood designed Solidity to support condition-oriented programming [54], a subdomain of contract-orientated programming with the principle to "Never mix transitions with conditions." When developing smart contracts, it is common to set complicated conditions to be satisfied to allow the execution of the transitions. Without a lot more comprehensive analyses of smart contract conditions, the detectors will report the transitions as vulnerable, even if the conditions can prevent them from being exploitable. Thus, these reported vulnerabilities are unexploitable. Section 4.5.4 presents several examples of methods in Solidity language that can defend attackers against exploiting the vulnerability. Although Solidity supports a few OO programming language principles, such as inheritance, its implementation of the OO features can differ. Two reasons for low exploitability explained in Sections 4.5.5 to 4.5.7 can also be coded to Scheme 2.

4.5.4 Overlooking access control (AccessControl)

Eight vulnerability types can be protected by access control, including UpS, TO, DC, UcC, RE, NC, TD, and TOD. Examples of extra access control checking the identities of critical functions' caller or controlling a critical variable are as follows.

① The *if/require* condition is set, such as *require(msg.sender == owner)*, before critical operations are called. For example, the access control on the function containing the UpS vulnerability prevents the attacker from exploiting it. In Listing 4, the function *closeStableCoin* checks the caller's identity using the *require* condition in line 2 before executing *selfdestruct*.

```
1 function closeStableCoin() public {
2     require(whitelist.isSuperAdmin(msg.sender), "Only
       SuperAdmin can destroy Contract");
3     selfdestruct(msg.sender); // admin is the admin address
4 }
```

Listing 4: A code with UpS vulnerability but is unexploitable

Another example is that: a smart contract labeled vulnerable to TO if *tx.origin* is used for authorization. The TO vulnerability is unexploitable because multiple-level access control defends the vulnerability from being exploited. The developer not only uses *tx.origin* for authorization but also checks the identity of the *msg.sender* and *recipient*, such as *require(owner == tx.origin && msg.sender == tx.origin, "Not token owner")*, which can reject external contracts calling the current contract and defend against the TO attack.

② In Solidity, modifiers are used to modify the behavior of a function. A modifier usually contains code, e.g., code to check the user's identity, and a special symbol "_". When executing the function claimed using the modifier, the functions' code will be inserted at the location of the symbol "_" in the modifier. If the modifier's code to check the user's identity is located before "_", the functions' code inserted will be protected by the identity checking. Otherwise, the functions' code can be called by any user. As shown in Listing 5, even though *tx.origin* is used for authorization, the modifier *onlyMain* checks the identity of *msg.sender* before executing the function *addBrick*, which prevents intermediate contracts from being used to call the current contract [47]. In 13 RE vulnerabilities, the code of the modifier checks whether the *msg.sender* is *tx.origin* or *owner*, which makes the recursive call from another contract impossible.

```
1 modifier onlyMain() {  require(msg.sender == main); _;}
2 function addBrick(uint _value) external onlyMain returns (
      bool success){
3     require(_value >= 10 ** 16);
4     require(owner == tx.origin);
5     return true;
6 }
```

Listing 5: A code with TO vulnerability but is unexploitable

③ Across control is performed across functions. Some NC vulnerabilities are unexploitable due to access control across multiple functions or modifiers. In Listing 6, the loop in line 7 is labeled with the NC vulnerability, and the max number of loop iterations is equal to the length of variable *landmarks* that the function *totalSupply* can access. The variable *landmarks* is global and can be modified through the function *createLandmark*. However, the function *createLandmark* is modified by the modifier *onlyCOO*, which requires that the caller is the authenticated user *coo*. Therefore, the attacker cannot arbitrarily increase the elements in the variable *landmarks* to exploit the NC vulnerability in function *buy*.

```
1  uint256[] private landmarks;
2  function totalSupply() public view returns (uint256) {
3      return landmarks.length;
4  }
5  function buy(uint256 _tokenId) public payable {
6      require(msg.sender != address(0));
7      for (uint i = 0; i < totalSupply(); i++) {
8          uint id = landmarks[i];
9          landmarkToOwner[id].transfer(feeGroupMember);
10     }
11 }
12 modifier onlyCOO() {  require(msg.sender == coo); _; }
13 function createLandmark(uint256 _tokenId) public onlyCOO {
14     ... ...
15     landmarks.push(_tokenId);
16 }
```

Listing 6: A code with NC vulnerability but is unexploitable

4.5.5 Neglecting constraints caused by factory patterns (FactoryPattern)

Factory pattern is one of the most used design patterns in Java. "In the factory pattern, instead of directly creating instances of objects, a single object (the factory) does it for you" [55]. Solidity supports the factory pattern, and smart contracts are the objects. A factory in Solidity is a contract (called **main contract**) that can deploy multiple instances of other contracts (called **template contracts** in this paper) at runtime. In Listing 7, *SwapperFactory* is a main contract that creates the template contract objects multiple times and destruct the objects by calling the function *destroy*.

When using the factor pattern, the *selfdestruct* in the instances created from the template contract cannot destruct the main contract *SwapperFactory*. The vulnerability detectors report six contracts vulnerable to UpS with the factory pattern constraint as shown in Listing 7. The *selfdestruct* in the reported vulnerable contract cannot be exploited by attackers.

```
contract SwapperFactory {//Main contract
  function performSwap(address payable user){
     Swapper swapper = createClone(user, srcToken, dstToken
     , uniqueId);
     swapper.destroy(user);
  }
  function createClone( ) private onlyAdmin() returns (
     Swapper) {... ...}
}
contract Swapper {  //Template contract
  function destroy(address payable user) external {
    selfdestruct(user);
  }
}
```

Listing 7: A code with UpS vulnerability but is unexploitable

4.5.6 Neglecting constraints caused by contract inheritance (ContractInheritance)

Solidity supports inheritance between smart contracts. A contract can inherit multiple contracts. The contract from which other contracts inherit is called a base contract, while the contract which inherits the features of the base contracts is called a derived contract. "With the inheritance construct, the derived contract inherits all the methods, functionality, and variables of the base contract and can extend a base contract with additional functionality" [47].

If the contract can receive ether but cannot transfer it by itself, the vulnerability detectors tag it vulnerable to FE. Fourteen contracts tagged as vulnerable to FE are base contracts containing no transfer operation. These base contracts do not have an account on Ethereum's main network and do not own ether. Other contracts inherit these base contracts and implement ether transferring. Therefore, these base contracts will not lock the ether, meaning the FE vulnerability is unexploitable.

4.5.7 Assuming all fallback functions receive ether (FallbackFunction)

66 of the 100 contracts tagged as vulnerable to FE will never lock ether because their fallback functions cannot or refuse to accept ether. To receive ether, the fallback function must be marked *payable*, as shown in line 1 in Listing 8. The contract cannot receive ether if it does not contain the fallback function or if the fallback function is not marked as *payable*, as shown in line 2 in Listing 8. The approach to refuse ether is to insert the *revert()* into the fallback function as shown in lines 3 and 4 in Listing 8. Once the contract receives the ether, this transaction will revert. Therefore, any transaction transferring ether to these contracts will fail, and these contracts will not receive any ether. Therefore, the FE vulnerability cannot be exploited to cause the ether lock.

```
function() payable public{ }//Accept ETH
function() public{ }//Don't accept ETH
function () payable public{ revert(); }//Don't accept ETH
receive() external payable { revert(); }//Don't accept ETH
```

Listing 8: Fallback functions

Scheme 3: Smart contract application scenarios reduce exploitability.

The application scenarios of the smart contracts are to transfer assets or ether between suppliers and clients. Without being able to manipulate or sabotage the asset transfer maliciously, the likelihood of security compromise or giving benefits to attackers by executing the code is low [56]. The reasons for low exploitability associated with Scheme 3 are presented in Sections 4.5.8 to 4.5.11.

