



All Your Tokens are Belong to Us: Demystifying Address Verification Vulnerabilities in Solidity Smart Contracts

Tianle Sun, Huazhong University of Science and Technology; Ningyu He, Peking University; Jiang Xiao, Huazhong University of Science and Technology; Yinliang Yue, Zhongguancun Laboratory; Xiapu Luo, The Hong Kong Polytechnic University; Haoyu Wang, Huazhong University of Science and Technology

<https://www.usenix.org/conference/usenixsecurity24/presentation/sun-tianle>

**This paper is included in the Proceedings of the
33rd USENIX Security Symposium.**

August 14-16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

**Open access to the Proceedings of the
33rd USENIX Security Symposium
is sponsored by USENIX.**

All Your Tokens are Belong to Us: Demystifying Address Verification Vulnerabilities in Solidity Smart Contracts

Tianle Sun¹, Ningyu He², Jiang Xiao¹, Yinliang Yue³, Xiapu Luo⁴
Haoyu Wang¹‡

¹ *Huazhong University of Science and Technology* ² *Peking University*

³ *Zhongguancun Laboratory* ⁴ *The Hong Kong Polytechnic University*

‡Corresponding author: haoyuwang@hust.edu.cn

Abstract

In Ethereum, the practice of verifying the validity of the passed addresses is a common practice, which is a crucial step to ensure the secure execution of smart contracts. Vulnerabilities in the process of address verification can lead to great security issues, and anecdotal evidence has been reported by our community. However, this type of vulnerability has not been well studied. To fill the void, in this paper, we aim to characterize and detect this kind of emerging vulnerability. We design and implement AVVERIFIER, a lightweight taint analyzer based on static EVM opcode simulation. Its three-phase detector can progressively rule out false positives and false negatives based on the intrinsic characteristics. Upon a well-established and unbiased benchmark, AVVERIFIER can improve efficiency 2 to 5 times than the SOTA while maintaining a 94.3% precision and 100% recall. After a large-scale evaluation of over 5 million Ethereum smart contracts, we have identified 812 vulnerable smart contracts that were undisclosed by our community before this work, and 348 open source smart contracts were further verified, whose largest total value locked is over \$11.2 billion. We further deploy AVVERIFIER as a real-time detector on Ethereum and Binance Smart Chain, and the results suggest that AVVERIFIER can raise timely warnings once contracts are deployed.

1 Introduction

After Satoshi Nakamoto launched Bitcoin [57], blockchain platforms have sprung up. Among them, Ethereum [17] is the most well-known one. Apart from these common characteristics among all blockchain platforms, the most eye-catching feature in Ethereum is the *smart contract*. It can be taken as a piece of unchangeable script that shall be executed in a determined way once the pre-defined conditions are met.

As billions of USD equivalent assets are stored within smart contracts, identifying and exploiting hidden vulnerabilities in them is the top priority for attackers. According to a recent report [65], there exist around 303 large attack events

against well-known Ethereum contracts within the past year, amounting to losses of approximately \$3.8 billion. Chen et al. [19] have summarized 26 major types of vulnerabilities in Ethereum smart contracts, and new types of vulnerabilities are still emerging in an endless stream along the constant feature introduction and deprecation in Ethereum [45].

The automatic identification of vulnerabilities in Ethereum smart contracts is a well-studied topic [14,35,44,53,67,69,75]. Considering the number of contracts and the economic loss caused by false negatives, static symbolic execution, which can guarantee soundness to a certain extent, is used as the main analysis method. For example, Mythril [67] is a static bytecode-level symbolic executor for Ethereum contracts, while Slither [29] works on source code, which claims to achieve high efficiency and effectiveness.

Verifying the validity of the input addresses is a common practice and a key step to ensure the secure execution of smart contracts. Vulnerabilities in the process of address verification can lead to great security issues, and anecdotal evidence has been reported by our community [11]. To be specific, functions in smart contracts can be addressed as arguments. If developers accidentally neglect the verification on the passed address, once it is taken as the target of an external call, arbitrary operations, including malicious ones, within an address can be invoked. Therefore, if any on-chain state modification relies on that external call, such modification may go against the developer's wishes.

Currently, no existing tools can detect the address verification vulnerability, and implementing such a vulnerability detector based on existing framework is challenging. On the one hand, according to the characteristics of the address verification vulnerability, to effectively identify it, the detector should perform an inter-procedural or even inter-contract analysis. Existing work, however, invariably suffers from either the efficiency or the effectiveness problem when conducting such analyses. For example, a pattern-based detector can barely handle such a complicated vulnerability pattern, while the symbolic execution suffers the path explosion issues and the bottleneck brought from the constraints solving. On the one

hand, most contracts are close-sourced, and bytecode lacks sufficient semantics, making it difficult to precisely identify hidden vulnerabilities. Specifically, the address verification vulnerability requires fine-grained tracking in both memory and storage areas. In bytecode, it is hard to distinguish different variables stored in these two areas. Moreover, tracking them via symbolic execution would suffer severe path explosion issue due to the large space of feasible slots.

This work. To fill the void, we aim to characterize and detect this kind of emerging vulnerability. To be specific, we design and implement AVVERIFIER, a novel lightweight static taint analysis framework that can efficiently and effectively identify the address verification vulnerability. Its taint analysis relies on the static simulation of the opcode sequence. In other words, without considering the feasibility of paths, AVVERIFIER maintains the values of data structures (e.g., stack and memory) and the taint propagation status. Such a static simulation traverses lots of paths, including the originally infeasible ones. Thus, according to the corresponding states collected, AVVERIFIER formally detects the address verification vulnerability through a three-phase detection. It can progressively rule out false positives and false negatives based on the intrinsic characteristics of the vulnerability.

Based on our crafted benchmarks, AVVERIFIER significantly outperforms state-of-the-art smart contract detection tools (i.e., Mythril, Ethainter, Jackal, and ETHBMC)¹ in both terms of efficiency and effectiveness. According to a comprehensive evaluation of over 5 million deployed Ethereum contracts, AVVERIFIER flags over 812 vulnerable contracts, and 348 open-source smart contracts were further verified, whose total value locked is over \$11.2 billion. Finally, we deploy AVVERIFIER as a real-time detector on Ethereum and BSC, an EVM-like blockchain platform. The results show that AVVERIFIER is capable of raising early warnings to developers and the community timely. A real-world case, where AVVERIFIER raises the warning 1.5 hours ahead of the attack, proves its effectiveness and efficiency.

This paper makes the following contributions:

- This is the first work on detecting address verification vulnerability. We have designed and implemented AVVERIFIER, an efficient and effective taint analyzer based on static EVM simulation.
- We have applied AVVERIFIER to over 5 million smart contracts on Ethereum, and uncovered hundreds of vulnerable smart contracts which were undisclosed by our community before. This is the first large scale in-the-wild characterization study of this kind of vulnerability.
- Compared to Mythril, Ethainter, Jackal, and ETHBMC, AVVERIFIER can improve the analysis efficiency around

¹Note that these tools cannot detect address verification vulnerability, we have implemented the same detection logic atop them, for a fair comparison.

2 to 5 times while achieving 94.3% precision and 100% recall on well-established benchmarks.

- We have deployed AVVERIFIER as a real-time detector on Ethereum and BSC. The results suggest that AVVERIFIER is able to raise early warnings once contracts are deployed before the attack is initiated by attackers.

2 Background

2.1 Ethereum Primer

There are two types of accounts in Ethereum: *external owned account (EOA)* and *smart contract*. Specifically, an EOA is an ordinary account, which is identified by a unique address and controlled by a private key. Furthermore, smart contracts can be regarded as scripts, which are mainly written in Solidity [26], a well-defined and easy-to-use programming language proposed by the Ethereum official. Interactions among accounts are achieved by initiating transactions, which carry the corresponding data. Smart contracts are executed in Ethereum Virtual Machine (EVM) [41], which is embedded in Ethereum client nodes. EVM is stack-based, and all data is stored either permanently or temporarily. Specifically, all operands of operators and intermediate values are pushed onto and popped from the *stack*. The *memory* area [40], the temporary one, only keeps data under the context of the current transaction. Only the data stored in the *storage* area [53] is permanent, i.e., is stored on-chain. We often denote a set of smart contracts which jointly achieve a specific functionality as a decentralized application (DApp).

DApps have demonstrated significant potential since their rise in 2016 [38]. Many genres of DApps have emerged, e.g., gambling [58], token swap [73], and lending [68]. Alongside the hundreds of billions of USD invested in Ethereum, the decentralized versions of traditional financial tools, e.g., exchanges and insurance, have appeared, which are called decentralized finance, i.e., DeFi. Taking advantage of the decentralization, permissionlessness, and transparency in Ethereum, DeFi starts to rise like a rocket. According to statistics, DeFi accounted for \$163 billion at the end of 2022 [51].

Except for the official token, *Ether*, Ethereum allows users to issue tokens as they wish, as long as these tokens meet the standard like ERC-20 [60]. The ERC-20 standard consists of six mandatory functions. Any smart contract implements these functions can issue valid tokens that can circulate in Ethereum, like USDT [15] and USDC [9]. Therefore, DeFi can also issue its own ERC-20 tokens, and take other ERC-20 tokens as valid ones. The interoperability between ERC-20 tokens and DeFi pushes the prosperity of Ethereum.

2.2 Whitelisted Address Verification

In Ethereum, examining the validity of the given addresses is a common practice, which is called *whitelisted address verifi-*


```

1 function deposit(uint256 amount, address token) external {
2     require(token == usdt, "not usdt token!");
3     token.safeTransferFrom(msg.sender, address(this), amount);
4 }

```

Listing 1: An example of hard-encoded address comparison.

```

1 address[] public addresses;
2
3 function contains(address _address) public view returns (bool)
4     ↪ {
5     for(uint i = 0; i < addresses.length; i++) {
6         if (addresses[i] == _address) {
7             return true;
8         }
9     }
10    return false;
11 }
12
13 function deposit(uint256 amount, address token) external {
14     require(contains(token), "not contain this token");
15     token.safeTransferFrom(msg.sender, address(this), amount);
16 }

```

Listing 2: An example of hard-encoded address enumeration.

ation. It is widely adopted in DeFi apps such as lending [74] and bank [6]. *Address verification is the cornerstone to ensure the safety of smart contracts.* Therefore, OpenZeppelin [61], a well-known standard library provider in Ethereum, offers a whitelisted verification method. Moreover, through a comprehensive study of the top 40 DeFi projects ranked by TVL (Total Value Locked) [50], which account for over 95% of the total DeFi market, we summarized the verification techniques they adopted. In short, three whitelisted verification methods are involved, i.e., *hard-encoded comparison*, *mapping validation*, and *hard-encoded address enumeration*. Note that, though we cannot guarantee all adopted address verification techniques are covered, we cover the most prevalent ones. Considering the extensive copy-and-use in Ethereum smart contracts [38], these three mechanisms are representative.

