

Fuzz on the Beach: Fuzzing Solana Smart Contracts

Sven Smolka
University of Duisburg-Essen
Essen, Germany
sven.smolka@uni-due.de

Jens-Rene Giesen
University of Duisburg-Essen
Essen, Germany
jens-rene.giesen@uni-due.de

Pascal Winkler
University of Duisburg-Essen
Essen, Germany
pascal.winkler@uni-due.de

Oussama Draissi
University of Duisburg-Essen
Essen, Germany
oussama.draissi@uni-due.de

Lucas Davi
University of Duisburg-Essen
Essen, Germany
lucas.davi@uni-due.de

Ghassan Karame
Ruhr University Bochum
Bochum, Germany
ghassan@karame.org

Klaus Pohl
University of Duisburg-Essen
Essen, Germany
klaus.pohl@uni-due.de

ABSTRACT

Solana has quickly emerged as a popular platform for building decentralized applications (DApps), such as marketplaces for non-fungible tokens (NFTs). A key reason for its success are Solana's low transaction fees and high performance, which is achieved in part due to its stateless programming model. Although the literature features extensive tooling support for smart contract security, current solutions are largely tailored for the Ethereum Virtual Machine. Unfortunately, the very stateless nature of Solana's execution environment introduces novel attack patterns specific to Solana requiring a rethinking for building vulnerability analysis methods.

In this paper, we address this gap and propose FUZZDELSOL, the first binary-only coverage-guided fuzzing architecture for Solana smart contracts. FUZZDELSOL faithfully models runtime specifics such as smart contract interactions. Moreover, since source code is not available for the large majority of Solana contracts, FUZZDELSOL operates on the contract's binary code. Hence, due to the lack of semantic information, we carefully extracted low-level program and state information to develop a diverse set of bug oracles covering all major bug classes in Solana. Our extensive evaluation on 6049 smart contracts shows that FUZZDELSOL's bug oracles finds impactful vulnerabilities with a high precision and recall. To the best of our knowledge, this is the largest evaluation of the security landscape on the Solana mainnet.

CCS CONCEPTS

• Security and privacy → Distributed systems security; Domain-specific security and privacy architectures; • Software and its engineering → Software safety.

KEYWORDS

Blockchain Security, Solana, Fuzzing

ACM Reference Format:

Sven Smolka, Jens-Rene Giesen, Pascal Winkler, Oussama Draissi, Lucas Davi, Ghassan Karame, and Klaus Pohl. 2023. Fuzz on the Beach: Fuzzing Solana Smart Contracts. In *Proceedings of Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3576915.3623178>

1 INTRODUCTION

Smart contracts are an essential part of the ecosystem in many modern blockchain platforms. Smart contracts allow developers to implement decentralized applications (DApps) that encode business logic on the blockchain, thereby facilitating numerous use cases. For instance, smart contracts enable the creation of non-fungible token (NFT) marketplaces. Artists use NFT marketplaces to auction their creations. Furthermore, established companies and sport franchises, like Nike [38], Budweiser [5], Lacoste [26], and the NBA [34] use these marketplaces to sell NFT collections to fans and investors all over the world.

The Solana blockchain [14] has become a key platform in the DApps and NFT space, because of its high-performance and low transaction fees. In comparison to the more established smart contract platform Ethereum [13], Solana can execute 100–1000 times more transactions per second [27, 45] while charging a fraction of a USD cent as a fee [14, 27]. As a result, the number of all transactions in the Solana network significantly exceeds the number of all transactions made in Ethereum by a factor of 85¹.

From a smart contract perspective, the Solana platform achieves a high transaction rate because its execution layer decouples program logic from state, i.e., smart contracts cannot store any dynamic state. This enables Solana to execute transactions operating on different data in parallel. However, this also introduces new attack patterns that are specific to Solana. In fact, attacks against Solana

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

CCS '23, November 26–30, 2023, Copenhagen, Denmark

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0050-7/23/11...\$15.00

<https://doi.org/10.1145/3576915.3623178>

¹On April 12, 2023, Ethereum processed 1933 million (<https://etherscan.io>) transactions, while Solana already reached 164 839 million transactions (<https://solscan.io>).

smart contracts already caused multi-million Dollar losses [53, 54]—the popular Wormhole attack induced losses of up to 320 million USD [18].

Smart contract security research ranges across different disciplines: from formal verification [47] and static analysis [33, 50] to dynamic analysis [8], with a high focus on the Ethereum platform. However, Solana suffers from different vulnerabilities than Ethereum, and the aforementioned techniques are not applicable to Solana due to its unique features: In comparison to Ethereum, the Solana blockchain is stateless and smart contracts have no direct association with the state. The stateless nature of Solana’s execution environment requires stricter handling of user input. However, vulnerabilities often come from developers not checking security-critical properties in smart contracts, like missing transaction *signer checks*. Unlike Ethereum smart contracts, which implicitly trust their state if it is not compromised. Moreover, the stateless approach of Solana and the impact on its security model is largely unexplored.

Research on Solana security and tooling is limited: At the time of writing, VRust [7] is the only existing static analysis approach that covers Solana smart contracts. VRust incorporates detection patterns for common vulnerabilities in Solana smart contracts and was able to detect 12 vulnerabilities in popular open-source smart contracts [7]. However, VRust suffers from several limitations: 1) it strictly requires source code to conduct analyses, 2) it suffers from a high number of false alarms, and 3) it does not provide an analyst with enough data to (re-)construct exploit transactions. In contrast, fuzzing is a technique that does not suffer from any of these limitations [3, 11, 12, 41]. The fuzzing input given to the analysis target can also usually be crafted into exploit transactions. Smart contract fuzzing is a valuable technique that has been extensively researched with promising results. [19, 37, 42].

Contributions. In this work, we propose a solution to detect bugs in Solana called FUZZDELSOL: the *first* binary-only coverage-guided fuzzer for Solana smart contracts. We developed a set of bug detection oracles to facilitate the detection of Solana-specific smart contract bugs, namely 1) missing signer checks, that is, the smart contract performing critical operations without checking for signatures, 2) missing owner checks, which allow a smart contract to use untrusted data, 3) arbitrary cross program invocation, i.e., a smart contract calls *any* other smart contract, 4) missing key checks, which, similar to the missing owner check, enables a smart contract to use spoofed accounts as system variables, and 5) integer bugs. In addition, we also design a generic oracle for FUZZDELSOL to detect vulnerabilities based on lamport gains. Lamports are the smallest denomination of Solana’s native currency SOL, and 1 SOL corresponds to 1 000 000 000 (one billion) lamports. We use this oracle to detect arbitrary leaking funds and lamport-theft.