4.5.8 Insufficient analysis of the values of the target contracts' addresses (TargetContractAddress)

Calling an external contract or transferring ether to an address can be dangerous if the attacker controls the contract or the target address. However, in some cases, the target contracts' addresses are hard-coded, fixed, or under the complete control of the contract owner. The specific target contracts' addresses can prevent the attacker from exploiting the RE, TD, and TOD vulnerabilities.

In 12 contracts that are reported as vulnerable to RE, the target contracts' addresses are hard-coded. The hard-coded address may be a global variable used in multiple functions. Defining the target address as *immutable* can also freeze the value of the addresses. Therefore, an unexpected call from the *fallback* function in the target contract will not happen. For TD and TOD, if the recipient address is fixed or fully controlled by the contract owner, the attack will not get profit by attacking the vulnerability even if an attacker can manipulate the timestamp or determine the order of function calls and transactions.

4.5.9 Omitting the case that the ether transfer initiator is the ether's initial owner (EtherOwner)

To detect TOD vulnerability, *Oyente* and *SoliDetector* focus on the ether flow because TOD may lead to undesirable outcomes when dealing with ether [2]. *Oyente* labels a contract as vulnerable to TOD if it sends out ether differently when the order of transactions changes. *Oyente*, *Slither*, and *SoliDetector* label the contract as vulnerable to TD if the block *timestamp* is used as the condition to send ether. However, *Oyente*, *Slither*, and *SoliDetector* all ignore the scenario that the ethers transferred to a user after the timestamp checking may come from the user himself. For example, many contracts vulnerable to TD or TOD are wallet contracts that support users to purchase or withdraw tokens within the specified time range. If the time range passes, the ether to purchase the token will be returned to the user. Such ether transfer after the timestamp checking is not harmful because the ether is returned to its initial owner. Listing 9 shows an example, in which users can send a message with *msg.value* to the contract *OpportyPresale* to purchase token. *OpportyPresale* contains a fallback function labeled as vulnerable to TD, and the vulnerability is in line 5. Once this contract receives ether, the fallback function will be triggered to verify whether the transaction meets the conditions regarding timestamp (*now >endDate)* and amount of ether (*msg.value >= 0.3 ether*). The contract will return the ether to the token purchaser if the transaction is not within the valid time. As a result, the ether is returned to its initial owner.

```
contract OpportyPresale is Pausable {
    function() whenNotPaused public payable {
        require(msg.value >= 0.3 ether);
        require(whiteList[msg.sender].isActive);
        if (now > endDate) {
            state = SaleState.ENDED;
            msg.sender.transfer(msg.value);
            return ;
        }
    }
}
```

Listing 9: A code with TD vulnerability but is unexploitable

4.5.10 Assuming critical operations after authorization (CriticalOperation)

One main characteristic of TO is having critical operations, e.g., sending ether, after successful authorization. If there is no critical operation following successful authentication, the TO vulnerability is not risky. Seven contracts vulnerable to TO are unexploitable because successfully bypassing the authentication will not bring security risks. The example vulnerable code is : *if(tx.origin == owner()) return;*.

4.5.11 Assuming status inconsistency when function call results are not checked (StatusInconsistency)

If there is no status change after calling the functions *send()* and *call()*, it is not risky even though the result of the message call is not checked. There is no status change following the message call in 14 contracts vulnerable to UcC. Listing 10 shows an example of such vulnerable codes. The function *executeCall()* transfers ether by the *call()* function in line 4. The variable *underExecution* is set to avoid the recursive calling from the *_target*. The initial value of *underExecution* is *false* (line 2). It will change to *true* (line 3) before executing transferring by *call()* and turn back to *false* after transferring (line 5). Therefore, no status change follows the execution of the function *call()* because the value of *underExecution* is always false regardless the function call *call()* fails or not. The failed call in line 4 does not cause any compromise.

```
function executeCall()external onlyAllowedManager(){
    require(underExecution == false);
    underExecution = true; // Avoid recursive calling
    _target.call.gas(_suppliedGas).value(_ethValue)(
     _transactionBytecode);
    underExecution = false;
  }
```

Listing 10: A code with UcC vulnerability but is unexploitable

4.6 Results of Verification Using Other Detectors

To select other detectors to verify exploitation analysis results, we first excluded unavailable fuzzing detectors, i.e., *Reguard* [33], *ContraMaster* [6], and *Ethploit* [34]. Then, we excluded the *ContractFuzzer* [5] and *sFuzz* [4] because the empirical evaluation [10] on nine detectors [2], [3], [4], [5], [12], [13], [14], [18], [43] demonstrates that *Mythril* [12] outperforms other fuzzing detectors. In addition to using *Mythril* [12], we choose to use two latest fuzzing detectors, i.e., *ConFuzzius* [26] and *Smartian* [27] for the verification.

Compared to the fast detectors, *Mythril* [12], *ConFuzzius* [26], and *Smartian* [27] apply more dynamic code analysis approaches and are good at identifying vulnerabilities in a runtime environment. For the exploitable smart contracts resulting from our manual analysis, we run these detectors to see if the exploitable vulnerabilities are reachable and triggerable. Results are in Table 5 and show that 56.72% of the 1,079 exploitable contracts are executable by *Mythril* [12], *ConFuzzius* [26], or *Smartian* [27]. Specifically for the RE vulnerability, 59 vulnerable contracts reported by Smartian are all exploitable. Such results give support to the conclusions of our manual analysis. However, *Mythril* [12], *ConFuzzius* [26], and*Smartian* cannot guarantee to explore all paths of the code. In addition, they may encounter execution errors due to missing external contract dependencies we cannot resolve. Therefore, they cannot verify all exploitable contracts we identify.

TABLE 5: Detection Results of *Mythril*, *ConFuzzius*, and *Smartian* on Exploitable Smart Contracts (ExploitableSC)

Vul.	Nr. of ExploitableSC	Mythril	ConFuzzius	Smartian	Total	P(%)
UpS	112	50	49	32	72	64.29
TO	24	8	-	20	20	83.33
IOU	63	1	-	17	18	28.57
DC	122	-	-	-	-	-
UcC	41	15	21	33	37	90.24
RE	71	54	22	59	67	94.36
FE	14	-	1	0	1	7.14
NC	182	45	-	85	94	51.65
TD	250	84	94	118	187	74.80
TOD	196	-	116	-	116	59.18
Total	1079	257	303	364	612	56.72

Note: The "Total" column shows the number of unique vulnerable smart contracts detected by *Mythril*, *ConFuzzius*, and *Smartian*. The "-" indicates that the tool does not support detection for the vulnerability type; the "P(%)" indicates the percentage of contracts detected out of all exploitable smart contacts.

As shown in Section 4.5, our manual exploitability analyses illustrate that *Oyente* [2], *SmartCheck* [13], *Slither* [14], and *SoliDetector* [24] are weak at analyzing infeasible path, preventative execution condition, and data flow control to identify unexploitable smart contracts. *Mythril* [12], *ConFuzzius* [26], and *Smartian* use more dynamic approaches and are usually better at covering complex paths. For the vulnerabilities that are labeled as unexploitable, we apply *Mythril* [12], *ConFuzzius* [26], and *Smartian* [27] to see if they can exploit them. Results are in Table 6 and show that those dynamic detectors are better at dealing with infeasible path, preventative execution condition, and data flow control than fast detectors. For example, a contract vulnerable to UpS with an infeasible path is found to be exploitable.