Listing 1 illustrates how **hard-encoded comparison** works. As we can see, the passed `token` at L2² is required to equal the address of `usdt`, otherwise it raises an exception. **Mapping validation** adopts a *mapping* structure that can dynamically maintain the whitelisted status of addresses, e.g., `mapping(address => bool) whitelist`. As for the **hard-encoded addresses enumeration**, it is a variant of the first one. As shown in Listing 2, an array named `addresses` keeps all whitelisted addresses. Therefore, once the `deposit` function is invoked, the argument `token` is passed to the `contains` function, defined at L3, which is basically a hard-encoded comparison wrapped by a loop. At the bytecode level, these three techniques perform similarly. The contract loads the address in arguments by `CALLDATALOAD` and performs examination via a conditional opcode `JUMPI`. If an address is whitelisted, the control flow will be directed to the *fallthrough*

²L2 refers to the second line, we will adopt this notation in the paper.

branch, and the following logic will be used. Otherwise, the *jumpdest* branch is responsible for handling failed assertions.

2.3 Taint Analysis on Smart Contracts

Taint analysis is a fundamental method of program analysis, used for detecting vulnerabilities [32] and tracking sensitive information flow [44]. Before performing the taint analysis, *sources* and *sinks* should be specifically defined, where *source* refers to input fields controlled by adversaries, and *sink* refers to any part of the system where potentially dangerous data can be used in an unsafe manner. Taint analysis will track data flow from sources to sinks and identify any operations or transformations on the data along the way.

In the context of Ethereum smart contracts, the source is often the function of smart contracts that accept transactions from other accounts, while the sink varies, depending on specific goals. For example, to examine if a contract can be destructed, Ethainter [14] takes the `SELFDESTRUCT` opcode as the sink. Moreover, Michael et al. [32] introduce a tool based on symbolic execution and taint analysis, designating `SSTORE` as the primary sink for its evaluations.

2.4 Threat Model

Adversaries in our study do not require any extra privileges. This is because Ethereum is a permissionless blockchain platform, which allows any non-privileged account, including malicious ones, to initiate transactions with sufficient gas, deploy valid smart contracts, and invoke any already deployed ones. However, certain limitations still exist. For instance, they cannot breach the integrity of the Ethereum network or manipulate the block generation process, and cannot access the private keys of legitimate accounts. In a nutshell, we can barely distinguish adversaries from well-behaved accounts.

3 Motivating Example and Challenges

3.1 Motivating Example

Listing 3 shows a smart contract owned by Visor Finance that is vulnerable to the address verification vulnerability. It was attacked on Dec. 21st, 2021 [11], causing \$8.2 million financial losses. As we can see, the `deposit` function takes three arguments, namely, the number of tokens to be deposited (`visrDeposit`), the payer (`from`), and the beneficiary (`to`). From L6 to L8, it performs sanity checks, i.e., a valid amount of deposit, and valid addresses for both payer and payee. After that, it translates the deposit into shares according to `totalSupply()` (L11 to L14), performs the corresponding token transfer from the `from` address to itself (L16 to L22), and mints some `visr` tokens to the `to` address (L24). The vulnerability is hidden in the `if` code block

```

1  function deposit(
2      uint256 visrDeposit,
3      address payable from,
4      address to
5  ) external returns (uint256 shares) {
6      require(visrDeposit > 0, "deposits must be nonzero");
7      require(to != address(0) && to != address(this), "to");
8      require(from != address(0) && from != address(this),
9          ↪ "from");
10
11     shares = visrDeposit;
12     if (vvisr.totalSupply() != 0) {
13         uint256 visrBalance = visr.balanceOf(address(this));
14         shares =
15             ↪ shares.mul(vvisr.totalSupply()).div(visrBalance);
16     }
17
18     if (isContract(from)) {
19         require(IVisor(from).owner() == msg.sender);
20         IVisor(from).delegatedTransferERC20(address(visr),
21             ↪ address(this), visrDeposit);
22     }
23     else {
24         visr.safeTransferFrom(from, address(this),
25             ↪ visrDeposit);
26     }
27
28     vvisr.mint(to, shares);
29 }

```

Listing 3: The vulnerable deposit function in vvisr.

at L16. Specifically, it allows the `from` address as a contract, and examines if its `owner` function returns the address of the transaction initiator (L17). If the assertion passes, then it invokes the `delegatedTransferERC20` function defined in `from`. Recalling the threat model mentioned in §2.4, attackers are able to deploy contracts and initiate transactions arbitrarily. More specific, if the `from` is actually provided by some malicious ones, they can control the behaviors of L17 and L18. To this end, the control flow will be successfully directed to L24, where `vvisr` finally issues tokens to `to`, which is also controlled by attackers, without receiving any tokens from `from` that is expected by developers of `Visor Finance`.

Through this example, we can summarize three principles related to the address verification vulnerability:

- P1** *The vulnerable function takes an address as a parameter, and performs insufficient authorization examination on that address.* Through the address, attackers can pass self-deployed and unauthorized contracts.
- P2** *The address in P1 is taken as the target of an external call.* Through the external call, the control flow is transferred to attackers. Thus, they can totally control the behavior of the external call, including the return value.
- P3** *On-chain states that are control-flow dependent on the return value mentioned in P2 are updated.* To this end, through an unauthorized control flow, adversaries can get profits by indirectly manipulating on-chain states, like initiating an external call or updating balance.

3.2 Challenges

In response to the address verification vulnerability, as outlined in the summarized principles and the motivating example in §3.1, we identify the following challenges.

C1: Lack of semantics. It is challenging to precisely identify if the address mentioned in **P1** is sufficiently verified due to the lack of semantics in bytecode. According to the statistics [56], more than 99% of Ethereum contracts have not released their source code. The bytecode format is quite unreadable and contains little semantic information. Moreover, there is no debug information to assist in recovering the semantics. Traditional bytecode-based analysis tools usually require some methods to overcome this challenge, like symbolic execution [67].

C2: Inter-procedural analysis on control flow and data flow. Detecting this vulnerability requires accurately extracting the control flow and data flow dependencies inter-procedurally. Specifically, in **P2**, there is an external call to an address, which is passed via the argument. Between the external call and the function entry, it will be propagated several times due to the authorization verification. Thus, we have to precisely identify if the callee address is one of the arguments through parsing data flow. Moreover, in **P3**, the on-chain state update depends on the return value of the external call in **P2** in terms of control flow, which requires us to identify control flow dependencies among variables. In addition, from Listing 2 we can conclude that some authorizations are verified in other functions, which requires inter-procedural analysis.

3.3 Limitations of Existing Tools

Considering the aim of implementing a lightweight and effective detector for address verification vulnerability, we exclude the analyzers that adopt dynamic analysis. The reasons for this decision are twofold. On the one hand, dynamic analysis requires a runtime environment for execution, which is resource-intensive and time-consuming, contrary to our lightweight goal. On the other hand, dynamic analysis identifies vulnerabilities with generated test cases and oracles. For intricate contracts, especially for inter-contract analysis, this method might not always cover all vulnerable paths, potentially resulting in false negatives. Consequently, static methods are considered, including *pattern-based matching*, *symbolic execution*, and *taint analysis*. To the best of our knowledge, no existing tools can be directly deployed to detect this vulnerability. Although we can extend them with the ability to detect the vulnerability, we observe some intrinsic limitations.

Pattern-based Matching. It relies on heuristic rules summarized by developers. Lots of preliminary tools adopt this method, identifying vulnerabilities according to opcode sequences [35], transaction histories [5], and call traces [21]. However, such a manual process is unsound and error-prone, which cannot even handle **C1**, i.e., identifying semantics of

bytecode sequences. Consequently, the constant update of Solidity syntax [72] and compilation toolchain [28] make this kind of analyzers ineffective.

Symbolic Execution & Model Checking. Both techniques are widely adopted in identifying vulnerabilities in software analysis. Taking advantages of abstraction on programs, symbolic executors and model checkers can recover some semantic information to overcome **C1**. However, they are inherently limited by the efficiency issue, mainly brought by the path/state explosion. For example, when dealing with the example shown in Listing 3, both of them should not only traverse all possible functions through the entry to find the `deposit`-like functions, but also get stuck from L16 to L18, where they would try all possible contracts. Because they cannot effectively and efficiently conduct the inter-contract or even inter-procedural analysis, **C2** is the main challenge hinders the adoption of them. Moreover, they try to model the memory area in a precise way, i.e., considering every bit in the memory. Although such a precise modeling benefits to the soundness of analyzing results, it comes at the cost of efficiency. In the context of address verification vulnerability, we only need to focus on the propagation of address parameters in the memory area instead of concrete data.

Taint Analysis. Existing taint analyzers cannot be directly adopted to identify this vulnerability. Firstly, some of them rely on the source code [29]. In Ethereum, such contracts only account for less than 1%. Secondly, some of them are limited by their adopted methods when collecting the taint propagation information and tracking taint propagation. For instance, Sereum [63] uses dynamic analysis to monitor execution at runtime. Furthermore, Mythril [67] and Osiris [70] both adopt symbolic execution to statically collect information. Mythril tries to traverse all feasible paths, while Osiris is only intra-procedural. Ethainter [14] is a strong contender. However, it exhibits proficiency in addressing **C1** and **C2**, but it suffers due to the intricacies of EVM’s linear memory model. Specifically, it cannot properly handle dynamic memory allocations and deallocations in EVM’s memory, which potentially compromises the accuracy of its taint tracking, especially when contracts execute complex memory operations or utilize memory for storing and manipulating intermediate data (like the three memory verification mechanisms mentioned in §2.2).

Our Key Idea. As aforementioned, symbolic execution and taint analysis can effectively recover the semantics and conduct the inter-procedural analysis on control flow and data flow dependencies to some extent. However, considering the soundness of the analysis, symbolic execution always suffers the efficiency issue. To this end, we decided to adopt *static EVM simulation based taint analysis*, which tracks the taint propagation through a static simulation on the bytecode. To reduce the false positives introduced by static simulation and filter out paths that can lead to exploitation, we design a three-stage detector corresponding to the three principles mentioned in §3.1. Moreover, to overcome the memory issue suffered

by Ethainter, we decide to design the memory model motivated by He et al. [39] to abstract the sparse linear memory into key-value pairs. Consequently, our design can address both **C1** and **C2**. Furthermore, AVVERIFIER simulates stack and memory through a self-implemented EVM simulator to capture the propagation of addresses, ensuring efficiency for non-dynamic data. As for the dynamic memory parameters, AVVERIFIER conservatively treats them as symbols and only guarantees the stack balance to avoid the negative influence brought by enumerating all possible values.