Our extensive evaluation of FUZZDELSOL consists of several experiments. In our first experiment, we test FUZZDELSOL with a set of vulnerable smart contracts provided by the community [35]. We demonstrate that our approach is capable of quickly detecting Solana-specific vulnerabilities. Compared to VRust, FUZZDELSOL does not report any false alarms for the dataset provided in [35]. In addition, FUZZDELSOL can also precisely trace back the vulnerability classes.

Next, we perform a large-scale bug-finding evaluation on all Solana smart contracts present on the mainnet on March 27, 2023, which amounts to a total of 6049 smart contracts. At the time of writing, this is the largest analyzed dataset of Solana smart contracts. FUZZDELSOL reports 92 bugs in these smart contracts. We analyzed 16 reports in-depth and confirmed that only 2 were false alarms, thus demonstrating the high accuracy of FUZZDELSOL in detecting bugs. FUZZDELSOL is the only analysis tool available that can analyze these contracts on the mainnet.

Third, our performance evaluation on 16 smart contracts from well-known bug bounty programs demonstrates that FUZZDELSOL can analyze complex smart contracts. Here, FUZZDELSOL’s generated transactions can consistently find new code paths in the programs. In this experiment, FUZZDELSOL reports a true bug and only has a single false alarm, i.e., a single wrongly reported vulnerability. We summarize our contributions as follows:

- We present the *first* fuzzing architecture for Solana smart contracts. We conceptualize FUZZDELSOL (Section 4) around the original Solana runtime to FUZZDELSOL faithfully model runtime specifics, such as smart contract interaction. Moreover, this design guarantees reproducibility and validity of transactions that FUZZDELSOL generates: Every transaction that generates a vulnerability report can be replayed, e.g., on a test network.
- We design and implement new bug oracles to detect vulnerabilities in Solana programs (Section 3.1). Due to our design choices, FUZZDELSOL detects impactful bugs in smart contracts regardless of source code availability.
- Our extensive evaluation on 6049 smart contracts shows that FUZZDELSOL’s bug oracles find bugs with a high precision and recall. This is the largest evaluation of the security landscape on the Solana mainnet.
- FUZZDELSOL detects the infamous Wormhole bug.

2 SOLANA’S EXECUTION ENVIRONMENT

In the following, we provide an overview of Solana’s execution environment [15, 16].

Solana Account Model. The Solana account model decouples accounts containing non-executable raw data from accounts containing executable code. In order to achieve this decoupling, Solana introduces an account layout consisting of the following six fields: (1) the public key of the account, (2) the public key of the account owner, (3) the executable flag, which indicates whether an account is executable (and accordingly a program), (4) the rent epoch, which specifies a point in time when the account must pay rent to remain deployed on the blockchain, (5) the funds that the account holds, in a unit called *lamports*, and (6) the data that the account contains.

Smart contracts in Solana are called *programs*. Executable on-chain programs contain extended Berkeley Packet Filter (eBPF) bytecode compiled as an executable and linkable format (ELF) shared object file in their data field and are stateless, i.e., they do not store runtime-modifiable data in their data field. Programs manage runtime-modifiable data in non-executable accounts whose owner field contains the program’s public key. Once a program is the owner of an account, only that program is able to modify the account’s data as well as deduct lamports from the account.




Figure 1: Serialization of an instruction as input of the eBPF VM

Solana distinguishes between two types of programs: *i) native programs* and *ii) on-chain programs*. Native programs are implemented in the Solana runtime and are not deployed on the blockchain. These programs typically perform tasks such as allocating new accounts on the blockchain or deploying on-chain programs on the blockchain. On-chain programs are written in C, C++, or Rust, compiled in eBPF, and deployed on the blockchain using a native program. Both native and on-chain programs are marked as executable. But only on-chain programs contain eBPF bytecode in their data field. Moreover, on-chain programs are only capable of modifying the data and lamports of non-executable accounts at runtime. However, some native programs are also capable of modifying the remaining four fields.

Solana Transactions. Transactions consist of, but are not limited to, (1) a list of signatures of accounts that signed the transaction (2) a recent blockhash used to determine if the transaction is too old, (3) a sorted list of accounts that can be used in the instructions of the transaction (4) and a list of *instructions*. An *instruction* is responsible for invoking a native or an on-chain program. Here, a native program is executed directly in the Solana runtime, while an on-chain program is executed in Solana’s eBPF VM using the eBPF bytecode stored in the program’s data field. Instructions consist of three elements: *i) the public key of the program called by the instruction* (also referred to as the *program id*), *ii) a list of accounts passed to the program*, which must be a subset of the sorted list of accounts in the transaction, and *iii) instruction data representing arbitrary data*. A called program accesses these three elements during execution.

Cluster Information. Solana programs receive cluster information by calling functions of *sysvar* accounts. For instance, the amount of lamports to pay for allocating one data byte for an account. According to the Solana programming model, *sysvar* accounts must also be passed as input to the program by including them in the instruction’s account list.

Program-Derived Addresses. In addition to addresses that are a pair of public and private keys, Solana introduces Program-Derived Addresses (PDAs). PDAs are addresses that are not located on the ed25519 curve, do not have a corresponding private key, and are associated with a program. A PDA receives its association with a program during its derivation, where the PDA is deterministically derived based on the program’s public key and a set of optional seeds. In order to ensure that a PDA does not lie on the ed25519 curve, a bump byte is iteratively determined that will uniquely “bump” the PDA out of the ed25519 curve.

Cross-Program Invocation. Solana programs can call other programs using Cross-Program Invocation (CPI). Here, the caller invokes the callee with a self-created instruction that contains at least the same privileges as the instruction which invoked the caller, i.e., accounts that signed the transaction. However, the Solana runtime allows the caller to delegate additional privileges to the callee using PDAs. These additional privileges are restricted to accounts whose public key is a PDA associated with the caller, i.e., derived from the caller’s public key along with a set of optional seeds. Thus, when invoking a program using CPI, the calling program can use PDAs to “sign” accounts in the CPI instruction.

eBPF VM. Solana’s eBPF VM [16] can execute on-chain programs in both *Just-in-time* compilation and *Interpreter* mode. The Solana runtime defines a number of environmental restrictions when executing on-chain programs in the eBPF VM. By default, the eBPF VM limits the resource consumption of an instruction to a maximum of a certain number of *compute units*. The Solana runtime accumulates compute units for all instructions within a transaction, with certain runtime operations such as function or system calls consuming a specific number of compute units. In addition, the Solana runtime enables the eBPF VM to reach a maximum depth of four Cross-Program Invocations. Lastly, it prohibits reentrancy during CPI.