However, *Mythril* [12], *ConFuzzius* [26], and *Smartian* are not perfect. Similar to the fast detectors, they are also weak at differentiating unexploitable vulnerabilities related to other reasons listed in Section 4.5. For instance, they all ignore the access control when detecting the UpS vulnerability. Listing 11 shows the code reported as vulnerable by *Mythril*, *ConFuzzius*, and *Smartian*. The *selfdestruct* in line 6 is labeled as vulnerable to UpS. However, an execution condition of *selfdestruct* is "tx.origin == O". The address *O* is the creator of this contract. Only the creator of the contract can destroy the contract, and all balances of the contract will be sent to the creator. The attacker can neither destroy the contract nor get ether. Therefore, the UpS vulnerability in line 6 is unexploitable.

TABLE 6: Detection Results of *Mythril*, *ConFuzzius*, and *Smartian* on Unexploitable Smart Contracts (UnExploitableSC)

Vul.	Nr. of UnExploitableSC	Mythril	ConFuzzius	Smartian
UpS	25	7	3	1
TO	21	3	-	6
IOU	37	0	-	0
DC	802	-	-	-
UcC	178	43	62	103
RE	26	3	3	0
FE	86	-	0	0
NC	291	33	-	73
TD	1106	390	336	562
TOD	717	-	72	-

```
contract GrungeTuesday{
    address O = tx.origin;
    function() public payable {}
    function multi_x() public payable {
        if (msg.value >= this.balance || tx.origin == O) {
            selfdestruct(tx.origin);
        }
    }
 }
```

Listing 11: A code with UpS vulnerability but is unexploitable

5 Results of RQ2

As mentioned in Section 3.2, we collect the transaction logs of the contracts that are reported as vulnerable and analyze them. The transaction log consists of several log blocks. Each block reflects the running state of EVM, in which:

- *pc* is the program counter.
- *op* represents a low-level machine language consisting of a series of instructions, each of them representing an operation.
- *gas* represents the remaining gas.
- *gasCost* refers to the gas consumption of the current opcode.
- *depth* of call stack indicates the depth of nested calls, which has a maximum value of 1024.
- *stack* is an internal place where temporary variables, such as local variables, intermediate calculation results, and return addresses, are stored.
- *memory* is a temporary place to store data, of which a contract obtains a freshly cleared instance for each message call [48].
- *storage* is a key-value store. Data in storage are stored permanently between function calls and transactions. For instance, the global variables declared in the smart contract are stored in the *storage*.

5.1 Analyzed Transaction Logs

In the analysis, we excluded transaction logs of TD and TOD vulnerabilities because these two types of vulnerabilities exploit the mining process. Therefore, information in the transaction log cannot reflect the exploitation. The 219 contracts vulnerable to UcC have 5,450,975 transactions, too many for us to replay and perform the ground theory analysis. We randomly selected 100 UcC contracts reported as vulnerable, which have 145,469 transactions. Of the 219 contracts, one contract had the largest number of transactions at 4,911,428, accounting for ninety percent of all transactions (5,450,975). Because we selected 100 contracts randomly, the contract that had the largest transaction number was excluded. Thus, the selected 100 UcC contracts only have 145,469 transactions. The numbers of transactions analyzed for the eight vulnerability types are as follows: UpS (33,920), TO (77,934), IOU (131,643), DC (1,683,694), UcC (145,469), RE (5,362), FE (396,127), and NC (1,631,985).

We designed analysis rules as shown in Figure 6 to analyze vulnerability exploitation. The analysis rules contain *opcode*, information to search in *stack* or *storage*, and additional constraints. Each of the rules is explained as follows.

5.1.1 Unprotected Suicide (UpS)

The EVM opcode SELFDESTRUCT destroys contracts. The SLEFDESTRUCT opcode used to be called SUICIDE, but SUICIDE was deprecated due to the negative associations of the word [47]. It is insecure if the attacker exploits SELFDESTRUCT. However, it is challenging to identify whether a user is malicious. In this study, we define the contract's creator as benign and assume any other users who destroy the contract they do not own are malicious.

5.1.2 TxOrigin (TO)

Attacks exploiting the TO vulnerability usually follow the attack process shown in Figure 3. We designed the TO analysis rule according to that attack process. To find a call from the *fallback* function in the attacker's contracts, we search the CALL instruction ①, in which the target address is the vulnerable contract address (1). Then, we look for the ORIGIN instruction ②, in which the origin address (2) is usually used for authorization. Finally, we identify the EQ instruction ③ for the authorization and expect the authorization result to succeed (3).

5.1.3 Arithmetic Overflow and Underflow (IOU)

The arithmetic overflow is usually caused by arithmetic instructions ADD or MUL, and the arithmetic underflow may happen when executing the SUB instruction. We first find the log block containing the arithmetic instruction ADD, MUL, or SUB ① and get two operands from the stack to calculate the expected result (1) of the arithmetic operation. Then we compare the actual value (2) with the expected value. If the actual value that has been pushed onto the stack is not equivalent to the expected value, it means that the arithmetic overflow or underflow has occurred. Torres et al. point out that not every overflow is considered harmful because the compiler may also introduce it for optimization purposes [26]. Thus, we only trace the overflow or underflow that has updated the blockchain state. If the error result flows into an SSTORE instruction ②, we will label the transaction as IOU exploitation.

5.1.4 DelegateCall (DC)

The *delegetacall* is insecure if the state of the called contract affects the calling contract. At the transaction log level, we recognize DC exploitation according to the storage state caused by the *delegatecall*. If a *DELEGATECALL* ① causes a

	Opcode	Stack/Storage	Additional Constraints
UpS	① SELFDESTRUCT		caller is not the contract creator
TO	① CALL② ORIGIN ③ EQ	(1) target_address=stack[-2]=contract_address (2) origin_address=stack[-1] (3) eq_result=stack[-1]=1	
IOU	① ADD/SUB/MUL ②SSTORE	(1) expected_value= ADD/SUB/MUL(stack[-1], stack[-2]) (2) storage_value=actual_value	actual_value=stack[1]!=expected_value
DC	① DELEGATECALL②SSTORE③SLOAD	(2) stack[-1]=key1, depth1, value1 (3) stack[-1]=key2, depth2, value2	depth1=depth2+1 *and* key2=key1 *and* value1=value2
UcC	① CALL ②¬ (ISZERO and JUMPI) ③SSTORE	(1) call_result=stack[-1]=0	
RE	① CALL ② CALL ③ SSTORE	(1) callee_address=stack[-2]=attacker_address (2) callee_address=stack[-2]=contract_address	
FE	①No CALL/ DELEGATECALL/ CREATER/SELFDESTRCT/SUICIDE		balance of the contract > 0
NC			error=out of gas

Fig. 6: Analysis rules of transaction logs

storage variable modification by SSTORE ② when executing an external contract, and this storage variable is used in the calling contract by SLOAD ③, we label the transaction as DC exploitation. When conducting an external call, e.g., CALL or DELEGATECALL, *depth* increases by one. To distinguish the external and the current contract call, we check if the depth value (*depth1*) of SSTORE (2) is exactly one bigger than the depth value (*depth2*) of SLOAD (3). When the *delegatecall* call ends, the *depth* turns back to the value it has had when DELEGATECALL is executed [19]. To identify the storage variable used in the calling contract but modified in the external contract, we check if the two storage variables share the same key and value at different depths. It is worth noting that the external call should be secure if the external contract address and the calldata are specified by the contracts' owners. Therefore, we must manually check if the initiator of the identified DC exploitation is the contract's owner to exclude false positive exploitation.

5.1.5 Unchecked Call (UcC)

The *send()* and *call()* functions are used to send ether and are compiled into the EVM CALL instructions. The CALL instruction results are pushed onto the stack, where 0 means failure and 1 means success. The CALL result is stored in the log block where the depth of the trace turns back to the value it has had when the CALL instruction is executed [19]. If CALL ① results in value 0 (1) and the SSTORE changes the storage status ③ without executing an opcode to check the CALL result ②, we will flag the transaction as an UcC exploitation.