4 Design of AVVERIFIER

This section elucidates the technical intricacies of AVVERIFIER, which is designed to detect the address verification vulnerability in Ethereum smart contracts. We firstly give a high-level overview in §4.1, and an introduction of adopted notations in §4.2. Then, we delve into the three components, respectively, from §4.3 to §4.5.

4.1 Overview

Fig. 1 illustrates the architecture and workflow of AVVERIFIER, which is composed of three main components, i.e., *code grapher* (denoted as *Grapher*), *EVM simulator* (denoted as *Simulator*), and *vulnerability detector* (denoted as *Detector*). Specifically, AVVERIFIER only takes the bytecode of a Solidity smart contract as input. The Grapher parses it into the control flow graph (CFG), filters out all suspicious functions as candidates, and delivers them to the Simulator. The Simulator maintains a state, consisting of two parts. One part is the data structures required by EVM, i.e., stack, memory, and storage (see §2.1); the other part is the collected taint information. According to the CFG, the Simulator updates fields in states according to the opcode sequence. It also adopts a heuristic-based path selection method to focus on the most valuable path, i.e., the ones that may lead to the vulnerability. Once the analysis against a path is finished, the corresponding state is sent to the Detector to determine if the current contract is vulnerable to the address verification vulnerability. The cascaded three-phase detection strategy in the Detector rules out false positives and false negatives based on the intrinsic characteristics (**P1** to **P3**).

4.2 Notations

To better explain the implementation of AVVERIFIER, we define some notations here:

- \mathcal{S} , the set of sources that can be controlled by users;
- \mathcal{T} , the set of tainted variables;
- $C_{\mathcal{T}}$, a mapping from a tainted variable to its sources;

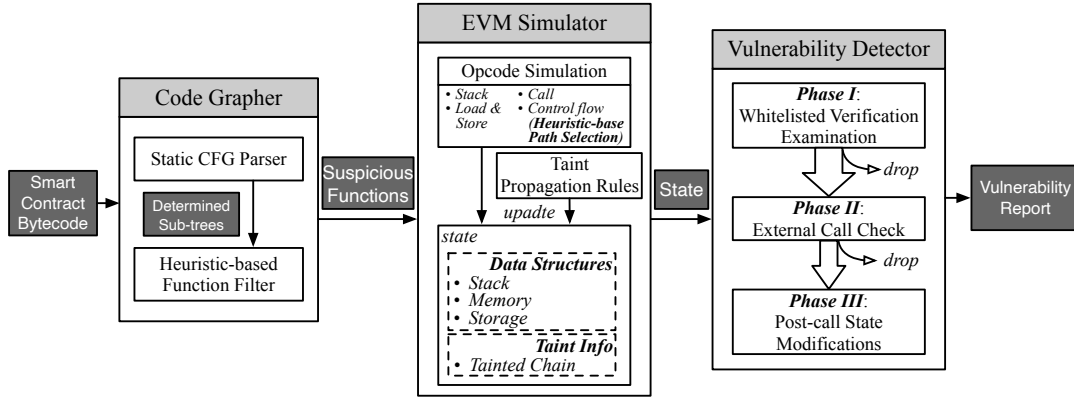


Figure 1: The workflow and architecture of AVVERIFIER.




Figure 2: The CFG of `foo`.

- \mathcal{F} , the set of suspicious functions;
- Mem, Sto , refer to the memory and storage area in EVM.
- V, EC , and SM , refer to the three-stage detection in the Detector, respectively. Each of them takes a function f and a parameter p as inputs.

4.3 Component#1: Code Grapher

Generally speaking, the Grapher is responsible for obtaining the sub-tree of CFG of the given function. Given a piece of bytecode, Grapher firstly extracts the runtime code, consisting of implementation of functions [38]. Then, Grapher parses it into basic blocks and constructs the CFG according to their jump relations. However, some jump relations are determined dynamically at runtime instead of statically at the compilation stage. Thus, the Grapher constructs the CFG only on statically determined jump relations considering the soundness. Take the Fig. 2 as an example, where bb_1 is the entry of the function `foo` and its jump relations to bb_2 and bb_3 can be statically determined. Though bb_3 and bb_4 can jump to bb_5 at runtime, they are determined dynamically. Thus, though the Grapher generates two trees whose roots are bb_1 and bb_5 , respectively, bb_5 is actually a subtree of bb_1 at runtime.

To filter out suspicious functions, i.e., the ones that may be vulnerable to the address verification vulnerability, the Grapher heuristically keeps functions that take addresses as

arguments (**P1**). Specifically, each address parameter undergoes a bitwise AND operation with $0xFF \dots FF$ (160-bit long). By identifying such a specific pattern, which is also widely adopted in previous work [7], we can extract functions that meet **P1**. These functions will be added to the set \mathcal{F} .

4.4 Component#2: EVM Simulator

Based upon the \mathcal{F} passed by the Grapher, the Simulator is responsible for: 1) maintaining the data structures required by EVM, and 2) updating the taint information. Specifically, \mathcal{S} is composed of all user-controllable variables, e.g., `ORIGIN`, `CALldata`, and `BALANCE`. For each opcode, the Simulator specifies a set of rules to update both data structures and taint information in a static way. Without loss of generality, EVM opcodes are classified into four categories: *stack related*, *load & store related*, *call related*, and *control flow related*.

4.4.1 Stack Related Opcodes

The stack-related opcodes include the ones that only interact with the stack and do not change the control flow, like `ADD`, `SWAP`, and `DUP`. Take an opcode that takes two arguments (op_{uid_1}, op_{uid_2}) and returns one (op_{uid_k}) as an example, like `ADD`. Except for conducting necessary update on the stack, we also formalize the taint propagation rules as follows:

$$\{op_{uid_1}, op_{uid_2}\} \cap (\mathcal{S} \cup \mathcal{T}) \neq \emptyset \vdash \mathcal{T} := \mathcal{T} \cup \{op_{uid_k}\}$$

$$C_{\mathcal{T}}[op_{uid_k}] := \begin{cases} \{op_{uid_1}\}, & op_{uid_1} \in \mathcal{S} \cup \mathcal{T}. \\ \{op_{uid_2}\}, & op_{uid_2} \in \mathcal{S} \cup \mathcal{T}. \\ \{op_{uid_1}, op_{uid_2}\}, & \{op_{uid_1}, op_{uid_2}\} \cap (\mathcal{S} \cup \mathcal{T}) \neq \emptyset. \end{cases}$$

The first rule implies that if any operand belongs to \mathcal{S} or \mathcal{T} , the return value shall be added to the \mathcal{T} , i.e., tainted set. The second rule updates $C_{\mathcal{T}}$, which keeps the tainted relation from op_{uid_k} to its direct taint predecessor. Note that, for other opcodes under this category but with different numbers of arguments and return value, the propagation rules are similar.

4.4.2 Load & Store Related Opcodes

As we mentioned in §2.1, memory and storage are maintained by EVM. Data in them are organized in different ways. Specifically, storage operates like a key-value dictionary, i.e., the values are retrieved by the given key [3]. As for the memory, it can be seen as a flat and contiguous array of bytes. Accessing to an item in memory is achieved by a calculated offset. However, in Ethereum smart contracts, a memory is always sparse. Thus, we are motivated by the method raised by He et al. [39], i.e., abstracting the memory area as a key-value pairs, where key is the offset and value is the corresponding data. To this end, data in memory and offset can be retrieved and indexed uniformly, like $Sto[key]$ or $Mem[offset]$.

Load Related Opcodes. This includes `MLOAD` and `SLOAD`, i.e., load from memory or storage, respectively. Both of them take a single argument, tar , i.e., the target address, and return ret , indicating the retrieved data. The taint rules are defined as:

$$\begin{aligned} tar \in \mathcal{S} \cup \mathcal{T} \vdash \mathcal{T} &:= \mathcal{T} \cup \{ret\} \\ C_{\mathcal{T}}[ret] &:= \{tar\}, \text{ if } tar \in \mathcal{S} \cup \mathcal{T} \end{aligned}$$

In other words, if tar is tainted, ret will be marked as tainted. Their dependency relation will be kept in $C_{\mathcal{T}}$. Moreover, there is another situation, i.e., the data designated by tar is tainted already. In this way, we set the taint propagation rules as (take retrieving from memory area as an example):

$$\begin{aligned} Mem[tar] \in \mathcal{S} \cup \mathcal{T} \vdash \mathcal{T} &:= \mathcal{T} \cup \{ret\} \\ C_{\mathcal{T}}[ret] &:= \{Mem[tar]\}, \text{ if } Mem[tar] \in \mathcal{S} \cup \mathcal{T} \end{aligned}$$

Store Related Opcodes. Both `SSTORE` and `MSTORE` take val and $dest$ as arguments, referring to the to-be-stored data and the destination, respectively. Thus, we can formalize the taint propagation rules as:

$$\begin{aligned} \{val, dest\} \cap (\mathcal{S} \cup \mathcal{T}) \neq \emptyset \vdash \mathcal{T} &:= \mathcal{T} \cup \{val\} \\ C_{\mathcal{T}}[val] &:= \{dest\}, \text{ if } dest \in \mathcal{S} \cup \mathcal{T} \end{aligned}$$

That is to say, if any of the val or $dest$ is tainted, the final stored val will be tainted. When maintaining the taint dependency relation, we only consider if val depends on $dest$. There does not exist an edge from val to val .

4.4.3 Call Related Opcodes

Several opcodes can conduct an external call, e.g., `CALL`, `DELEGATECALL`, and `SELFDESTRUCT`. Among all arguments, the address is what we concerned, denoted as $param_{addr}$. Therefore, we customize the taint propagation rules as:

$$\begin{aligned} C_{\mathcal{T}}.ancestor(param_{addr}) = \text{CALLDATALOAD} \vdash \mathcal{T} &:= \mathcal{T} \cup \{suc\} \\ C_{\mathcal{T}}[suc] &:= \{param_{addr}\} \end{aligned}$$

Specifically, our main concern is whether the ancestor of $param_{addr}$ can be tainted from `CALLDATALOAD`, which is the

only source in \mathcal{S} that can parse the address-typed variables. Thus, if $C_{\mathcal{T}}.ancestor(param_{addr}) = \text{CALLDATALOAD}$ holds, it means that attackers can deliver a malicious address through the function entry to that external call finally.