When executing an instruction in the eBPF VM, the instruction is serialized and passed to the VM as input for the program. The program input starts at a fixed address of the VM memory layout. Figure 1 depicts the layout of the serialized instruction, which is divided into three parts: accounts, instruction data, and program id which corresponds to the input of the program. The first 8 bytes of each account contain information about whether the account has signed the transaction, is read-only, or is executable. Followed by the public key of the account and the account’s owner, the amount of lamports, the data length, and other information.

3 SOLANA PROGRAM SECURITY AND CHALLENGES

There are different types of vulnerabilities in Solana programs. Attackers exploiting these vulnerabilities may compromise accounts managed by vulnerable programs, i.e., by stealing an account’s funds or manipulating an account’s data. The attacker typically accomplishes this by crafting and executing sequences of instructions that allow the attacker to maliciously select the order of accounts and the instruction data in each instruction. As a result, an attacker gains access to accounts owned by or authorized to the vulnerable program, and can steal lamports or violate the accounts’ data integrity.

3.1 Solana Program Vulnerabilities

In the following, we describe five key vulnerabilities [7, 18, 35] of Solana programs, which we all address in this paper.

Missing Signer Check. A *missing signer check* (MSC) vulnerability exists when a program does not verify that an account that should have signed the transaction according to the business logic actually signed the transaction. This allows an attacker to gain unauthorized access to program behavior.

Missing Owner Check. A *missing owner check* (MOC) vulnerability exists when a program reads and processes data from an account that, according to the business logic, should be owned by the program without verifying it is actually owned by the program. Thus, an attacker could create his own account with data and pass it to the program as input.

Arbitrary CPI. An *arbitrary CPI* (ACPI) vulnerability exists when a program does not verify the program id of the invoked program during CPI. Hence, an attacker can deploy a malicious program on the blockchain and pass it as input to the program possessing the arbitrary CPI vulnerability. As a result, the program possessing the vulnerability invokes the malicious program using CPI, giving the attacker control over the execution. This is especially critical if the program grants additional privileges to the invoked program by signing PDAs.

Missing Key Check. A *missing key check* (MKC) vulnerability exists when a program expects a specific account and processes its data without verifying that the passed account is actually the expected specific account. This vulnerability resulted in a loss of up to 320 million USD in the Wormhole program [18].

Integer Bugs. An *integer bug* (IB) exists when values underflow or overflow in arithmetic operations. This vulnerability can be exploited by an attacker when a program transfers lamports from account a to account b . Here, the attacker selects the value to be transferred so high that the lamports field of account b overflows, while the lamports field of account a underflows. The Solana runtime allows this transfer because the total amount of lamports remains the same before and after executing the instruction.

3.2 Solana Security Analysis

In contrast to the available tools and approaches for Ethereum smart contract security analysis, only very few approaches and tools support Solana program security analysis. The existing ecosystem merely consists of bytecode lifters [30, 39] that are capable to lift the eBPF bytecode of Solana programs to another, tool-specific intermediate language (IL), such as Ghidra's [2] or Binary Ninja's [1] IL. These plugins help in reverse engineering, but do not conduct any sophisticated security analysis on their own.

Currently, VRust [7] is the only existing static analysis approach that covers Solana programs with a focus on security. VRust covers a wide range of vulnerability patterns for common vulnerabilities in Solana smart contracts. As a result, VRust was able to detect 12 vulnerabilities in popular open-source smart contracts [7]. VRust uses the Rust compiler to analyze Solana programs on the Rust Mid-Level Intermediate Representation (MIR) [32]. The MIR is a typed language that has complete information about the memory layout and types of a variable's data structure.

The Rust compiler uses MIR as a compiler-internal representation that models control-flows and data-flows for further optimization of the ownership and borrow checking rules, hence it is designed to model control-flows and data-flows precisely. In general, it is possible to map MIR-level objects backwards to higher-level representations, which also provides access to high-level information like variable names. As a result, today's decompiling and lifting approaches [30, 39] cannot achieve the same precision for data-flows as MIR. Further, custom analysis passes can be implemented as a

visitor on top of the MIR-level control flow graph for any given Solana program written in the Rust programming language. However, relying on the MIR for vulnerability analysis strictly requires source code, because the MIR can only be generated by the Rust compiler from the higher-level representations.

Solana programs can be written in C or Rust, and it is possible to differentiate between programs written in either language by checking which system calls a program uses. As this can be done on bytecode level, we count the occurrences of the C vs. Rust variants of the `sol_invoke_signed` system call. We find that 97% of Solana programs are written in the Rust programming language, whereas the remaining programs use the C programming language. Thus, VRust cannot analyze these programs. The *Solana security.txt* [36] feature allows program authors and developers to provide information on where to find the source code of a program and whom to contact for security issues. We use this information to find out about source code availability of Solana programs. Our findings show that less than 2% of Solana programs come with source code. As a result, VRust cannot be applied to a vast majority of Solana programs.

3.3 Challenges of Fuzzing Solana Programs

A popular approach to uncover bugs is coverage-guided fuzzing [11, 12, 48, 57]. This technique mutates the inputs based on instruction-coverage data, or feedback information, collected during the target's execution, to uncover new paths in the application. Fuzzing Solana contracts is yet to be explored.