5.1.6 Reentrancy (RE)

As shown in Figure 4, a reentrancy attack usually calls ether transferring functions in the vulnerable contract to trigger the malicious fallback function in the malicious contract. Therefore, an RE exploitation has at least two CALL instructions ①② in one transaction. The target address of the first CALL is the attack's contract address (1), and the target address of the second CALL is the vulnerable contract's address (2). The state in the contract is updated by executing the SSTORE ③ instruction.

5.1.7 Frozen Ether (FE)

There are several reasons for funds being locked in a contract. Perez et al. [19] focus on the case that the contract relies on an external contract to transfer ether, but the external contract does not exist any longer. In this study, the vulnerable contracts are labeled by detectors that cannot know whether the contract being relied on to transfer ether is destroyed. Therefore, if the contract balance is not 0 and the contract's transaction logs do not contain any instruction supporting transferring ①, we flag the contract as a FE exploitation. A potential issue of our analysis rule is that no ether transfer in the contract's transaction does not mean the contract cannot transfer ether. Thus, the analysis rule may report false positives of FE exploitation.

TABLE 7: Information about Vulnerability Exploitation

	UpS	TO	IOU	DC	UcC	RE	FE	NC	TD	TOD	Total
VC	137	45	100	924	219	97	100	473	1,356	913	4,364
UnExploitableSC	25	21	33	802	178	26	86	291	1,106	717	3,285
ExploitableSC	112	24	67	122	41	71	14	182	250	196	1,079
ExploitedC	19	0	15	0	2	0	2	28	-	-	67

VC: The number of contracts reported as vulnerable.
UnExploitableSC: The number of contracts labeled as unexploitable.
ExploitableSC: The number of contracts labeled as exploitable.
ExploitedC: The number of contracts that were exploited.

5.1.8 Nested Call (NC)

In Ethereum, when transactions fail due to gas shortage, the transaction log will contain error messages, such as *"error":"out of gas"*. To search NC exploitation, we first look for the transaction log containing an error tag, and the reason is "out of gas." Gas shortage can also be caused by other failed transfer transactions irrelevant to NC. These unrelated transactions usually do not call any function, and their transaction logs contain mostly only one PUSH1 instruction. To exclude unrelated transactions and reduce false positives, we manually check if the nested call is the reason for the transaction failures.

5.2 Identified Vulnerability Exploitation

If at least one contract's transaction on Ethereum's main network matches our analysis rules, we count the contract as exploited. The numbers of identified exploitations are shown in Table 7. We do not find any exploitation of contracts that we labeled as unexploitable, which confirms our manual analysis in answering RQ1. For the identified exploitable contracts, their exploitations are explained below.

Unprotected Suicide (UpS). We found 19 contracts vulnerable to UpS that have been exploited, meaning 19 contracts were destructed but not by their owners. Listing 12 shows an example of the exploited contracts. The contract is related to a game, and the function *takeAGuess()* of the contract is to let users take a guess after transferring 0.0001 ether. If the number inputted by the user equals the *winningNumber*, the function will transfer 90% of the contract's balance to the user, then kill the contract and return the remaining balance to the contract owner. The transaction details of the exploited UpS are shown in Figure 8 in Appendix A, which shows that an attacker input the number nine that satisfied the if condition in line 3 and then executed the *selfdestruct()* function successfully and got 0.0468 ether from the contract.

```
1 function takeAGuess(uint _myGuess) public payable {
2     require(msg.value == 0.0001 ether);
3     if (_myGuess == winningNumber) {
4         msg.sender.transfer((this.balance*9)/10);
5         selfdestruct(owner);
6     }
7 }
```

Listing 12: A code with UpS that was exploited

An interesting finding is that some contracts still have balances even though the contracts have been self-destructed. The reason is that the contract account will not disappear even if the contract is destroyed. The contract account can receive ether but does not support transferring ether, leading to the locked ether. An example in Figure 9 in Appendix A shows the transactions of a self-destructed contract, in which 0.033 ether are locked.

TO. The transaction log analysis does not reveal exploitation of the TO vulnerabilities.

IOU. We found 15 occurrences of arithmetic overflows. Listing 13 shows the source code of an example contract that is exploited. The transaction has an invocation of the function *transport()*, in which an arithmetic operation was conducted based on the function *addDungeonRewards()*, which calculated the reward for different *originDungeonId* in the *dungeons* without using SafeMath.

```
1  DungeonToken public dungeonTokenContract;
2  function transport() external payable {
3      // ** STORAGE UPDATE **
4      // Increment the accumulated rewards for the dungeon.
5      dungeonTokenContract.addDungeonRewards(originDungeonId,
         requiredFee);
6      ......
7  }
8  contract DungeonToken {
9      function addDungeonRewards(uint _id, uint _Rewards) {
10         dungeons[_id].rewards += uint128(_Rewards);
11     }
12 }
```

Listing 13: A code with arithmetic overflow

DC. We did not find any exploitation of the 924 contracts reported as vulnerable to DC.

UcC. We found three UcC exploitations, which contained failed calls and storage status changes after the call failures. An example of the exploited contracts is shown in Listing 14. In the contract's transaction, at a certain point in time, there was insufficient ether in the contract supporting the transfer in line 7. Therefore, the log shows that a transaction calling the function *sendTokensManager* did not call the *send()* function in line 7 successfully. However, the contract, e.g., *Exxcoin*, calling the function *sendTokensManager* still changed the storage variable *balances* in line 8 after the invocation of the *send()* function failed.

```
1  contract ExxStandart is ERC20 {
2    mapping (address => uint) balances;
3  }
4  contract Exxcoin is owned, ExxStandart {
5      function sendTokensManager(address _to, uint _tokens)
         onlyManager public{
6          require(manager != 0x0);
7          _to.send(_tokens);
8          balances[_to] = _tokens;
9          Transfer(msg.sender, _to, _tokens);
10     }
11 }
```

Listing 14: A code containing failed functions

RE. We found no exploitations of the 71 contracts with RE vulnerabilities.

FE. Among the 100 FE vulnerable contracts we choose to analyze, four of them have ether. We analyzed the transaction logs of these four contracts and found two contracts had never transferred ether to other accounts. Thus, we label these two contracts as exploited.

NC. 28 out of the 182 contracts with NC vulnerabilities were exploited when executing the functions containing a *for* loop. Listing 15 shows an example of the exploited contract, in which the *for* loop in the function *distribute* iterates over the input parameter *addresses* of the function. The total size of *addresses* in this transaction is 202 and the gas limit of this transaction is 2,417,107 as shown in Figure 10 in Appendix A. This transaction used up all the given gas and failed due to the gas shortage.

```
1 function distribute(address[] calldata addresses, uint256[]
    calldata amounts) payable external {
2     require(addresses.length > 0);
3     require(amounts.length == addresses.length);
4     for (uint256 i; i < addresses.length; i++) {
5         uint256 value = amounts[i];
6         address _to = addresses[i];
7         address(uint160(_to)).transfer(value);
8     }
9 }
```

Listing 15: A code related to out-of-gas transactions

5.3 Results of Analyzing Vulnerability Exploitation

As shown in Table 7, only 6% (66 out of 1,079) exploitable contracts were exploited. We categorize the reasons for unexploitation through open and axial coding. The results are shown in Figure 7.

By performing selective coding, we acquired **Scheme 4: the smart contracts' application scenario and the execution environment and cost on blockchain demotivate or prevent vulnerable contracts from being exploited.**

5.3.1 *Application scenarios demotivate exploitations*

We found that attackers may not be motivated to exploit the vulnerability because their gains or the attacks' impact are trivial.