4.4.4 Control Flow Related Opcodes

According to the specification, some of them do not take arguments, like `RETURN`, `STOP`, `REVERT`, and `INVALID`. Therefore, no taint marks are propagated when simulating these instructions. The Simulator only performs the control flow simulation on them. For example, when the Simulator encounters the `INVALID` instruction, it halts the simulation of the current path and moves to the next one.

`JUMP` and `JUMPI` opcodes are crucial in handling taint propagation as they take the jump destination to explicitly alter the control flow, where `JUMPI` takes an argument as the condition. Because `JUMP` can be seen as a special edition of `JUMPI`, we demonstrate the handling on `JUMPI` in the following. Specifically, `JUMPI` takes two arguments, i.e., $dest$ and $cond$, referring to the jump destination and the jump condition, respectively. It returns no value and directly performs the control flow jump. For each `JUMPI`, there are two following branches, $fallthrough$ and $jumpdest$. The former one corresponds to the successive opcode, which is executed once the condition is not met, while the latter one is executed when the condition meets. According to the specification [63], $dest$ should not rely on any arguments, i.e., determined statically during the compilation. Therefore, according to whether the $cond$ is determined dynamically, the Simulator adopts the following rules:

1. If the $cond$ is a concrete number, either $fallthrough$ or $jumpdest$ is chosen by the Simulator deterministically.
2. Otherwise, two paths are all feasible.
 - (a) If the ancestor of the $cond$ is `CALLDATALOAD`, prioritize the $jumpdest$ path, i.e., $cond = \text{True}$, to emulate attackers have successfully bypassed the examination.
 - (b) Otherwise, it takes two paths into consideration.

The branch prioritization in step 2(a) is heuristic. Specifically, once a condition is tainted from a source, it means that it is totally *controllable* for attackers, like L17 in Listing 3, i.e., attackers can always bypass the assertion by constructing the `owner` function. On bytecode level, there is a conditional jump at L17, one is to L18 associated with a `met` condition, another is not shown in Listing 3, indicating the `require` at L17 fails. Thus, the heuristic in step 2(a) prioritizes the branch to L18, i.e., attackers successfully bypass verifications. Compared to symbolic execution, which collects path conditions along the simulation and queries the back-end SMT solver to obtain a concrete set of solutions, the advantages

of this heuristic are twofold. On the one hand, collecting path conditions and asking for solving is time- and resource-consuming, which is the greatest bottleneck [24]. On the other hand, symbolic execution sometimes even cannot handle complex address verification logic. For example, when handling the L17 in Listing 3, the symbolic executor has to conduct an inter-contract analysis to `IVisor` to obtain the value of `owner()`, which cannot be completed within an acceptable time. Though step 2(b) may introduce some false positives, it is a tradeoff between the efficiency and effectiveness. In §5.2, it proves its effectiveness and practicability.

As for the taint propagation rules, we should note that no values are returned by these opcodes. Moreover, they have no side effects on both memory and storage. In summary, no taint marks are propagated through these opcodes.

4.5 Component#3: Vulnerability Detector

Based on the information collected from the Simulator, i.e., \mathcal{F} and $C_{\mathcal{T}}$, the Detector is able to determine whether a contract is vulnerable. Specifically, as Fig. 1 illustrates, the risk detector is composed of three sequential phases, corresponding to the three principles mentioned in §3.1 (P1 to P3). We detail these three phases in the following.

4.5.1 Phase I: Whitelisted Verification Examination

The first phase is to examine whether a contract adopts a whitelisted verification method. Therefore, we design a function $V(f_i, param_j)$. The function takes the i -th function from \mathcal{F} and its j -th parameter as input, verifies if the parameter is examined by any of the methods described in §2.2. The logic of $V(f_i, param_j)$ is defined as:

1. If $param_j$ is not typed as an address, we do not regard this situation to be valid. $V(f_i, param_j)$ returns *True*;
2. Against each `JUMPI` instruction, we examine if its *cond* is tainted by the $param_j$, where $param_j$ is further tainted by the `CALLDATALOAD`. Formally, it is behaved like $param_j \in C_{\mathcal{T}}[cond] \wedge CALLDATALOAD \in C_{\mathcal{T}}[param_j]$. If it is hold, $V(f_i, param_j)$ returns *True*;
3. Otherwise, the tuple $(f_i, param_j)$ returns *False*.

Note that, the first two steps return *True*, indicating a whitelisted verification is inapplicable or conducted normally. In other words, only the states with the *False* return are kept and sent to the phase II check.

4.5.2 Phase II: External Call Check

According to P2, in the second phase, we further investigate if the $param_j$ can be used as the target of any external call related instructions. Therefore, we implement a function, named

$EC(f_i, param_j)$, which returns *True* if the parameter is used as the target of an external call instruction in the function.

In Ethereum, an external call takes an address as a parameter, and allows the current contract to interact with them. The implementation of $EC(f_i, param_j)$ is also intuitive. Specifically, if the tainted predecessor of the *tar* of an external call instruction is $param_j$, or the $param_j$ itself, $EC(f_i, param_j)$ returns *True*. Otherwise, it returns *False*:

$$EC(f, p) := \begin{cases} True, & p \in C_{\mathcal{T}}[tar] \vee tar = p. \\ False, & \text{otherwise.} \end{cases}$$

Similarly, to avoid meaningless resource consumption, only the states that correspond to the *True* return value are passed to the third phase check. We regard the ones with *False* return value as worthless vulnerable contracts.

4.5.3 Phase III: Post-call State Modifications

According to P3, we concern whether on-chain states can be updated according to the return value of the external call instructions focused by the phase II detection. Therefore, we implement $SM(f_i, param_j)$. In Ethereum, on-chain state update can be achieved by two ways. On the one hand, some on-chain states can be seen as ordinary variables and modified directly, e.g., `BALANCE`. On the other hand, `SSTORE` instructions can also be used to alter on-chain state (see §2.1). Therefore, for all these valuable targets, i.e., on-chain states keywords and *val* and *dest* of `SSTORE`, we put them in a set \mathcal{V} . SM returns *True* if any element in \mathcal{V} has a tainted predecessor that is related to the return value, i.e., *suc*, of the external call concerned by the phase II. Formally,

$$SM(f, p) := \begin{cases} True, & \exists e \in \mathcal{V}. suc \in C_{\mathcal{T}}[e] \\ False, & \text{otherwise.} \end{cases}$$

Consequently, for all states with returned value as *True*, they are kept as final ones.

4.5.4 Address Verification Vulnerable Contracts

In a nutshell, through such a three-phase detection, the Detector can effectively identify a state that can be exploited due to the existence of address verification vulnerability. We can formally summarize our detection strategy as follows.

By parsing states passed from Simulator, Detector can obtain a set of tuples, consisting of potential victims:

$$Tuples = \{(f_i, param_j), \dots\} := parse(states)$$

Through a three-phase detection, only the valuable and vulnerable states are remained:

$$Remained = \{(f, p) \in Tuples \mid \neg V(f, p) \wedge EC(f, p) \wedge SM(f, p)\}$$

If a contract has a state that is corresponded to a tuple in *Remained*, the contract is vulnerable to the address verification vulnerability.

5 Evaluation

5.1 Experimental Setup & Research Questions

Baselines. To the best of our knowledge, no existing smart contract analysis framework supports the detection of address verification vulnerability. However, to illustrate the effectiveness of AVVERIFIER, we have extended Mythril [67] (commit: f5e2784), Ethainter [14], Jackal [37] (commit: 3993e5c), and ETHBMC [31] (commit: e887f33), four popular contract analyzers, as baselines. Specifically, Mythril is a well-known and widely adopted symbolic executor that is specifically designed to detect vulnerabilities in Ethereum smart contracts. Thus, we firstly integrate the same taint propagation rules into Mythril. Then, we employ the same three-stage detection logic as we stated in §4.5. Consequently, Mythril is modified as a static symbolic execution based taint analyzer. As for Ethainter, it is a source-code level reputable, extensively utilized, and scalable taint analyzer based on Datalog. After being given a source code, Ethainter will first perform a complete analysis and extract the control flow and data flow dependencies in the contract. Therefore, we also implement a three-stage detector in Datalog to filter out the ones that comply with the rules mentioned in §4.5 from all generated states. Regarding Jackal, it decompiles Ethereum smart contracts' bytecode to an intermediate representation for constructing control flow graphs. We integrate the same taint propagation rules in Jackal, and employ the three-stage detection logic outlined in §4.5. This modification enhances Jackal's capability to focus on identifying address verification vulnerabilities in smart contracts. Last, ETHBMC utilizes bounded model checking and symbolic execution. We also manually integrate the three-stage detection logic into it to sharpen its focus on identifying address verification vulnerabilities. Because detectors in all these tools follow the same set of principles on semantic level, this can reflect the distinction among them in terms of effectiveness and efficiency.

Implementation of AVVERIFIER. AVVERIFIER is written in Python and consists of 1.3K lines of code. As shown in Fig. 1, it is composed of three main modules:

Code Grapher. It is responsible for disassembling the given bytecode into opcodes, and constructing the CFG according to the statical function call opcodes as we mentioned in §4.3. Furthermore, in the function selector, the Grapher selects suspicious functions according to the heuristic in §4.3, and obtains the corresponding entry basic block.

EVM Simulator. The body of the Simulator is basically a two-layer nested loop, the outer one iterates all suspicious functions collected from the Grapher, and the inner one iterates all opcodes. According to the taint propagation rules

defined in §4.4, a *state*, which is composed of data structures like EVM stack, memory, storage, and taint information, is updated. When `JUMPI` is encountered, AVVERIFIER employs the heuristic-based path selection approach. The final state of each path will be sent to the Detector. Once a function is labelled as vulnerable, the inner loop will break to the next function to improve the analysis efficiency.

Vulnerability Detector. Leveraging the state yielded by the Simulator, all states undergo a three-phase check as introduced in §4.5. If any state can pass all three phases, i.e., vulnerable, it will be returned to Simulator immediately.

Experimental Setup. The experiments are conducted on a 48 core server equipped with two Intel Xeon 6248R processors, accompanied by 256GB RAM, while its time limitation for each contract is 10 minutes.

We answer the following research questions (RQs):

- RQ1** Is AVVERIFIER efficient and effective in identifying the address verification vulnerability?
- RQ2** How many smart contracts are vulnerable in the wild and what are their characteristics?
- RQ3** Can AVVERIFIER be deployed as a real-time detection system?

To answer RQ1, we launch AVVERIFIER and four baseline tools on a well-constructed benchmark and real-world contracts in the wild, ranging from the genesis block to the one with the height of 17,421,000, created at Jun. 6th, 2023, around 59 million contracts in total. Among these contracts, we regard the ones that are involved in at least one piece of transaction as the worth being analyzed one. Thus, 5,158,101 smart contracts remain as candidates. To answer RQ2, we characterize the vulnerable contracts from different perspectives, e.g., number of transaction and tokens involved. To answer RQ3, we have deployed AVVERIFIER on both Ethereum and BSC, a well-known EVM-like blockchain platform with the market cap as \$37.7 billion [22]. We comprehensively its usability and scalability of being a real-time detector.