Ethereum Smart Contract Fuzzing. Fuzzing Ethereum smart contracts is a heavily researched area [23, 25, 37, 42, 51, 55]. Ethereum fuzzers like EFCF [42] model interaction between smart contracts to detect complex and hard-to-find Ethereum bugs, e.g., compositional reentrancy bugs. However, these approaches are not feasible or applicable to Solana, because of its unique programming model (cf. Section 2). Modelling the interaction between Ethereum smart contracts does not require in-depth information about the type of other smart contracts. In order to model the interaction between programs and accounts in Solana, it is necessary to know the domain of a program. For instance, the program in Listing 1 interacts with three different types of accounts: a wallet, a vault, and an authority. Ethereum encapsulates the state within the smart contract, i.e., an Ethereum fuzzer [25] does not need to consider the different types of data to detect bugs. Nevertheless, to detect more complex bugs, like delegated re-entrancy bugs, Ethereum fuzzers [42] must understand the semantics of the respective contracts. In contrast, it is hard to detect which type an account assumes in a Solana program. Therefore, modelling the blockchain state and the content of accounts is essential to detect real, reproducible and impactful bugs in smart contracts. Thus, to fuzz Solana programs, a Solana fuzzer has to solve the following challenges:

Challenge 1: Modeling Ledger Snapshots. As Solana requires programs to store data in external, non-executable accounts, it is necessary to model a valid ledger snapshot which consists of multiple accounts. Moreover, programs must be able to change this snapshot across multiple transactions as well as operate on the changed snapshots to be able to execute business logic which depends on a specific state of the ledger snapshot.

Challenge 2: Reproducibility of Transactions. To identify vulnerabilities exploitable in the real production blockchain, it is necessary to generate transactions that are reproducible in the production blockchain.

Challenge 3: Cluster Information. Several Solana programs require cluster information at runtime, such as the amount of lamports to allocate a byte of data. Solana programs receive this cluster information using functions provided by *sysvar* accounts, which must be passed to the program as input. Hence, it is necessary to ensure that Solana programs receive *sysvar* accounts as input and can call functions of these *sysvar* accounts.

Challenge 4: Solana-specific Vulnerabilities. Since Solana programs comprise specific vulnerabilities resulting from the Solana programming model, it is necessary to develop detection mechanisms to detect these Solana-specific vulnerabilities.

Challenge 5: Program semantics need to be retrieved. Solana programs manage accounts associated with the program using PDAs with a program-specific seed structure. Since programs verify whether the public key of a passed account corresponds to the program-specific PDA seed structure, it is necessary to *i)* determine the program-specific PDA seed structure and *ii)* to pass accounts to the program, whose public key is derived from the program-specific PDA seed structure.

Challenge 6: Faithful CPI. Solana programs can invoke other programs at runtime using CPI. Hence, the runtime environment must ensure that programs can use the CPI mechanism to invoke both on-chain programs—in a separate eBPF VM—and native programs in the Solana runtime.

4 OVERVIEW OF FUZZDELSOL

In this section, we introduce the design of FUZZDELSOL and its main components. Furthermore, we describe how our design tackles the challenges mentioned in Section 3.3.

Intended Use of FUZZDELSOL. FuzzDelSol aims to explore the prevalence of vulnerabilities in Solana programs. Although similar bytecode-based security studies have been conducted for other blockchain platforms, such as Ethereum [17, 47], or other application domains like Android [9], there does not yet exist any comprehensive study about the security of Solana programs. Hence, for the first time, we aim to raise awareness for Solana program security with FUZZDELSOL. For 98% of Solana programs, no source code is available (see Section 3.2). Hence, we argue that source code-based analysis techniques like VRust [7] are not applicable to analyze the vast majority of Solana programs. Moreover, FUZZDELSOL can be used to find vulnerabilities with the intention of forming a better understanding of the vulnerability types. This is necessary to develop appropriate countermeasures for Solana-specific vulnerabilities. Further, Solana program developers may use FUZZDELSOL to vet closed-source third-party programs interacting with their own programs. The same applies to users of Solana programs: FUZZDELSOL helps in ensuring that closed-source programs a user wants to invest in are secure, before investing funds.

Overview. Our high-level architecture is shown in Figure 2. The main idea of FUZZDELSOL is to ① create a valid blockchain snapshot using the *blockchain emulator*; comprising the Solana program to

analyze, an *attacker* account and additional accounts, e.g., user and non-executable data accounts, or executable programs. The next component is the *transaction generator* ②, which receives random and mutated bytes from a fuzzer to generate valid transactions. FUZZDELSOL executes these transactions in an instrumented Solana runtime ③ called *RunDelSol*. In particular, we extended the original Solana runtime with patches to detect Solana-specific vulnerabilities (cf. Section 3.1) induced by the generated transactions. Finally, the *transaction evaluator* ④ analyzes the aftermath of the transactions, and extracts valuable insights for following fuzzing iterations. However, if the transactions signal an erroneous ledger snapshot, FUZZDELSOL generates a vulnerability report with information to reproduce this ledger snapshot.

Blockchain Emulator (Section 5.1). A challenge for fuzzing Solana programs is that program execution is largely dependent on the ledger snapshot. Thus, we developed a component called *blockchain emulator* ① to prepare the snapshot of the ledger available for analysis. The modeled ledger snapshot contains the program being fuzzed as well as additional accounts that are relevant to the execution context. Moreover, this component provides all public keys that can be passed to the program as input. Operating on a valid ledger snapshot allows a program, for example, to manage and modify lamports and account data at runtime across multiple transactions. FUZZDELSOL uses the blockchain emulator to model a ledger snapshot for programs to operate on during program execution, thereby addressing Challenge 1. In addition, the blockchain emulator—along with RunDelSol—also enables FUZZDELSOL to address Challenge 3, since the public keys provided by the blockchain emulator include those of the *sysvar* accounts used by the program to retrieve cluster information. Furthermore, the blockchain emulator incorporates an account generator that creates attacker-controlled accounts containing malicious data to trigger Solana-specific vulnerabilities. Therefore, the blockchain emulator supports addressing Challenge 4. Finally, the blockchain emulator uses PDA seed structures obtained from the transaction evaluator to derive valid PDAs, which assists in tackling Challenge 5.

Transaction Generator (Section 5.2). For each fuzzing iteration, FUZZDELSOL *mimics* real Solana transactions to find reproducible bugs. However, Solana transactions contain structured data. Thus, FUZZDELSOL incorporates a *transaction generator* ②, which *transforms* the randomly generated bytes from a fuzzer into valid Solana transactions. As the transaction generator produces valid and reproducible transactions, FUZZDELSOL covers Challenge 2.

RunDelSol (Section 5.3). Effective fuzzing requires coverage feedback [11, 57] to guide the generation of test inputs. Source code is commonly used to instrument a program for achieving accurate results in this feedback mechanism. However, given the absence of source code for the large majority of Solana programs, we cannot rely on the source code. Therefore, instead of instrumenting programs, FUZZDELSOL implements a specialized Solana runtime environment ③, called *RunDelSol*. This environment—besides instantiating and executing Solana programs on the previously generated ledger snapshot—includes instrumentation to measure coverage. Furthermore, it allows programs to invoke functions of passed *sysvar* accounts to retrieve *cluster information* (addressing Challenge 3 along with the blockchain emulator).