Very little or no financial benefits (LowBenfit). By exploiting the UpS, TO, RE, TD, and TOD vulnerabilities, the attacker may get profit. However, many vulnerable contracts contain no or very little ether. For example, 70 out

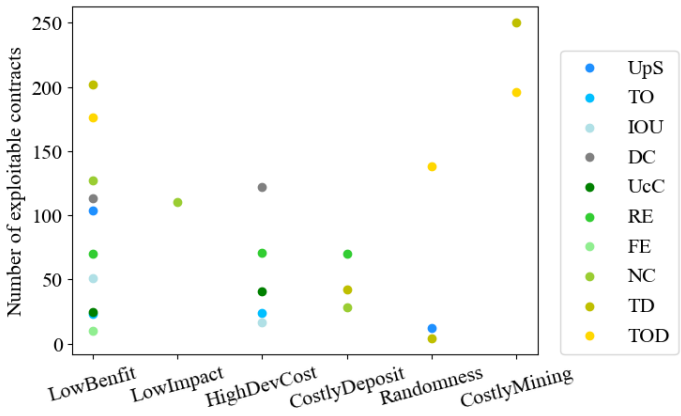

Fig. 7: Reasons to demotivate and prevent the attackers from exploiting the vulnerabilities

of 71 contracts vulnerable to RE have no ether, and only one contract vulnerable to RE holds 0.001307 ether. There are four game contracts vulnerable to TD from which users may win some rewards. However, three of them contain no balance, and the fourth one contains only 0.03 ether.

Insignificant impacts (LowImpact). Attackers may cause contract function failure or user losses by exploiting FE vulnerabilities to freeze ether or exploiting NC vulnerabilities to use up gas. Very few of the contracts vulnerable to FE have ether. Thus, locking a limited amount of ether will not bring significant impacts. Although the 28 NC vulnerability exploitations, few of them bring severe impacts. The for loops in 27 contracts iterate over function input parameters, meaning the size of for iteration is controlled by the function caller. If the caller sets a small number for iterations or gives enough gas, the function will execute successfully. Therefore, most reported NC vulnerabilities will not lead to the permanent failure of the functions, and the failure of transfer functions caused by NC will not lock ether in contracts.

5.3.2 Execution environment and cost imply defense

Some characteristics of Solidity language and Ethereum mechanisms demand attackers to putting in extra effort and investment and be lucky to exploit the vulnerabilities, which may demotivate their exploitation.

Attacker must develop attack contracts (HighDevCost). To exploit some vulnerabilities, e.g., RE, DC, UcC, and TO, attackers must develop attack contracts, which should be customized according to different vulnerabilities. For instance, designing an attacking contract containing a specified fallback function is vital to trigger the contracts vulnerable to RE. To exploit the contracts vulnerable to DC, the attacker needs an attacking contract that controls different global variables and modifies the storage values in vulnerable contracts.

Attacker must deposit ether is a prerequisite (CostlyDeposit). Some vulnerable contracts are used for bidding, games, or wallets. 42 smart contracts vulnerable to TD are wallet contracts supporting users to purchase or withdraw tokens within a limited time. Users must deposit ether into the contract and exchange it for other digital assets. These contracts usually require users to send ether to contracts first to get an authenticated identity. Therefore, sending ether is a prerequisite for an attacker to call the contract. In the example contract in Listing 16, line 11 has a TD vulnerability. Suppose an attacker wants to exploit the vulnerability and get all ether of this contract. In that case, the attacker must meet the condition in line 12, i.e., sending more than 0.001 ether (line 3) and getting a correct *randomNumber* (line 4) to start the exploitation.

```
1  function () public payable {
2      require(msg.sender == tx.origin);
3      require(msg.value >= 0.001 ether);
4      uint256 randomNumber = uint256(keccak256(blockhash(
       block.number - 1)));
5      if (randomNumber > highScore) {
6          currentWinner = msg.sender;
7          lastTimestamp = now;
8      }
9  }
10 function claimWinnings() public {
11     require(now > lastTimestamp + 1 days);
12     require(msg.sender == currentWinner)
13     msg.sender.transfer(address(this).balance);
14 }
```

Listing 16: A code containing a puzzle

Attacker must be lucky in random number competition (Randomness). We found that setting a puzzle as a deciding condition for some critical operations is a popular defense method in our studied contracts. An attacker cannot get permission to run critical operations or get benefits unless the attacker is lucky enough to solve the puzzle successfully. As shown in Listing 16, the *randomNumber* is calculated by the *keccak256* function in line 4 based on the *block.number*, which cannot be controlled by the attacker.

Attacker must be mining winner (CostlyMining). If attackers want to exploit TD and TOD vulnerabilities, they must monitor the transaction pool to capture critical transaction information, such as a puzzle answer. After that, the attacker can initiate a new transaction to compete with the old transaction. The attack will not succeed unless the attacker is a mining winner and the attacker's transaction is successfully packaged.

6 Discussion

Results of RQ1 and RQ2 bring novel insights to vulnerability detector development and evaluation.

6.1 Comparison with related work

Durieux et al. [16] hypothesize that the detectors they evaluate report a considerable number of false positives because the percent of vulnerable contracts (44,589/47,518, 93%) is high. Perez and Livshits [19] also hypothesized that most reported vulnerabilities are either false positives or unexploitable. They identified one factor affecting the actual exploitation of smart contracts, e.g., ether distribution. However, no follow-up study tried to identify and understand other possible reasons for the low exploitation rate.

Different from [19], we uncover that vulnerabilities reported by the detectors are possibly unexploitable and there are 11 reasons for that. These reasons are not limited to the imperfect implementation of the vulnerability detectors but are also relevant to overlooking the characteristics of the

Solidity programming language and smart contract application scenarios. For the exploitable vulnerabilities, beyond ether distribution mentioned in [19] that may demotivate exploitations, our results of RQ2 uncover five other reasons, which are related to financial aspects and blockchain mechanisms, that can help defend the exploitable smart contract against exploitation.

6.2 Implication to Vulnerability Detector Development

Results of RQ1 reveal that state-of-the-art smart contract vulnerability detectors using pattern matching, symbolic execution, data flow analysis, and fuzzing approaches are limited to adapting classical static and dynamic OO application analysis approaches to check smart contracts. Vulnerability detectors for smart contracts need to better consider the unique characteristics of blockchain and smart contract programming languages to identify vulnerabilities and differentiate between vulnerable and exploitable ones. Several issues that cause low exploitability, as shown in Section 4.5, are blockchain specific and require novel or unique detection technologies. Thus, identifying exploitable vulnerabilities needs a more comprehensive analysis of programing language characteristics, contract application scenarios, and digital assets associated with the contract.

Results of RQ2 show that several factors could influence vulnerability exploitation possibilities and vulnerability criticality. When reporting and ranking the vulnerabilities, vulnerable contracts with no ether can be lower ranked. The extra effort and investments needed from attackers and the attackers' chance to execute the attack shall be considered in vulnerability criticality evaluation.