5.2 RQ1: Effectiveness & Efficiency

5.2.1 Evaluating Results on Benchmark

Crafting the Benchmark. After comprehensively collecting technical reports from well-known blockchain security companies [12], we have identified six confirmed vulnerable contracts, as \bar{P} . As all their source code files are available, we manually patch each of them to compose \bar{P} . Moreover, we manually sample four benign contracts from widely-adopted DeFi, i.e., Aave [2], Compound [23], ParaSpace Lending [66], and a yield protocol [30], to form the set N . The reason of selecting these four contracts is twofold. On the one hand, they all require the input of external contract address. Specifically, Aave and Compound primarily employ tokens as collateral to

Table 1: Performance comparison among AVVERIFIER, Mythril, Ethainter, Jackal, and ETHBMC on the benchmark.

Metrics	AVVERIFIER				Mythril*				Ethainter*				Jackal*				ETHBMC*			
	P	\bar{P}	N	\bar{N}	P	\bar{P}	N	\bar{N}	P	\bar{P}	N	\bar{N}	P	\bar{P}	N	\bar{N}	P	\bar{P}	N	\bar{N}
Avg. Time (s)	7.98	6.85	6.74	7.72	24.36	31.62	29.39	28.47	9.74	10.32	12.36	12.15	20.21	18.40	18.05	20.04	0.33	0.35	1.64	1.52
# Timeout	0	0	0	0	2	2	2	2	1	1	1	1	2	2	1	1	0	0	0	0
True Positives	6	-	-	4	2	-	-	1	4	-	-	3	4	-	-	3	0	-	-	0
True Negatives	-	6	4	-	-	4	2	-	-	5	3	-	-	4	3	-	-	6	4	-
False Positives	-	0	0	-	-	0	0	-	-	0	0	-	-	0	0	-	-	0	0	-
False Negatives	0	-	-	0	2	-	-	1	1	-	-	0	0	-	-	0	6	-	-	4
Precision	100%				100%				100%				100%				0%			
Recall	100%				50%				87.5%				100%				0%			

*The address vulnerability detector is implemented by ourselves.

borrow another valuable token, while ParaSpace uses NFTs as collateral. The yield protocol uses the input external contract address to generate collateral yield. Because all these four contracts perform the necessary verification on the passed address, they meet both **P1** and **P2**. On the other hand, after executing certain on-chain operations, the valuable tokens in their contracts are transferred to the caller, thus posing potential risks as **P3**. Similarly, we deliberately remove their verification on addresses to make them vulnerable, denoting this set as \bar{N} . Consequently, we obtained 20 ground truth cases. Table 1 illustrates the results, where the highlighted rows refer to the mis-detected results.

Average Time. It takes AVVERIFIER around 7.34s on average, while Mythril lags considerably, taking about 28.3s on average. Ethainter and Jackal sit in between, with times ranging from 10.86s to 19.19s on average. AVVERIFIER is approximately 3.86x, 1.48x, and 2.61x faster than Mythril, Ethainter, and Jackal, respectively. When considering the timeout cases, both AVVERIFIER and ETHBMC recorded zero, while Mythril, Ethainter and Jackal encountered timeout in 8, 4, and 6 instances, respectively. Additionally, we can easily observe that ETHBMC performs well in terms of executing time. However, after a comprehensive code audit and analysis on output logs, we believe it is because its bounded model checking approach prioritizes efficiency over thoroughness. In other words, paths may be overlooked, which may compromise the accuracy in complex scenarios, like the address verification vulnerability focused in this work.

Precision & Recall. Precision and recall are two critical metrics for evaluating an analyzer’s effectiveness, where AVVERIFIER outperforms other tools. Specifically, AVVERIFIER achieves 100% precision and 100% recall on the benchmark. In the case of Mythril and ETHBMC, the main issue is false negatives. For the cases that can be completed within the time limit, Mythril and ETHBMC have a 50% and 100% false negative rate, respectively. We speculate that the primary reason for ETHBMC’s non-ideal results is twofold. On the one hand, ETHBMC’s bounded model checking strategy inherently focuses on a specific range of states and paths within contracts, potentially missing the complexities involved in address verification due to its limited scope. On the other hand, ETHBMC necessitates a pre-defined initial state for analysis.

However, this state is very likely not optimal for detecting the address verification vulnerability, potentially affecting its performance. Ethainter also has a worse performance compared to AVVERIFIER in terms of recall, with a false negative rate of around 12.5%. We think the most critical factor is the adoption of Gigahorse [34], a toolchain for binary analysis. According to its implementation, one of its limitations is its inability to perfectly handle dynamic memory, affecting the performance of Ethainter in identifying functions that extensively use dynamic memory. Consequently, this limitation leads to the false negatives.

Root Causes. Considering the differences in metrics when conducting analysis on the benchmark among these five tools, we speculate that there are four reasons for their distinctions on the performance on the benchmark. First, AVVERIFIER fully leverages the characteristics summarized from **P1**. In the Grapher, it filters suspicious functions as candidates, which significantly reduces the number of possible states, a predicate affects these tools. Second, as detailed in §4.4.4, the path-searching strategy employed by the Simulator is specifically designed for the address verification vulnerability. This strategy prioritizes paths that may lead to vulnerabilities. Thirdly, when handling the dynamic memory, the other four tools struggle to accurately analyze vulnerable functions that extensively employ complex dynamic memory allocation. In contrast, AVVERIFIER leverages an EVM simulator, enabling it to precisely track address parameters without explicitly modeling dynamic memory behaviors, thereby enhancing its capability to identify such functions. Last, AVVERIFIER adopts a straightforward simulation approach, rather than static symbolic execution. This choice contributes to its efficiency. Previous studies, like KLEE [18], suggest that backend SMT solvers can be significant drags on performance.

5.2.2 Real-world Contracts Results

To further illustrate the effectiveness of AVVERIFIER on real-world contracts, we perform the analysis on all collected contracts, 5,158,101 ones in total. Consequently, 812 of them are marked as vulnerable by AVVERIFIER. To evaluate the effectiveness of AVVERIFIER, we again use Mythril, Ethainter, Jackal, and ETHBMC as baselines. However, because the un-

Table 2: Performance comparison among AVVERIFIER, Mythril, Ethainter, Jackal, and ETHBMC on real-world contracts.

Metrics	AVVERIFIER	Mythril*	Ethainter*	Jackal*	ETHBMC*
Avg. Time(s)	6.34	33.69	8.74	29.96	5.43
# Timeout	0	42	6	60	164
True Positives	348	16	147	172	2
True Negative	0	2	4	11	21
False Positives	21	8	17	10	0
False Negatives	0	301	195	116	182
Precision	94.3%	66.7%	89.6%	94.5%	100%
Recall	100%	5.1%	43.0%	59.7%	1.1%

*The address vulnerability detector is implemented by ourselves.

readability of the bytecode, for a more effective comparison, we tried to obtain their source code from Etherscan. Finally, we collected 369 pieces of source code. The final scanning results for all these five tools on 369 open-source contracts are shown in Table 2.

Average Time. As we can see, for all 369 cases, AVVERIFIER achieves the second-best performance in terms of average analysis time. Moreover, there are no timeout cases within the 10-minute limit. Ethainter is one place behind, with 8.74s and 6 timeouts. Mythril’s efficiency lags far behind, averaging 33.69s per case and suffering 42 timeouts exceeding the 10-minute threshold, the highest among compared tools. Jackal averages around 29.96s with 60 timeouts within. Finally, though ETHBMC has a decent 5.43s average time, it suffers 164 timeout cases, which significantly impacts its effectiveness. By comparing Table 2 and Table 1, we can observe that the performance among these tools is roughly consistent, except that ETHBMC has more timeout cases on the real world cases. We speculate that this is because it needs to try different initial states when the current one cannot explore paths to exploitations, leading to a huge efficiency issue.

Precision & Recall. After manually rechecking all these contracts, the numbers of false positives and false negatives are also shown in Table 2. We can easily observe that there are 21 false positives generated by AVVERIFIER, leading to 94.3% precision. The main reason for that is there are *unconventional verification methods on addresses*. Except for the three mechanisms we summarized in §2.2, some of them delegate address verification to other contracts, which is not a widely adopted verification method. Moreover, some contracts perform verification via digital signatures [16] or Merkle proofs [47]. Currently, due to efficiency issues, AVVERIFIER does not integrate such patterns. Moreover, as inter-contract analysis is always a huge obstacle for smart contract analysis [55], it is a compromise must be made. In contrast, all the other four baselines suffer from severe false negative issues. Although Mythril employs a similar path filtering approach to AVVERIFIER, its recall is only 5.1%, because the symbolic execution cannot effectively find feasible paths to exploit vulnerable contracts. Ethainter and Jackal, both of which use the GigaHorse framework, achieve recall rates of only 43.0% and

59.7%, respectively. As we mentioned in §5.2.1, Gighorse struggles to accurately construct complete CFGs when handling some contracts, which stems from its less optimized handling on the dynamic memory. The recall of ETHBMC is only 1.1%, whose reason is mainly due to its adopted initial states as we stated above. We have conducted a case study to illustrate how a case is mislabeled as false negative by these four tools, please refer to our open-source repo at [link](#).

RQ-1 Answer

Compared to Mythril, Ethainter, Jackal, and ETHBMC, AVVERIFIER can improve the efficiency 2 to 5 times. AVVERIFIER can achieve at least 94% precision and 100% recall on the well-constructed benchmark and real-world contracts, while the others suffer a severe false negative issue due to their design and implementation.

5.3 RQ2: Characteristics of Real-world Vulnerable Contracts

To characterize the ecosystem of real-world vulnerable contracts, we firstly illustrate the overall results on all vulnerable ones among over 5M collected ones (see §5.3.1). Then, we focus on the behavioral characteristics, including *financial related* and *activity related*, of these vulnerable contracts (see §5.3.2 and §5.3.3).

5.3.1 Overall Results

As we mentioned in §5.2.2, we have applied AVVERIFIER on all 5.2M contracts. In total, we have identified 812 vulnerable ones. Among them, we found 443 of them are close-sourced. According to the MD5 results, we have successfully identified 131 unique close-sourced contracts. To recheck the identified results, we decompiled the unique close-sourced ones by contract library tools [1], a well-known decompiler, and asked two Ph.D. students who major in this area to recheck the results. The manual recheck has not revealed any false positives. Due to the unreadability of the close-sourced bytecode, even for the decompiled ones, we can only confirm that 17 of close-sourced contracts are related to Ethereum tokens according to the function signature.