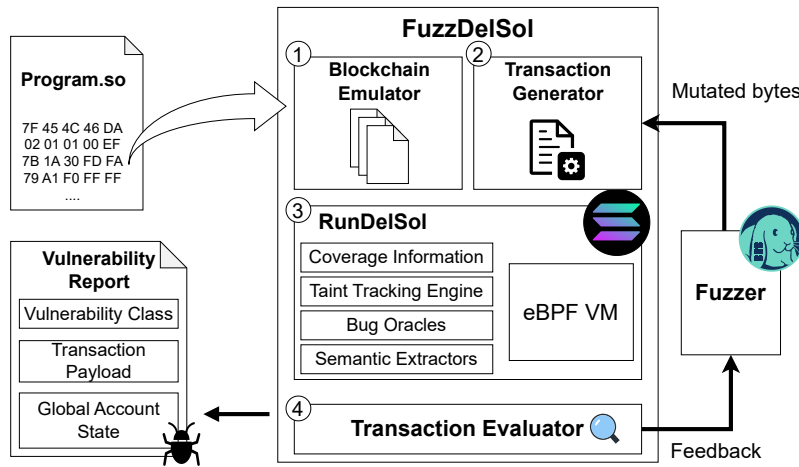


Figure 2: FUZZDELSOL Design

Moreover, RunDelSol features a *taint tracking engine* to trace the data-flow during the execution of a program. This enables us to implement Solana program vulnerability detectors or *bug oracles*, as well as Solana-specific runtime information extractors which extract PDA seed structures. Each oracle aims to detect a potential vulnerability without source code information and uses the taint tracking engine differently, and therefore tackling [Challenge 4](#). Moreover, the Solana-specific runtime information extractors of RunDelSol (along with the transaction evaluator, see below) allow FUZZDELSOL to address [Challenge 5](#) by extracting PDA seed structures.

Moreover, RunDelSol can invoke multiple programs using CPI and run them in separate eBPF VMs for on-chain programs. In the case of native programs, RunDelSol executes them in the Solana runtime. Otherwise, for invoking on-chain programs, RunDelSol traces their data-flow using the taint tracking engine independently to the callee program. This allows the oracle to detect vulnerabilities across multiple CPI invocations. Enabling programs to call other programs using CPI and analyzing their interaction allows FUZZDELSOL to solve [Challenge 6](#).

Transaction Evaluator (Section 5.4). Transactions can impact the state of the blockchain. For correct adjustment of the ledger snapshot and preparation of the next fuzzing iteration, the *transaction evaluator* ④ extracts relevant information from RunDelSol after the execution of transactions, including PDA seed structures, eBPF VM signals, and feedback information for the fuzzer, e.g., coverage. Next, the transaction evaluator forwards the information to the fuzzer and the blockchain emulator and decides whether it should re-generate the ledger snapshot for subsequent fuzzing iterations, taking into account the newly received semantic program information. As a result, the transaction evaluator, together with RunDelSol, allows addressing [Challenge 5](#).

5 FUZZDELSOL INTERNALS

In this section, we detail the implementation of FUZZDELSOL. FUZZDELSOL uses the state-of-the-art Libafl [12] fuzzer. Libafl’s design

allows the FUZZDELSOL to include its own feedback mechanism in the fuzzing mutation. The bytes generated by Libafl are based on feedback obtained from previous fuzzing iterations, which helps uncover new paths and overcome barriers, e.g., such as public key comparisons, in the Solana program.

We describe how each of the four components depicted in [Figure 2](#) is used to detect bugs. Moreover, we leverage an exemplary program that suffers from two impactful bugs to describe FuzzDelSol for ease of presentation. [Listing 1](#) shows the code of this example program. We chose this example because it showcases two of the more popular bugs in Solana. This function allows a user to withdraw funds from a *vault* managed by the program. [Line 13](#) in the function checks if the *authority* account is authorized to withdraw funds from account *vault*. In contrast to Ethereum, Solana programs need a separate wallet account to store data. In [Line 14](#), the function also checks whether the *vault* provided as an input is associated with the *wallet_info* account. Finally, if account *from* contains sufficient lamports, the transfer proceeds. However, this program suffers from a *missing signer check*: an attacker can provide account *authority* without the actual owner of *authority* knowing about this. Hence, the attacker can transfer funds on behalf of the actual *authority* without proper authorization. This vulnerability can be fixed by adding a signer check to the program, which is shown in [Line 20](#). Furthermore, the program lacks an owner check to verify the integrity of the information stored in the *wallet_info* account. In this case, an attacker can provide a forged *wallet_info* account with fake data that refers to any *vault* managed by the program while supplying his own public key as the *authority*. This means that, even if the code includes the signer check in [Line 20](#), an attacker could drain any *vault* associated to the program. The check in [Line 17](#) prevents this, as it verifies that the information in *wallet_info* is trusted. Therefore, even if an attacker passes a fake account to the program, the program recognizes that this account does not contain the program’s *program_id* as owner. Thus, this check mitigates illegal lamport withdrawal. Note that FUZZDELSOL is detecting these bugs without access to the program’s source code.

```

1 fn withdraw(program_id: Pubkey, accounts: [AccountInfo], amount: u64) ->
  ProgramResult {
2
3     // The accounts from the transaction are attacker-controlled,
4     // including wallet_info and wallet
5     let wallet_info = accounts[0];
6     let wallet = deserialize(wallet_info.data);
7
8     let vault = accounts[1];
9     let authority = accounts[2];
10
11    // An attacker can forge the authority and vault fields
12    // of the wallet account, to pass both of these checks
13    assert(wallet.authority == authority.key);
14    assert(wallet.vault == vault.key);
15
16    + // FIX: check owner of account from
17    + assert(wallet_info.owner == program_id);
18
19    + // FIX: the following line adds the required signer check
20    + assert(authority.is_signer);
21
22    // check for sufficient funds
23    if amount > from.lamports {
24        raise InsufficientFundsException;
25    }
26
27    // transfer lamports
28    vault.lamports -= amount;
29    authority.lamports += amount;
30
31    // XXX: Missing signer and owner check vulnerability
32    // Funds can be transferred to and from unauthorized accounts
33    Ok(())
34 }

```

Listing 1: Solana program to withdraw funds from a vault. This program contains a missing signer and owner check. We mark additional checks with + and green highlighting.