6.3 Implication to Vulnerability Detector Evaluation

Existing studies, e.g., [10], [11], [16], [19], applied three main methods to construct evaluation benchmarks, namely, collecting vulnerable contracts with manual labels, crawling real-world contracts, and injecting vulnerabilities into contracts. Durieux et al. [16] and Ren et al. [10] collected vulnerable contracts with clear labels to evaluate different detectors. However, the number of vulnerable contracts is small, e.g., 69 contracts in [16] and 214 contracts in [10]. These vulnerable contracts are often short and have no complex business logic. Ghaleb et al. [11] constructed a dataset containing 50 contracts with 9,369 injected vulnerabilities. However, the vulnerability injection is limited to known characteristics of vulnerabilities. SolidiFi [11] provides 50 vulnerability patterns for each vulnerability, many of which share the same code logic and only differ in function or variable names. Although studies [10], [16], [19] construct the dataset using real-world contracts, the type and amount of vulnerabilities in these contracts are unknown. This study collected unique 4,364 real-world smart contracts which are cross-labeled by at least two tools as vulnerable. In addition, we categorized the contracts into exploitable and unexploitable. Thus, our dataset can be used as a novel benchmark to evaluate the vulnerability detectors' capability to identify exploitable smart contracts.

6.4 Threats to Validity

In this study, we only focus on the types of vulnerabilities covered by at least two fast vulnerability detectors to avoid bias caused by a single tool. This filtering excluded several types of vulnerabilities that may have different reasons for the low exploitability than those we identified. We found that more than one reason caused the low exploitability of a contract. We give only one reason for each smart contract because we focus on understanding the reasons rather than counting their numbers. The percentage numbers in Figure 5 are calculated based on this strategy. However, such a strategy will not impact the main findings of RQ1, i.e., the 11 reasons. For RQ2, there are probably false negatives due to unknown attacks and exploitations because our log analysis is limited to the rules presented in Table 6.

When manually analyzing the source code of the vulnerable smart contract to label their exploitability, a possible risk is mislabelling. However, we believe that such a risk is low because there are often easy-to-distinguish code features associated with the reasons for the low exploitability, as shown below:

- **Access control**: There is the modifier, i.e., *onlyOwner, onlyAdmin, onlyManager* or *if/require* statement.
- **Constraints caused by factory patterns**: The contracts using the factory Pattern use the same template as shown in Listing 7.
- **Constraints caused by contract inheritances**: The contract inheritance relationship is declared with the contract name.
- ***fallback* functions refuse ether**: If the *revert* statement is in the *fallback* function, the contract refuse to receive ether.
- **Ether transfer initiator**: The amount of ether is *msg.value* that comes from the caller.
- **Critical operation**: Critical operations include function calls or assignment statements, etc.
- **Status inconsistency**: There is a statement after the *call* function that may cause inconsistency between the transfer and the state variable change.
- **Target contract's address**: The address is hard-coded, fixed, or under the control of the contract owner.

7 Conclusion and Future Work

As smart contracts' security is critical, many vulnerability detectors have been proposed. Several empirical studies show that the detectors report many vulnerabilities and the exploitation rate is low, and, therefore, hypothesize many reported vulnerabilities are either false positives or unexploitable. In this study, we have analyzed the exploitability of 4,364 unique real-world smart contracts that are reported as vulnerable by multiple detectors. We identified 11 reasons causing low exploitability. In addition, we discover six aspects that may demotivate attackers to exploit vulnerable contracts. The results of this study delight the need to consider more the characteristics of smart contract programming languages and smart contract application scenarios and execution environments to analyze and rank the smart contract vulnerabilities.

Besides recognizing exploitable vulnerability, another critical aspect of improving vulnerability detectors is reducing false negatives. Our future work will focus on analyzing the false negative results of the vulnerability detectors and give more suggestions to improve the detectors.

ACKNOWLEDGMENTS

This work is jointly supported by the National Key Research and Development Program of China (No. 2019YFE0105500) and the Research Council of Norway (No. 309494) and the Key Research and Development Program of Jiangsu Province (No. BE2021002-3). TY.H thanks the Chinese Scholarship Council (CSC) for financial support (202106090057).

REFERENCES

[1] N. Atzei, M. Bartoletti, and T. Cimoli, "A survey of attacks on ethereum smart contracts (sok)," in *International conference on principles of security and trust*. Springer, 2017, pp. 164–186.

[2] 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.

[3] P. Tsankov, A. Dan, D. Drachsler-Cohen, A. Gervais, F. Bünzli, and M. Vechev, "Securify: Practical security analysis of smart contracts," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 67–82.

[4] T. D. Nguyen, L. H. Pham, J. Sun, Y. Lin, and Q. T. Minh, "sfuzz: An efficient adaptive fuzzer for solidity smart contracts," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 778–788.

[5] B. Jiang, Y. Liu, and W. K. Chan, "Contractfuzzer: Fuzzing smart contracts for vulnerability detection," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. Association for Computing Machinery, 2018, pp. 259–269.

[6] H. Wang, Y. Liu, Y. Li, S. Lin, C. Artho, L. Ma, and Y. Liu, "Oracle-supported dynamic exploit generation for smart contracts," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 03, pp. 1795–1809, may 2022.

[7] J. Chen, X. Xia, D. Lo, J. Grundy, X. Luo, and T. Chen, "Defectchecker: Automated smart contract defect detection by analyzing evm bytecode," *IEEE Transactions on Software Engineering*, pp. 1–1, 2021.

[8] L. Jin, Y. Cao, Y. Chen, D. Zhang, and S. Campanoni, "Exgen: Cross-platform, automated exploit generation for smart contract vulnerabilities," *IEEE Transactions on Dependable and Secure Computing*, pp. 1–1, 2022.

[9] C. F. Torres, M. Steichen, and R. State, "The art of the scam: Demystifying honeypots in ethereum smart contracts," in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1591–1607.

[10] M. Ren, Z. Yin, F. Ma, Z. Xu, Y. Jiang, C. Sun, H. Li, and Y. Cai, "Empirical evaluation of smart contract testing: What is the best choice?" 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. 566–579. [Online]. Available: https://doi.org/10.1145/3460319.3464837

[11] A. Ghaleb and K. Pattabiraman, "How effective are smart contract analysis tools? evaluating smart contract static analysis tools using bug injection," in *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2020.

[12] (2018) Mythril: an open-source security analysis tool for ethereum smart contracts. [Online]. Available: https://github.com/ConsenSys/mythril

[13] S.Tikhomirov, E.Voskresenskaya, I.Ivanitskiy, R.Takhaviev, E. Marchenko, and Y. Alexandrov, "Smartcheck: Static analysis of ethereum smart contracts," in *2018 IEEE/ACM 1st International Workshop on Emerging Trends in Software Engineering for Blockchain*, 2018, pp. 9–16.

[14] J. Feist, G. Grieco, and A. Groce, "Slither: A static analysis framework for smart contracts," in *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain*, 2019, pp. 8–15.

[15] (2019) Manticore: a symbolic execution tool for analysis of smart contracts and binaries. [Online]. Available: https://github.com/trailofbits/manticore

[16] T. Durieux, J. a. F. Ferreira, R. Abreu, and P. Cruz, "Empirical review of automated analysis tools on 47,587 ethereum smart contracts," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, ser. ICSE'20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 530–541. [Online]. Available: https://doi.org/10.1145/3377811.3380364

[17] I. Nikolic, A. Kolluri, I. Sergey, P. Saxena, and A. Hobor, "Finding the greedy, prodigal, and suicidal contracts at scale," in *Proceedings of the 34th Annual Computer Security Applications Conference*. New York, NY, USA: Association for Computing Machinery, 2018, pp. 653–663.

[18] C. F. Torres, J. Schütte, and R. State, "Osiris: Hunting for integer bugs in ethereum smart contracts," in *Proceedings of the 34th Annual Computer Security Applications Conference*, ser. ACSAC '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 664–676. [Online]. Available: https://doi.org/10.1145/3274694.3274737

[19] D. Perez and B. Livshits, "Smart contract vulnerabilities: Vulnerable does not imply exploited," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 1325–1341.