As for the remaining 348 true positive cases mentioned in Table 2, we obtained their source code from Etherscan. Similarly, by deduplicating the source code, we finally obtained 299 unique ones. After a manual recheck, the classification results are shown in the Table 3. We can observe that nearly half of them are related to ERC-20, indicating potential financial impacts of this vulnerability. Moreover, around one-third are in the DApp category, like lending market or swap router, which may also bring in impacts to the Ethereum ecosystem.

We believe that these open-source contracts can gain more user trust and are more representative. In the following §5.3.2

Table 3: For 348 identified true positive open-source cases, the classification results according to their functionalities.

Contract Type	ERC-20	ERC-721	DApp
Open-source Contract	177	61	110
Unique Open-source Contract	153	45	101




Figure 3: Distribution of vulnerable ones by creation time.

and §5.3.3, we characterize these 348 open source smart contracts, as a lower-bound of the overall landscape.

5.3.2 Activity Related Metrics

To depict the activity of them, we first illustrate the distribution of their deployment time, as shown in Fig. 3. As we can see, the figure illustrates a general upward tendency, which aligns with our intuition as Ethereum has been growing and the propagation of a vulnerability is a temporal phenomenon. There are two noticeable troughs in late-2019 and late-2021. Delving deeper, we discerned that the first trough was primarily influenced by stringent cryptocurrency policies, leading to a liquidity crunch. Meanwhile, the second one can be attributed to the aftermath of the *Luna Event* [48], which precipitated a significant liquidation of assets, further causing a depletion in liquidity. We can also observe a peak located in the January 2019. This is because the birth of Uniswap [71] in November 2018, which played a pivotal role in the DeFi prosperity, leading to a substantial increase in the number of contract creations as well as the vulnerable ones.

We also observe the lifespan of these 348 vulnerable smart contracts in Fig. 4, where the y-axis and x-axis represent the block height when a contract was created and the time when the last transaction occurred respectively. The size of the bubble is proportionally to its historical transaction count. As we can see, there are lots of tiny bubbles locate along or near the diagonal red line. This implies that a significant number of contracts had a very short lifespan. Such a trend suggests a plethora of transient contracts, potentially due to testing exercises, spamming campaigns, or temporary endeavors within the Ethereum ecosystem. Interestingly, while larger bubbles scattered across the figure signify contracts with considerable transactional activity, they may still be susceptible to the address verification vulnerability. We speculate their persistent activity suggests that they may correspond to low balance, which likely diminishes the motivation for attack-




Figure 4: Relations between creation time and lifespan for vulnerable contracts.




Figure 5: Distribution of the number of transactions involved in vulnerable contracts.

ers to exploit their vulnerabilities. Additionally, we observe some large bubbles distribute near the red line. This indicates a burst of transactions in their earlier phases but have since transitioned into a dormant state. Such an intriguing contrast between their past vibrancy and current inactivity prompts us to further investigate these contracts. Therefore, we have filtered out the transactions of the top 50 contracts in terms of the number of involved transactions, i.e., the size of bubbles. Through a detailed and comprehensive transactional analysis, we found that 3 of them have been exploited already, while 47 of them are at risk. We speculate the reason is that the balance is small and attackers have not noticed yet or are waiting for the opportunity to make a large profit.

5.3.3 Financial Related Metrics

We further evaluate the financial impact of these vulnerable smart contracts. Intuitively, by measuring how many transactions are involved in a contract can reflect its financial impact to some extent. Fig. 5 illustrates the distribution of the number of transactions involved in them. We can easily observe that the distribution follows the Pareto principle [27], i.e., there exists a long tail in the distribution. More than 55.46% cases are involved within 5 transactions, while the case named AnySwap [8] is involved in around 237,000 transactions. The results follow the Oliva et al.'s [59], i.e., most contracts in Ethereum are inactive, a small portion of contracts greatly contribute to the prosperity of the Ethereum ecosystem.

In addition, we desire to evaluate how many assets are directly involved in them. For DeFi projects, TVL (Total Value

Locked) is one of the most representative metrics. Thus, for each case, we retrieve its historical TVL from a third-party API, DeFiLlama, a DeFi data browser known for providing metrics like TVL and market cap. For each project, we retrieve their peak TVL as it corresponds to the timestamp when attackers can obtain the most profit. In total, around \$11.2 billion was considered to be directly locked into these vulnerable contracts. The top-3 Ethereum projects in terms of compromised assets include Visor³ [11], TempleDao⁴ [13], and AnySwap⁵ [62]. Interestingly, except for Anyswap, other two projects immediately stop providing services after being attacked. This is because hackers exploit a vulnerability in Anyswap to indirectly steal tokens from users. Conversely, in the other two projects, the funds within the vulnerable contracts were directly accessed and stolen, leading to an immediate cessation of services. Moreover, the time window between the creation and the corresponding exploitation suggests the opportunity for vulnerability detection and remediation. While Visor saw exploitation within 3 months, Anyswap remained uncompromised for 210 days. Such a variation might result from factors like the implementation complexity, public visibility, or the inherent vulnerability's nature. However, it proves that there is often a time window to detect and patch vulnerabilities before an attack happens.

RQ-2 Answer

Around 68.4% vulnerable contracts are ERC-20 or ERC-721 tokens. We further reveal that attackers tend to launch attacks dozens of days after the deployment for greater benefits, suggesting there exists a time window for vulnerability detection and remediation.

5.4 RQ3: Real-time Detection

We aim to deploy AVVERIFIER as a real-time detector. Thus, we measured several performance metrics (see §5.4.1). Additionally, we give a case study to illustrate how a vulnerability can be detected by AVVERIFIER before an attack (see §5.4.2).

5.4.1 Quantitative Analysis

We measure two real-world performance metrics. First, we compare the rate of contract creation along block generation to the performance of AVVERIFIER, which can shed light on the responsiveness and real-time applicability of AVVERIFIER. Second, we illustrate the correlation between bytecode length and the consumed time taken for analysis. This metric indicates the scalability of AVVERIFIER with the increasing complexity and size of contracts.

Rate on Contract Creation vs. Detection. We directed our attention to the data from Nov. 2022 to Jan. 2023, a period

³0xC9f27A50f82571C1C8423A42970613b8dBDA14ef

⁴0xd2869042E12a3506100af1D192b5b04D65137941

⁵0x6b7a87899490EcE95443e979cA9485CBE7E71522




Figure 6: The relationship between the bytecode length and the time consumed on each case.

of time when contracts are heavily deployed (illustrated in Fig. 3). According to our statistics, these three months cover blocks with the height from 15,870,000 to 16,518,000, accounting for 289,238 deployed contracts. In other words, 0.45 contract is deployed on average within each block. As for BSC, according to a widely known BSC browser, BscScan, we can calculate that a block is generated every 3s, and there are 2.1 contracts deployed in each BSC block on average. According to the results in RQ1, each Ethereum contract takes around 6.42s. Therefore, considering the number of contract deployed in each block and the speed of block generation in Ethereum, a single-core processor can be used to deploy AVVERIFIER as a real-time detector. As for BSC, each block spends around $6.42s \times 2.1 = 13.48s$, greater than the time taken by the block generation. However, the methodologies adopted by AVVERIFIER can be paralleled easily, like analyzing multiple suspicious functions simultaneously. Therefore, a multi-core machine is sufficient.

Scalability. To evaluate the scalability of AVVERIFIER, we randomly sample 1,000 contracts from the ones deployed within the recent year. Figure 6 presents the relation between the bytecode length and the consumed time. Clearly, there does not exist a linear correlation or even an exponential one between these two metrics. We can also observe that most cases can be finished within 20s. Such a high detection efficiency can be attributed to two points. On the one hand, the detection logic is very efficient. For example, the Detector can effectively screen suspicious functions, and can stop the analysis in time when a vulnerability is encountered. On the other hand, the detection method has few performance bottlenecks. Unlike static symbolic execution techniques, the Simulator can quickly and accurately traverse paths that could lead to vulnerabilities. Therefore, the spent time of AVVERIFIER on each case is not directly proportional to bytecode length, illustrating its scalability.

5.4.2 Case Study: A Real-world Early-warning Case

We illustrate a real-world case that is marked as vulnerable when AVVERIFIER is deployed as a real-time detector on BSC. As there is no source code for the case, Listing 4 illustrates its decompiled version. Moreover, due to the non-disclosure principle, we make a slight change syntactically

```

1 function 0xbc423deb(uint256 varg0, uint256 varg1, uint256
  → varg2) public payable {
2     require((address(varg2)).code.size);
3     v0 = address(varg2).slip(_ilk, msg.sender, 0 -
  → varg1).gas(msg.gas);
4     require(v0); // checks call status, propagates error
  → data on error
5     v1, v2 = _gem.transfer(address(varg0), varg1).gas(msg.gas);
6     require(v1); // checks call status, propagates error
  → data on error
7     emit Exit(address(varg0), varg1);
8 }

```

Listing 4: A case that is detected by our real-time detector.

without modifying its original semantics.

As we can see, L2 indicates that the `varg2` is an address passed from the external environment (satisfying **P1**). At L3, the function invokes `slip`, which takes the `varg2` as the target address (satisfying **P2**). Then, at L4, a `require` checks the returned value (satisfying **P3**). Finally, at L5, the contract invokes `transfer` to transfer tokens to the address referred by `varg0`. Therefore, attackers can deploy a contract to bypass the verification on `v0` to drain this contract out.

Notably, we detected this vulnerability on May 18th, 2023, at 6:10 UTC. An attack transaction was initiated 1.5 hours after, at 7:41 UTC. Such a time gap highlights that capability of AVVERIFIER. However, the absence of an automated exploit response mechanism within AVVERIFIER prevented timely intervention, leading to a user loss of \$ 30K USD. This incident underscores the importance of designing an automated response tool to mitigate potential financial damages resulted from the address verification vulnerability.

RQ-3 Answer

Though a quantitative analysis, it illustrates the usability and scalability of AVVERIFIER as being a real-time detector. A real-world case, worth approximately \$30,000, proves its ability to raise early-stage warnings.

6 Discussion

6.1 Threats to Validity

Scalability of AVVERIFIER. AVVERIFIER is specifically designed for the address verification vulnerability. However, due to the efficient taint analysis we implemented in AVVERIFIER, it can be easily extended to detect other kinds of vulnerabilities, like unchecked external calls and inadequate access controls. Moreover, AVVERIFIER can also be adopted on other blockchain platforms, i.e., EVM-compatible chains, like BSC, TRON, and AVAX. Consequently, AVVERIFIER is not a tool limited to the address verification vulnerability on Ethereum.