5.1 Blockchain Emulator

The blockchain emulator allows FUZZDELSOL to model a valid ledger snapshot. When creating a ledger snapshot, the blockchain emulator deploys several accounts: (1) a *user* and *attacker* wallet account, which are used by the oracles to determine whether a transaction triggered a potential security vulnerability (2) on-chain programs such as the program to fuzz or the *SPL Token program* (3) *sysvar* accounts, enabling programs to use their functions to receive *cluster information* (4) attacker-controlled accounts, which attempt to exploit potential missing owner check vulnerabilities.

The blockchain emulator provides all these accounts and their public keys to the transaction generator for generating transactions. The example in Listing 1 requires three accounts: the *wallet_info*, *vault* and *authority* accounts with their respective data. If these accounts are not available, the transaction fails. FUZZDELSOL generates all three accounts and deploys them on the blockchain. The blockchain emulator additionally provides the *user*'s and *attacker*'s private keys, as well as the public keys of Solana's native programs—which do not need to be deployed on the blockchain as they are integrated in the Solana runtime. Finally, it also implements a *PDA generator* and an *attacker-controlled accounts generator*, which receive information about the runtime from the transaction evaluator.

PDA Generator. The PDA generator derives PDAs based on the information provided by the PDA seed structures. The seed structures are extracted by the *PDA seed structure extractor* (cf. Section 5.3.4) and received by the transaction evaluator. The generator creates user-related and attacker-related PDAs by inserting the user's and

the attacker's public key at the precise seed positions at which the program expects a public key. Furthermore, the generator inserts statically extracted seeds in places where the PDA seed structure extractor could not determine the origin of the seed. In the case that no seed exists in a PDA seed structure that originates from a public key, the PDA generator derives a single PDA without reference to the user or attacker. Finally, the PDA generator provides each derived PDA to the transaction generator for generating transactions.

Attacker-controlled Accounts Generator. This generator uses information about the location of public keys from the account data—extracted by the *account data structure extractor* (cf. Section 5.3.4) and received by the transaction evaluator—to create attacker-controlled accounts populated with malicious data and deploys them on the ledger snapshot. Here, the generator populates the attacker-controlled account data with public keys of user-related accounts and the public key of the attacker at positions at which the program expects public keys. Hence, an extracted account structure must contain at least two public keys for the generator to generate attacker-controlled accounts. By generating attacker-controlled accounts containing malicious data, FUZZDELSOL aims to detect potential missing owner check vulnerabilities. In Listing 1, this applies to the wallet account *wallet_info* with an account *authority* by comparing *wallet_info*'s data with the public key of *authority*. Since the program does not contain an owner check, FUZZDELSOL treats the *wallet_info* account as attacker controlled. Lastly, the attacker-controlled accounts generator provides the public key of each generated attacker-controlled account to the transaction generator.

5.2 Transaction Generator

The transaction generator allows FUZZDELSOL to create valid and reproducible Solana transactions, effectively simulating real transactions. This component receives generated bytes from the fuzzer, and blockchain information from the blockchain emulator. This includes information regarding the public keys of the selectable accounts, the last blockhash, and accounts capable of signing a transaction

As each transaction consists of at least one instruction, the generator first generates instructions from the received bytes, and then includes the remaining elements of a transaction, including the signature list, blockhash, and sorted account list (see Section 2).

FUZZDELSOL transforms the generated bytes of the fuzzer into a valid Solana transaction. The transformed bytes can be mainly divided into four steps: *i*) number of accounts, *ii*) information about the account signing the transaction containing the instruction, *iii*) information about the accounts contained in the instruction, *iv*) and information about the instruction data contained in the instruction. Regarding Listing 1, in step *iii*), the blockchain emulator extracts the accounts used in the instruction, and deploys accounts for *wallet_info*, *vault*, and *authority*.

For example, the number of accounts generated for the transaction is defined by a one-byte field. This field defines how many accounts the generator should insert in the instruction, considering the maximum number of selectable accounts provided by the blockchain emulator. Similarly, the remaining bytes provided by the fuzzer are structured to create valid transactions and accounts. Here,

FuzzDelSol takes dependencies between these fields into account: For instance, the number of generated accounts directly affects the account and signer information fields. The instruction bytes information field uses the raw bytes generated by the fuzzer as instruction data. After transforming the bytes into a valid instruction, with all the required information, the transaction generator creates an associated transaction for that instruction.

5.3 RunDelSol

RunDelSol executes the previously generated transactions in the Solana runtime environment. Since RunDelSol uses the real Solana environment, FUZZDELSOL can provide cluster information to the executed Solana program to fuzz and allow invoking other programs deployed on the blockchain (on-chain programs) and integrated in the runtime (native programs) using CPI. Moreover, RunDelSol instantiates the Solana program to fuzz in an instrumented eBPF VM in Interpreter mode and executes it based on the ledger snapshot generated by the blockchain emulator. RunDelSol includes an instrumented eBPF VM, which extends the execution environment with *coverage information*, data-flow tracing with *taint tracking*, the implementation of six *bug oracles* to detect potential vulnerabilities, and the extraction of *Solana-specific program semantics*.

5.3.1 Coverage Information. The instrumented eBPF VM of RunDelSol extracts coverage information during program execution. Here, RunDelSol examines each transition in the control flow graph, i.e., each JMP, CALL, and RET eBPF-instruction, and computes an index for that transition as follows: Let src be the program counter of the eBPF JMP, CALL or RET instruction and dst be the program counter of the target instruction of the transition, then the index of the transition is $i \leftarrow (src + dst) \bmod s$, where s is the size of the *coverage array*. The indices of the *coverage array* provide information about whether the program executed this transition. After the instrumented eBPF VM terminates, RunDelSol forwards the received coverage information of the executed instruction to the transaction evaluator.