[20] S. Kalra, S. Goel, M. Dhawan, and S. Sharma, "Zeus: Analyzing safety of smart contracts," in *Network and Distributed System Security Symposium*, 2018, pp. 18–33.

[21] J. Krupp and C. Rossow, "Teether: Gnawing at ethereum to automatically exploit smart contracts," in *27th USENIX Security Symposium*, 2018, pp. 1317–1333.

[22] N. Grech, M. Kong, A. Jurisevic, L. Brent, B. Scholz, and Y. Smaragdakis, "Madmax: Surviving out-of-gas conditions in ethereum smart contracts," *Proceedings of the ACM on Programming Languages*, vol. 2, no. OOPSLA, pp. 1–27, 2018.

[23] (2017) Vulnerability vs. exploitability: Why they're different. [Online]. Available: https://cloudtweaks.com/2017/07/vulnerability-vs-exploitability/

[24] T. Hu, B. Li, Z. Pan, and C. Qian, "Detect defects of solidity smart contract based on the knowledge graph," *IEEE Transactions on Reliability*, pp. 1–17, 2023.

[25] K.-J. Stol, P. Ralph, and B. Fitzgerald, "Grounded theory in software engineering research: A critical review and guidelines," in *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*, 2016, pp. 120–131.

[26] C. Ferreira Torres, A. K. Iannillo, A. Gervais *et al.*, "Confuzzius: A data dependency-aware hybrid fuzzer for smart contracts," in *European Symposium on Security and Privacy, Vienna 7-11 September 2021*, 2021.

[27] J. Choi, D. Kim, S. Kim, G. Grieco, A. Groce, and S. K. Cha, "Smartian: Enhancing smart contract fuzzing with static and dynamic data-flow analyses," in *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2021, pp. 227–239.

[28] Y. Xue, M. Ma, Y. Lin, Y. Sui, J. Ye, and T. Peng, "Cross-contract static analysis for detecting practical reentrancy vulnerabilities in smart contracts," in *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 1029–1040.

[29] L. Brent, N. Grech, S. Lagouvardos, B. Scholz, and Y. Smaragdakis, "Ethainter: a smart contract security analyzer for composite vulnerabilities," in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2020, pp. 454–469.

[30] Y. Zhuang, Z. Liu, P. Qian, Q. Liu, X. Wang, and Q. He, "Smart contract vulnerability detection using graph neural networks," ser. IJCAI'20, 2021.

[31] W. Wang, J. Song, G. Xu, Y. Li, H. Wang, and C. Su, "Contractward: Automated vulnerability detection models for ethereum smart contracts," *IEEE Transactions on Network Science and Engineering*, vol. 8, no. 2, pp. 1133–1144, 2020.

[32] F. Mi, Z. Wang, C. Zhao, J. Guo, F. Ahmed, and L. Khan, "Vscl: Automating vulnerability detection in smart contracts with deep

learning," in *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2021, pp. 1–9.

[33] C. Liu, H. Liu, Z. Cao, Z. Chen, B. Chen, and B. Roscoe, "Reguard: Finding reentrancy bugs in smart contracts," in *Proceedings of the 40th International Conference on Software Engineering: Companion Proceeedings*. Association for Computing Machinery, 2018, pp. 65–68.

[34] Q. Zhang, Y. Wang, J. Li, and S. Ma, "Ethploit: From fuzzing to efficient exploit generation against smart contracts," in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2020, pp. 116–126.

[35] (2013) World wide web consortium, sparql 1.1 update. [Online]. Available: https://www.w3.org/TR/sparql11-update

[36] J. Ye, M. Ma, Y. Lin, Y. Sui, and Y. Xue, "Clairvoyance: Cross-contract static analysis for detecting practical reentrancy vulnerabilities in smart contracts," in *2020 IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, 2020, pp. 274–275.

[37] (2021) Wikipedia, datalog: a declarative logic programming language. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Datalog oldid=1053711548

[38] H. Jordan, B. Scholz, and P. Subotić, "Soufflé: On synthesis of program analyzers," in *International Conference on Computer Aided Verification*. Springer, 2016, pp. 422–430.

[39] H. Liang, X. Pei, X. Jia, W. Shen, and J. Zhang, "Fuzzing: State of the art," *IEEE Transactions on Reliability*, vol. 67, no. 3, pp. 1199–1218, 2018.

[40] J. a. F. Ferreira, P. Cruz, T. Durieux, and R. Abreu, "Smartbugs: A framework to analyze solidity smart contracts," in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE'20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1349–1352. [Online]. Available: https://doi.org/10.1145/3324884.3415298

[41] (2022) Solhint: an open source project for linting solidity code. [Online]. Available: https://www.npmjs.com/package/solhint/

[42] (2019) Ethereum (eth) blockchain explorer. [Online]. Available: https://etherscan.io

[43] J. He, M. Balunović, N. Ambroladze, P. Tsankov, and M. Vechev, "Learning to fuzz from symbolic execution with application to smart contracts," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 531–548.

[44] (2022) Jaccard index: a statistic used for gauging the similarity and diversity of sample sets. [Online]. Available: https://en.wikipedia.org/wiki/Jaccard_index

[45] (2021) Securify 2.0: a security scanner for ethereum smart contracts supported by the ethereum foundation and chainsecurity. [Online]. Available: https://github.com/eth-sri/securify2

[46] "Smart contract weakness classification and test cases," https://swcregistry.io/, 2020, accessed 28 May 2022.

[47] A. M. Antonopoulos and G. Wood, *Mastering ethereum: building smart contracts and dapps*. O'reilly Media, 2018.

[48] "Solidity: a statically-typed curly-braces programming language designed for developing smart contracts that run on ethereum," https://soliditylang.org/, 2022, accessed 25 July 2022.

[49] J. Gao, H. Liu, C. Liu, Q. Li, Z. Guan, and Z. Chen, "Easyflow: Keep ethereum away from overflow," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, 2019, pp. 23–26.

[50] (2020) Delegatecall to untrusted callee. [Online]. Available: https://swcregistry.io/docs/SWC-112

[51] P. Praitheeshan, L. Pan, X. Zheng, A. Jolfaei, and R. Doss, "Solguard: Preventing external call issues in smart contract-based multi-agent robotic systems," *Information Sciences*, vol. 579, pp. 150–166, 2021.

[52] C. F. Torres, H. Jonker, and R. State, "Elysium: Automagically healing vulnerable smart contracts using context-aware patching," *arXiv preprint arXiv:2108.10071*, 2021.

[53] 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," *arXiv preprint arXiv:1809.03981*, 2018.

[54] G. Wood. (2016, Jun.) Condition-orientated programming. [Online]. Available: https://gavofyork.medium.com/condition-orientated-programming-969f6ba0161a

[55] E. Gamma, R. Helm, R. Johnson, R. E. Johnson, J. Vlissides *et al.*, *Design patterns: elements of reusable object-oriented software*. Pearson Deutschland GmbH, 1995.

[56] Y. Xue, J. Ye, W. Zhang, J. Sun, L. Ma, H. Wang, and J. Zhao, "xfuzz: Machine learning guided cross-contract fuzzing," *IEEE Transactions on Dependable and Secure Computing*, pp. 1–14, 2022.

Tianyuan Hu is currently working toward the Ph.D. degree with the School of Computer Science and Engineering, Southeast University under the supervision of Dr. Bixin Li. Her research interests include program analysis, vulnerability detection, blockchain security, and software engineering.

Jingyue Li is a Professor at the Computer Science Department, Norwegian University of Science and Technology (NTNU). He received his Ph.D. degree in software engineering from the Department of Computer Science, NTNU, in 2006. His research interests include software engineering, software security, and blockchain technologies.