Candidate Contracts. When answering RQ1 and RQ2, we only choose around 5 million contracts as candidates. According to the statistics, there still exist millions of contracts. However, we think those already deployed ones without trans-

actions are the worthless targets for attackers. However, in answering RQ3, i.e., deploying AVVERIFIER as a real-time detector, we take all newly deployed contracts as candidates.

6.2 Limitations

Dynamic Parameters Verification Mechanism. Some smart contracts employ dynamic parameters verification mechanisms, like digital signatures, which require loading data into memory. To this end, the size of `CALLDATA` is unknown beforehand, posing challenges to AVVERIFIER as the variable length of `CALLDATA` leads to unpredictable offsets, making it difficult to track variable arguments in dynamic memory. Currently, for such dynamic parameters, AVVERIFIER concretizes their length to obtain the final results. Due to its possibility of introducing false negatives, we take this as one of our future research directions.

Scalability. AVVERIFIER is a tool specifically designed to identify the address verification vulnerability. As such, currently, it cannot be used for identifying other smart contract vulnerabilities directly. However, because we have implemented a comprehensive set of taint propagation rules, AVVERIFIER can be easily extended to other smart contract security problems.

Auto-exploitation. While AVVERIFIER can efficiently and accurately identify the address verification vulnerability and the entry point, it cannot automatically generate exploitation yet. Currently, all identified contracts still require a comprehensive manual confirmation to determine if they are exploitable. However, we can take advantage of other techniques, like fuzzing and symbolic execution, to help AVVERIFIER generate exploitations and `CALLDATA` based on the generated results. We believe it is a meaningful research area.

6.3 Ethical Consideration

Due to the anonymity of blockchain, it is only possible to contact the developers of open-sourced projects. Thus, we have tried our best to contact developers of vulnerable contracts through all possible and usable media, like mail, official website, and twitter, with more than 500 transactions in history. Unfortunately, after we warned 48 developers once after the identification and manual recheck, we did not receive any responses from them within 2 weeks. Where we deploy a real-time detector on Ethereum and BSC, we observed that AVVERIFIER has raised alarms several times, like the one shown in §5.4.2. However, it is impossible to reach its developers privately. Moreover, we cannot initiate a transaction to notify the developer, because it will also notify malicious users that there is a possible victim. Thus, we urge collaboration with blockchain security companies for the secure and efficient disclosure of vulnerabilities.

7 Related Work

Smart Contract Vulnerability Detection. Vulnerability detection in smart contracts employs varied methodologies dependent on input types and detection principles. Analyzing contracts can either occur at the high-level source code [33,52] or through the bytecode interfacing with the EVM [4, 36]. Further categorizing based on analytical techniques, static analysis delves into code structure and inherent semantics to identify vulnerabilities, often applied to source code evaluations. Dynamic analyses, more prevalent in bytecode assessments, utilize strategies such as fuzz testing to spot anomalies by bombarding contracts with randomized inputs [10, 42]. In addition to these, hybrid analysis methods are also gaining traction, combining static and dynamic analysis techniques for a more thorough examination of smart contracts. Through this combined approach, hybrid analysis helps in identifying a wider range of vulnerabilities [46, 54]. Furthermore, trace-based evaluations provide a unique perspective by scrutinizing historical transaction patterns, unearthing vulnerabilities from real-world usage patterns [21, 76].

Taint Analysis in Smart Contracts. Pioneering efforts in taint analysis have led to the development of tools tailored specifically for Ethereum contracts, adept at detecting common vulnerabilities such as unchecked send and reentrancy [6, 63]. Other research efforts have heightened the granularity of taint analysis to uncover intricate data leaks and permission oversights [69, 70]. Furthermore, several advanced methods have integrated both static and dynamic analysis techniques, blending the advantages of both to provide a more comprehensive security assessment [33, 44]. These collective advancements underscore the pivotal role of taint analysis in shaping a robust ecosystem for smart contracts.

Permission Checks and Access Control. Permission checks underpin smart contract security, preventing unauthorized actions which can lead to financial or data losses [20, 49]. Many tools and frameworks now aid developers in verifying permissions [43, 64]. However, some advanced attacks can bypass traditional checks, underscoring the need for context-aware analysis [25, 77]. Rigorous permission validation remains crucial for a secure smart contract environment.

8 Conclusion

In this work, we present AVVERIFIER, a taint analyzer based on static EVM opcode simulation, which is designed for identifying the address verification vulnerability hidden in Ethereum smart contracts. With the help of the heuristic-based path selection method and taint propagation rules in Simulator, as well as the three-phase formal detection rules implemented in Detector, AVVERIFIER significantly outperforms Mythril in both terms of efficiency and effectiveness. According to a comprehensive evaluation on over 5 million contracts, as well as the behaviour characteristics they illustrate, it proves

the necessity of implementing AVVERIFIER. Additionally, AVVERIFIER is proven efficient and effective enough to be a real-time detector on EVM-like blockchain platforms to raise early warnings once contracts are deployed.

Availability

We have released AVVERIFIER and the benchmark at [link](#).

Acknowledgements

We are deeply grateful to our shepherd and the anonymous reviewers for their insightful comments. This work was supported in part by the National Key R&D Program of China (2021YFB2701000), the National Natural Science Foundation of China (grant No.62072197), the Key R&D Program of Hubei Province (2023BAB017, 2023BAB079), the Knowledge Innovation Program of Wuhan-Shuguang and Hong Kong RGC Projects (PolyU15224121, PolyU15231223). The full name of Tianle Sun and Haoyu Wang's affiliation is Hubei Key Laboratory of Distributed System Security, Hubei Engineering Research Center on Big Data Security, School of Cyber Science and Engineering, Huazhong University of Science and Technology.

References

- [1] Dedaub decompiler. <https://app.dedaub.com/decompile>.
- [2] Aave. Aave documentation hub. <https://docs.aave.com/hub/>, 2023.
- [3] Elvira Albert, Pablo Gordillo, Alejandro Hernández-Cerezo, and Albert Rubio. A max-smt superoptimizer for evm handling memory and storage. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 201–219. Springer, 2022.
- [4] Elvira Albert, Pablo Gordillo, Benjamin Livshits, Albert Rubio, and Ilya Sergey. Ethir: A framework for high-level analysis of ethereum bytecode. In *International symposium on automated technology for verification and analysis*, pages 513–520. Springer, 2018.
- [5] Amir Ali, Zain Ul Abideen, and Kalim Ullah. Sescon: Secure ethereum smart contracts by vulnerable patterns' detection. *Security and Communication Networks*, 2021:1–14, 2021.
- [6] Ayman Alkhalifah, Alex Ng, Paul A Watters, and ASM Kayes. A mechanism to detect and prevent ethereum blockchain smart contract reentrancy attacks. *Frontiers in Computer Science*, 3:598780, 2021.

- [7] Sidney Amani, Myriam Bégel, Maksym Bortin, and Mark Staples. Towards verifying ethereum smart contract bytecode in isabelle/hol. In *Proceedings of the 7th ACM SIGPLAN international conference on certified programs and proofs*, pages 66–77, 2018.
- [8] Anyswap. Anyswap. <https://github.com/anyswap>, 2023.
- [9] Douglas W Arner, Raphael Auer, and Jon Frost. Stablecoins: risks, potential and regulation. 2020.
- [10] Imran Ashraf, Xiaoxue Ma, Bo Jiang, and Wing Kwong Chan. Gasfuzzer: Fuzzing ethereum smart contract binaries to expose gas-oriented exception security vulnerabilities. *IEEE Access*, 8:99552–99564, 2020.
- [11] Beosin. Two vulnerabilities in one function: The analysis of visor finance exploit. <https://beosin.medium.com/two-vulnerabilities-in-one-function-the-analysis-of-visor-finance-exploit-a15735e2492>, 2023.
- [12] BlockSecTeam. Blocksecteam blog on medium. <https://blocksecteam.medium.com/>, 2023.
- [13] BlockSecTeam. Tweet by blocksecteam. <https://twitter.com/BlockSecTeam/status/1579843881893769222>, 2023.
- [14] Lexi Brent, Neville Grech, Sifis Lagouvardos, Bernhard Scholz, and Yannis Smaragdakis. Ethainter: a smart contract security analyzer for composite vulnerabilities. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 454–469, 2020.
- [15] Dirk Bullmann, Jonas Klemm, and Andrea Pinna. In search for stability in crypto-assets: are stablecoins the solution? Available at SSRN 3444847, 2019.
- [16] Vitalik Buterin. Ethereum: platform review. *Opportunities and Challenges for Private and Consortium Blockchains*, 45, 2016.
- [17] Vitalik Buterin et al. A next-generation smart contract and decentralized application platform. *white paper*, 3(37):2–1, 2014.
- [18] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *OSDI*, volume 8, pages 209–224, 2008.
- [19] Huashan Chen, Marcus Pendleton, Laurent Njilla, and Shouhuai Xu. A survey on ethereum systems security: Vulnerabilities, attacks, and defenses. *ACM Computing Surveys (CSUR)*, 53(3):1–43, 2020.
- [20] Jiachi Chen, Xin Xia, David Lo, and John Grundy. Why do smart contracts self-destruct? investigating the self-destruct function on ethereum. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(2):1–37, 2021.
- [21] Ting Chen, Zihao Li, Yufei Zhang, Xiapu Luo, Ang Chen, Kun Yang, Bin Hu, Tong Zhu, Shifang Deng, Teng Hu, et al. Dataether: Data exploration framework for ethereum. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 1369–1380. IEEE, 2019.
- [22] CoinMarketCap. Bnb (binance coin) statistics. <https://coinmarketcap.com/currencies/bnb/>, 2023.
- [23] Compound. Compound protocol documentation. <https://docs.compound.finance/>, 2023.
- [24] Filippo Contro, Marco Crosara, Mariano Ceccato, and Mila Dalla Preda. Ethersolve: Computing an accurate control-flow graph from ethereum bytecode. In *2021 IEEE/ACM 29th International Conference on Program Comprehension (ICPC)*, pages 127–137. IEEE, 2021.
- [25] Jason Paul Cruz, Yuichi Kaji, and Naoto Yanai. Rbac-sc: Role-based access control using smart contract. *Ieee Access*, 6:12240–12251, 2018.
- [26] Chris Dannen. *Introducing Ethereum and solidity*, volume 1. Springer, 2017.
- [27] Rosie Dunford, Quanrong Su, and Ekraj Tamang. The pareto principle. 2014.
- [28] Ethereum. Ethereum ide and tools for the web. <http://remix.ethereum.org/>, 2020.
- [29] Josselin Feist, Gustavo Grieco, and Alex Groce. Slither: a static analysis framework for smart contracts. In *2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WET-SEB)*, pages 8–15. IEEE, 2019.
- [30] Convex Finance. Convex finance official documentation. <https://docs.convexfinance.com/convexfinance/>, 2023.
- [31] Joel Frank, Cornelius Aschermann, and Thorsten Holz. {ETHBMC}: A bounded model checker for smart contracts. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2757–2774, 2020.
- [32] Michael Fröwis, Andreas Fuchs, and Rainer Böhme. Detecting token systems on ethereum. In *Financial Cryptography and Data Security: 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18–22, 2019, Revised Selected Papers 23*, pages 93–112. Springer, 2019.