5.3.2 Taint Tracking Engine. The taint tracking engine provides fine-grained tracing of data-flows during program execution. RunDelSol instruments memory addresses and register indices used by LOAD, STORE and MOV instructions. The Data-flow tracing is enabled for the following three events. First, FUZZDELSOL traces data of accounts located on the *program input* (cf. Figure 1) and account public keys read by the program. In Listing 1 FUZZDELSOL traces the data of account *wallet_info*, because its data is read in Lines 13 and 14. FUZZDELSOL also traces the public keys of the *vault* and *authority* accounts for the same reason. Second, the taint tracking engine starts tainting memory addresses where the program stores the return values of Solana-specific syscalls for deriving PDAs. Tainting these memory addresses is crucial, as the program can use them in its execution instead of the public keys of the accounts located on the program input. Third, we taint register indices where the program stores an overflow or underflow value as a result of an arithmetic eBPF-instruction.

5.3.3 Bug Oracles. RunDelSol implements six bug oracles in the instrumented eBPF VM, to find the following vulnerabilities: *i*) missing signer check, *ii*) missing owner check, *iii*) arbitrary cross

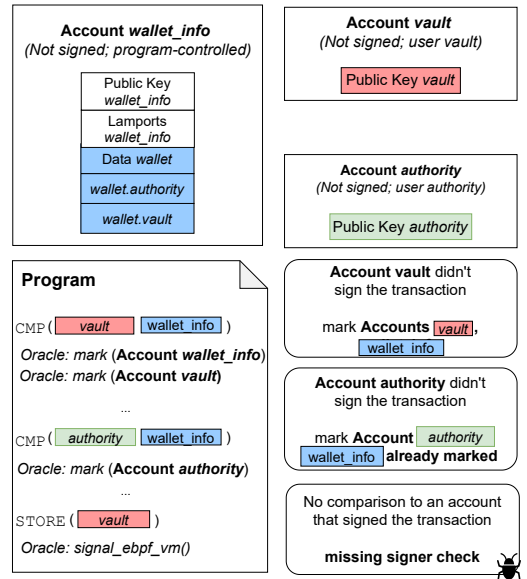


Figure 3: Interaction of the taint tracking engine and the missing signer check oracle

program invocation, *iv*) missing key check, *v*) integer bugs, *vi*), and a lamports-based bug oracle to detect lampion-theft. These oracles investigate eBPF-instructions, leverage the taint tracking engine to trace data-flows between eBPF-instructions, and signal the instrumented eBPF VM if a transaction triggers a potential vulnerability. When signaling the eBPF VM, each oracle specifies the reason for the crash and the program counter at which the crash occurs. In the following, we will describe how these oracles work. A more detailed presentation of our algorithms utilized in our fuzzing framework can be found in the extended version of this paper [49].

Missing Signer Check Oracle. Listing 1 suffers from a missing signer check. FUZZDELSOL traces and analyzes the data-flow of this program to detect this vulnerability. Figure 3 describes this process: The program reads data from the *wallet_info*, and public keys from the *vault* and *authority* accounts. Algorithm 1 details the bug-detecting process for a missing signer check. The oracle keeps track of mainly two things: accounts that the oracle marks as *vulnerable*, and accounts that are marked as *not vulnerable*. Further, the oracle marks accounts as vulnerable if their public key or data are compared with another account without any of the accounts having signed the transaction (see Line 10). On the other hand, the oracle marks accounts as *not vulnerable* if a comparison against an account that signed the transaction happens (see Line 14). The oracle signals the VM if the program writes to the lamports or data fields of an account that is marked as *vulnerable* (see Line 21 to 23).

Regarding the example in Listing 1, because of the checks in Lines 13 and 14, FUZZDELSOL marks the *wallet_info* account as *vulnerable*. Similarly, FUZZDELSOL marks the *vault* and *authority* accounts because their public keys are compared to the data of another account which also did not sign the transaction. The oracle subsequently detects that the program writes a decreased amount

Algorithm 1 Missing Signer Check Oracle

```

1:  $k \leftarrow$  amount of accounts
2:  $V \leftarrow [\emptyset_1, \emptyset_2, \dots, \emptyset_k]$ 
3:  $T \leftarrow \emptyset$ 
4: while program running do
5:   if program performs register comparison then
6:      $a \leftarrow$  account in src register
7:      $b \leftarrow$  account in dst register
8:      $i \leftarrow$  index of account  $a$  in  $V$ 
9:     if  $a$ .data is compared with  $b$ .pubkey and  $a \notin T$  then
10:      if not  $a$  signed tx and not  $b$  signed tx then
11:         $V[i] \leftarrow V[i] \cup \{a\} \cup \{b\}$ 
12:      else
13:        for each  $c \in V[i]$  do
14:           $T \leftarrow T \cup \{c\}$ 
15:        end for
16:         $V[i] \leftarrow \emptyset$ 
17:      end if
18:    end if
19:  end if
20:  if program writes to lamports or data of account  $a$  then
21:    for each  $v \in V$  do
22:      if  $a \in v$  then
23:        signal VM
24:      end if
25:    end for
26:  end if
27: end while

```

of lamports in account *vault*'s lamports field. Since the account is marked as *vulnerable* and is not previously compared to at least one public key of an account which signed the transaction, the oracle signals the VM. However, after enabling the signer check in [Line 20](#) of [Listing 1](#), FUZZDELSOL transitively marks all three accounts as non-vulnerable.

Missing Owner Check Oracle. The *missing owner check oracle* checks whether an eBPF CMP instruction compares data of an account that is *not owned* by the program with the public key of another account. In case the program performs such a comparison, the oracle marks the latter account as potentially vulnerable to a missing owner check, because the former account could be an attacker-controlled account in which the attacker could assemble data in such a way to reference arbitrary accounts.

Consider our example from [Listing 1](#). Note that the code includes [Line 17](#) to show a valid owner check, i.e., in the example we assume that the code *does not contain* the check at [Line 17](#). Due to the missing check, the oracle marks the *vault* and *authority* accounts as vulnerable, because in [Lines 13](#) and [14](#) their keys are compared to data read from the *wallet_info* account. As [Listing 1](#) does not check the ownership of the *wallet_info* account, an attacker can forge this account such that his own public key assumes authority over the funds stored in any given *vault* account. To facilitate this, FUZZDELSOL uses the *attacker-controlled accounts generator* (cf. [Section 5.1](#)) to create attacker-controlled accounts that have references to accounts managed by the program as well as to the attacker to

trigger operations on the accounts the program manages without verifying the owner of the account. In this example, the attacker constructs the *wallet_info* account to contain his public key in the *wallet.authority* field and the public key any given *vault* account managed by the program in the *wallet.vault* field.