Bixin Li received his bachelor's degree and master's degree both in mathematics from Anhui University in 1991 and 1994, respectively, and received his doctor' degree in software engineering from Nanjing University in 2001. He is a full Professor of School of Computer Science and Engineering of Southeast University, he is the chairman of Technology Committee of Software Engineering Standards of Jiangsu Province, and he is also the header of Software Engineering Institute of Southeast University in that he is working hard together with more than 50 young people on software architecture and blockchain security projects etc. His main research interests include program slicing and its application, software evolution and maintenance, software testing and verification, software safety and security techniques etc. He has published over 180 research papers and patented more than 80 inventions of china up to now.

André Storhaug is a Ph.D. student in the Department of Computer Science at the Norwegian University of Science and Technology (NTNU). His research interests include machine learning, software engineering, software security, and blockchain technologies.

APPENDIX

TABLE 8: Reasons for Excluding Some Fast Vulnerability Detectors

Year and ref.	Tool Name	SCI	SV	VL	Availability
Pattern Matching					
2018 [13]	Smartcheck				
2021 [24]	SoliDetector				
Symbolic Execution					
2016 [2]	Oyente				
2018 [20]	ZEUS				•
2018 [18]	Osiris	•			
2018 [12]	Mythril				
2018 [3]	Securify		•		
2018 [21]	TEETHER	•			
2018 [17]	MAIAN			•	
2019 [9]	HONEYBADGER	•			
2021 [7]	DefectChecker		•		
2022 [8]	EXGEN				•
Data Flow Analysis					
2018 [22]	MadMax	•			
2019 [14]	Slither				
2020 [28]	Clairvoyance				•
2020 [29]	Ethainter	•			
Machine Learning					
2019 [30]	GNN-based			•	
2020 [31]	ContractWard			•	
2021 [32]	VSCL	•		•	

Note: • means that the tool does no meet the criterion.

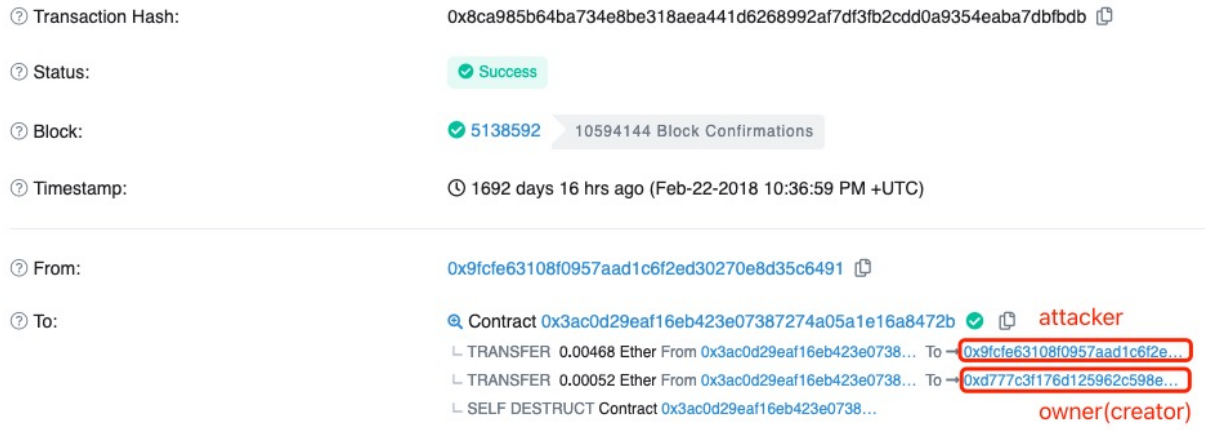

Fig. 8: The transaction details of the exploited UpS

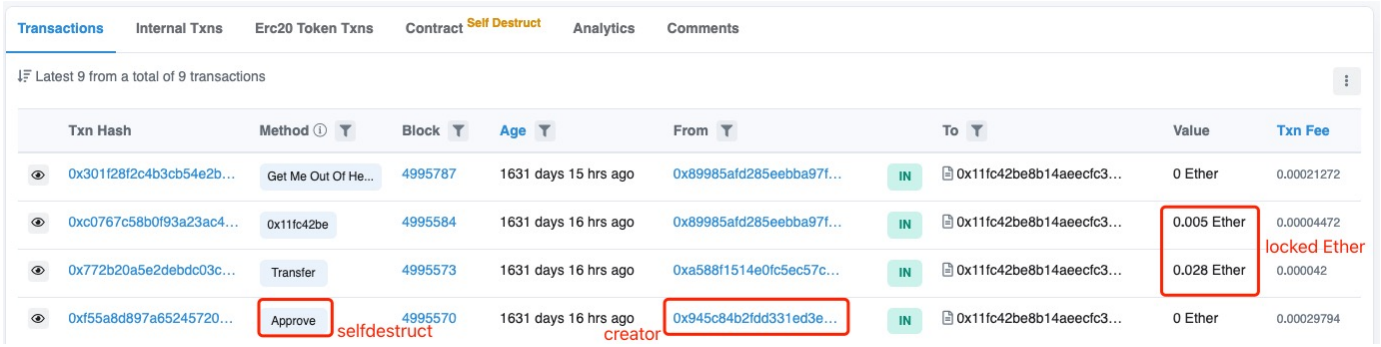

Fig. 9: Locked ether in a self-destructed contract

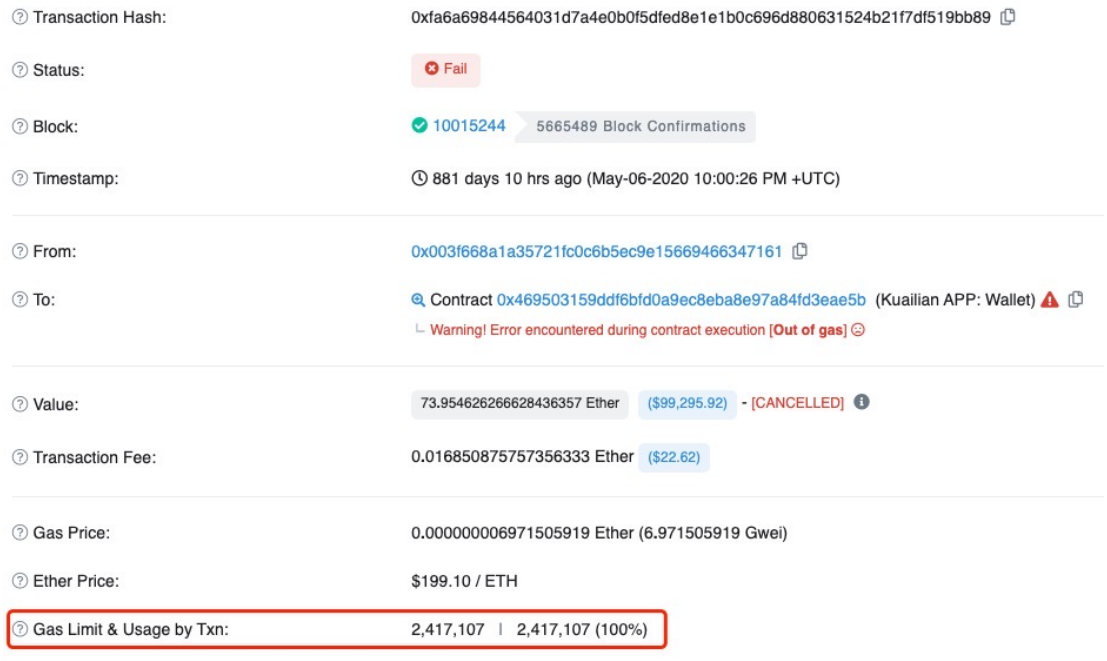

Fig. 10: Transaction information on Etherscan