- [33] Jianbo Gao, Han Liu, Chao Liu, Qingshan Li, Zhi Guan, and Zhong Chen. Easyflow: Keep ethereum away from overflow. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*, pages 23–26. IEEE, 2019.
- [34] Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. Gigahorse: thorough, declarative decompilation of smart contracts. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pages 1176–1186. IEEE, 2019.
- [35] Neville Grech, Michael Kong, Anton Jurisevic, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. Madmax: Surviving out-of-gas conditions in ethereum smart contracts. *Proceedings of the ACM on Programming Languages*, 2(OOPSLA):1–27, 2018.
- [36] Ilya Grishchenko, Matteo Maffei, and Clara Schneidewind. Ethertrust: Sound static analysis of ethereum bytecode. *Technische Universität Wien, Tech. Rep.*, pages 1–41, 2018.
- [37] Fabio Gritti, Nicola Ruaro, Robert McLaughlin, Priyanka Bose, Dipanjan Das, Ilya Grishchenko, Christopher Kruegel, and Giovanni Vigna. Confusum contractum: confused deputy vulnerabilities in ethereum smart contracts. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1793–1810, 2023.
- [38] Ningyu He, Lei Wu, Haoyu Wang, Yao Guo, and Xuxian Jiang. Characterizing code clones in the ethereum smart contract ecosystem. In *Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers 24*, pages 654–675. Springer, 2020.
- [39] Ningyu He, Ruiyi Zhang, Haoyu Wang, Lei Wu, Xiapu Luo, Yao Guo, Ting Yu, and Xuxian Jiang. {EOSAFE}: security analysis of {EOSIO} smart contracts. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1271–1288, 2021.
- [40] Everett Hildenbrandt, Manasvi Saxena, Nishant Rodrigues, Xiaoran Zhu, Philip Daian, Dwight Guth, Brandon Moore, Daejun Park, Yi Zhang, Andrei Stefanescu, et al. Kevm: A complete formal semantics of the ethereum virtual machine. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pages 204–217. IEEE, 2018.
- [41] Yoichi Hirai. Defining the ethereum virtual machine for interactive theorem provers. In *Financial Cryptography and Data Security: FC 2017 International Workshops, WAHC, BITCOIN, VOTING, WTSC, and TA, Sliema, Malta, April 7, 2017, Revised Selected Papers 21*, pages 520–535. Springer, 2017.
- [42] Bo Jiang, Ye Liu, and Wing Kwong Chan. Contract-fuzzer: Fuzzing smart contracts for vulnerability detection. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, pages 259–269, 2018.
- [43] Priyanka Kamboj, Shivang Khare, and Sujata Pal. User authentication using blockchain based smart contract in role-based access control. *Peer-to-Peer Networking and Applications*, 14(5):2961–2976, 2021.
- [44] Satpal Singh Kushwaha, Sandeep Joshi, Dilbag Singh, Manjit Kaur, and Heung-No Lee. Ethereum smart contract analysis tools: A systematic review. *IEEE Access*, 10:57037–57062, 2022.
- [45] Satpal Singh Kushwaha, Sandeep Joshi, Dilbag Singh, Manjit Kaur, and Heung-No Lee. Systematic review of security vulnerabilities in ethereum blockchain smart contract. *IEEE Access*, 10:6605–6621, 2022.
- [46] Xiaoqi Li et al. Hybrid analysis of smart contracts and malicious behaviors in ethereum. 2021.
- [47] Haojun Liu, Xinbo Luo, Hongrui Liu, and Xubo Xia. Merkle tree: A fundamental component of blockchains. In *2021 International Conference on Electronic Information Engineering and Computer Science (EIECS)*, pages 556–561. IEEE, 2021.
- [48] Jiageng Liu, Igor Makarov, and Antoinette Schoar. Anatomy of a run: The terra luna crash. Technical report, National Bureau of Economic Research, 2023.
- [49] Ye Liu, Yi Li, Shang-Wei Lin, and Cyrille Artho. Finding permission bugs in smart contracts with role mining. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 716–727, 2022.
- [50] DeFi Llama. Ethereum defi statistics. <https://defillama.com/chain/Ethereum>, 2023.
- [51] DeFi Llama. Hacks. <https://defillama.com/?tv1=true>, 2023.
- [52] Fuchen Ma, Ying Fu, Meng Ren, Mingzhe Wang, Yu Jiang, Kaixiang Zhang, Huizhong Li, and Xiang Shi. Evm: From offline detection to online reinforcement for ethereum virtual machine. In *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 554–558. IEEE, 2019.
- [53] Fuchen Ma, Meng Ren, Ying Fu, Mingzhe Wang, Huizhong Li, Houbing Song, and Yu Jiang. Security reinforcement for ethereum virtual machine. *Information Processing & Management*, 58(4):102565, 2021.

- [54] Fuchen Ma, Meng Ren, Lerong Ouyang, Yuanliang Chen, Juan Zhu, Ting Chen, Yingli Zheng, Xiao Dai, Yu Jiang, and Jiaguang Sun. Pied-piper: Revealing the backdoor threats in ethereum ERC token contracts. *ACM Transactions on Software Engineering and Methodology*, 32(3):1–24, 2023.
- [55] Fuchen Ma, Zhenyang Xu, Meng Ren, Zijing Yin, Yuanliang Chen, Lei Qiao, Bin Gu, Huizhong Li, Yu Jiang, and Jiaguang Sun. Pluto: Exposing vulnerabilities in inter-contract scenarios. *IEEE Transactions on Software Engineering*, 48(11):4380–4396, 2021.
- [56] Pengxiang Ma, Ningyu He, Yuhua Huang, Haoyu Wang, and Xiapu Luo. Abusing the ethereum smart contract verification services for fun and profit. *arXiv preprint arXiv:2307.00549*, 2023.
- [57] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, 2008.
- [58] Robert Norvill, Beltran Borja Fiz Pontiveros, Radu State, Irfan Awan, and Andrea Cullen. Automated labeling of unknown contracts in ethereum. In *2017 26th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–6. IEEE, 2017.
- [59] Gustavo A Oliva, Ahmed E Hassan, and Zhen Ming Jiang. An exploratory study of smart contracts in the ethereum blockchain platform. *Empirical Software Engineering*, 25:1864–1904, 2020.
- [60] OpenZeppelin. ERC20 tokens. <https://docs.openzeppelin.com/contracts/4.x/erc20>, 2023.
- [61] OpenZeppelin. Openzeppelin contracts: Access control. <https://github.com/OpenZeppelin/openzeppelin-contracts/tree/master/contracts/access>, 2023.
- [62] PeckShieldAlert. Tweet by peckshieldalert. <https://twitter.com/PeckShieldAlert/status/1483363515411099651>, 2023.
- [63] Michael Rodler, Wenting Li, Ghassan O Karame, and Lucas Davi. Sereum: Protecting existing smart contracts against re-entrancy attacks. *arXiv preprint arXiv:1812.05934*, 2018.
- [64] Sina Shahab and Zaheer Allam. Reducing transaction costs of tradable permit schemes using blockchain smart contracts. *Growth and Change*, 51(1):302–308, 2020.
- [65] SlowMist. Slowmist 2022 blockchain security and aml analysis annual report. *Medium*, 2022.
- [66] Para Space. Para space official documentation. <https://docs.para.space/para-space/>, 2023.
- [67] Mythril Team. Mythril: Security analysis tool for ethereum smart contracts. <https://github.com/ConsenSys/mythril>, 2018.
- [68] Huang Teng, Wayne Tian, Haocheng Wang, and Zhiyuan Yang. Applications of the decentralized finance (defi) on the ethereum. In *2022 IEEE Asia-Pacific Conference on Image Processing, Electronics and Computers (IPEC)*, pages 573–578. IEEE, 2022.
- [69] Sergei Tikhomirov, Ekaterina Voskresenskaya, Ivan Ivanitskiy, Ramil Takhaviev, Evgeny Marchenko, and Yaroslav Alexandrov. Smartcheck: Static analysis of ethereum smart contracts. In *Proceedings of the 1st international workshop on emerging trends in software engineering for blockchain*, pages 9–16, 2018.
- [70] Christof Ferreira Torres, Julian Schütte, and Radu State. Osiris: Hunting for integer bugs in ethereum smart contracts. In *Proceedings of the 34th annual computer security applications conference*, pages 664–676, 2018.
- [71] Uniswap. Uniswap official documentation. <https://docs.uniswap.org/>, 2023.
- [72] Maximilian Wohrer and Uwe Zdun. Smart contracts: security patterns in the ethereum ecosystem and solidity. In *2018 International Workshop on Blockchain Oriented Software Engineering (IWBOSE)*, pages 2–8. IEEE, 2018.
- [73] Pengcheng Xia, Haoyu Wang, Bingyu Gao, Weihang Su, Zhou Yu, Xiapu Luo, Chao Zhang, Xusheng Xiao, and Guoai Xu. Trade or trick? detecting and characterizing scam tokens on uniswap decentralized exchange. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 5(3):1–26, 2021.
- [74] Jiahua Xu and Nikhil Vadgama. From banks to defi: the evolution of the lending market. *Enabling the Internet of Value: How Blockchain Connects Global Businesses*, pages 53–66, 2022.
- [75] Jiaming Ye, Mingliang Ma, Yun Lin, Lei Ma, Yinxing Xue, and Jianjun Zhao. Vulpedia: Detecting vulnerable ethereum smart contracts via abstracted vulnerability signatures. *Journal of Systems and Software*, 192:111410, 2022.
- [76] Mengya Zhang, Xiaokuan Zhang, Yinqian Zhang, and Zhiqiang Lin. {TXSPECTOR}: Uncovering attacks in ethereum from transactions. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2775–2792, 2020.
- [77] Yuanyu Zhang, Shoji Kasahara, Yulong Shen, Xiaohong Jiang, and Jianxiang Wan. Smart contract-based access control for the internet of things. *IEEE Internet of Things Journal*, 6(2):1594–1605, 2018.