Lastly, the oracle signals the VM as soon as an eBPF STORE instruction writes to the lamports or data field of any vulnerable account. Note that when writing to a lamports field, the oracle signals the VM only if the STORE instruction decreases the account's lamports. In our example, the oracle signals the VM at [Line 28](#), because the program debits the *vault* account, which is marked as vulnerable.

Lamports-based Oracle. The *lamports-based oracle* checks for user-related accounts that could lose lamports to an attacker-related account. We illustrate the process for the *lamports-based oracle* in [Algorithm 2](#). Before executing the program, the runtime initializes user-related and attacker-related accounts with PDAs (cf. [Section 5.1](#)) (see [Line 1](#) and [2](#)).

This oracle is executed only if the issuer and signer of the transaction is the attacker (see [Line 5](#)). Consider a program where the attacker passes user-related PDA seeds as a parameter: An attacker could transfer lamports from user-related accounts to attacker-related accounts by specifying arbitrary seeds, despite proper owner and signer checks.

The lamports-based oracle checks eBPF STORE instructions to the lamports fields of accounts (see [Line 7](#)). Thus, the oracle checks each time if the program subtracts lamports from a user-related account (see [Line 10](#) and [11](#)) and credits lamports to an attacker-related account (see [Line 12](#) and [13](#)). The oracle signals the VM if a program transfers lamports from a user-related account to an attacker-related account (see [Line 17](#) and [18](#)). The lamports-based oracle only reports an error if neither the missing owner check, nor the missing signer check reports an error, since both oracles already check the lamports field.

Arbitrary CPI Oracle. The *arbitrary cross program invocation oracle* checks whether an eBPF CALL instruction calls the Solana-specific syscalls for invoking other programs using CPI. In case the program executes such an instruction, the program checks whether a previously specified malicious key was passed to the syscall as the target program, which implies that an attacker is able to invoke his malicious program. If a previously specified malicious key was passed as a target program to the syscall, the oracle signals the VM.

Missing Key Check Oracle. The *missing key check oracle* requires specifying the following two parameters before starting fuzzing: *i*) the address of a function f for which a key check is expected *ii*) and an account a , where the program is expected to compare the public key of an account it passes to function f with the public key of account a before calling function f . In the case of Wormhole, f would represent the function *load_instruction_at* and a would represent the *KeccakSecp256k* sysvar account. At program runtime, the oracle checks whether an eBPF CALL instruction calls the function f and signals the VM if the program does not compare the public key of the account passed to the function f with the public key of account a before calling function f .

Integer Bugs Oracle. The *integer bugs oracle* checks whether an eBPF STORE instruction writes an overflowed or underflowed value

Algorithm 2 Lamports-based Oracle

```

1:  $U \leftarrow$  user related accounts
2:  $A \leftarrow$  attacker related accounts
3:  $\text{lampor}^{\text{lose}}_{\text{user}} \leftarrow \text{false}$ 
4:  $\text{lampor}^{\text{gain}}_{\text{attacker}} \leftarrow \text{false}$ 
5: while program running and attacker signed tx do
6:
7:   if program writes to lampor field of account  $a$  then
8:      $l_{\text{prev}} \leftarrow$  lampor of  $a$  before write
9:      $l_{\text{after}} \leftarrow$  lampor of  $a$  after write
10:    if  $a \in U$  and  $l_{\text{after}} < l_{\text{prev}}$  then
11:       $\text{lampor}^{\text{lose}}_{\text{user}} \leftarrow \text{true}$ 
12:    else if  $a \in A$  and  $l_{\text{after}} > l_{\text{prev}}$  then
13:       $\text{lampor}^{\text{gain}}_{\text{attacker}} \leftarrow \text{true}$ 
14:    end if
15:  end if
16:
17:  if  $\text{lampor}^{\text{lose}}_{\text{user}} = \text{lampor}^{\text{gain}}_{\text{attacker}} = \text{true}$  then
18:    signal VM
19:  end if
20: end while

```

to the lampor field of an account. Subsequently, the oracle signals the VM. To do so, this oracle leverages the taint tracking engine. Note that the oracle is not applicable to the accounts' data field, since the Borsh serializer²—which programs frequently use for serializing and deserializing account data—calculates overflow or underflow values when serializing data into the account's data field.

5.3.4 Program Semantic Extractors. FUZZDELSOL does not have access to a program's source code, and therefore RunDelSol implements a *PDA seed structure extractor* and an *account data structure extractor* to determine semantic program information.

PDA Seed Structure Extractor. The PDA seed structure extractor extracts the structure and origin of the seeds used by the program to derive PDAs. For this purpose, we use the taint tracking engine to determine the origin of the arguments of Solana-specific syscalls that derive PDAs. This allows the PDA seed structure extractor to determine, for example, whether a seed originates from an account's public key. In case no origin for a seed is available, we assume that the seed has been statically compiled into the program. Hence, the PDA seed structure extractor does not record the origin of the seed (e.g., a public key) but the actual bytes of the seed.

Account Data Structure Extractor. The account data structure extractor extracts positions of public keys in the data structures of accounts. The extractor checks whether the program writes a public key to the data field of an account located on the *program input* (cf. Figure 1). In Listing 1, the *account data structure extractor* extracts the offsets of the *wallet.authority* and *wallet.vault* fields in the *wallet_info* account. Note that this information is extracted during account initialization of the *wallet_info* account. Further, FUZZDELSOL leverages this information in subsequent fuzzing runs to populate the data fields of *wallet_info*-like accounts with existing public keys that may refer to either attacker-controlled accounts or

²<https://borsh.io/>

program-managed accounts. Regarding the program in Listing 1, this means that FUZZDELSOL creates a *wallet_info* account with the exact keys required to trigger a bug that leads to illegitimate gain of lampor. Thus, the information persists across the independent transactions created by the fuzzer. The blockchain emulator uses this information to deploy accounts on the blockchain which can trigger potential missing owner check vulnerabilities.

5.4 Transaction Evaluator

After RunDelSol finishes the execution of the transaction, RunDelSol forwards the coverage information to the transaction evaluator. This includes the PDA seed structures, account data structures, and—if a bug was detected—the oracle signals, received at runtime. The coverage information and oracle signals determine the input bytes of the next fuzzing iteration. If the oracle signaled the eBPF VM, the transaction evaluator also creates a vulnerability report, that includes information to reproduce the transaction. In case RunDelSol extracted new PDA seed structures or account data structures, the transaction evaluator informs the blockchain emulator to generate a new ledger snapshot before starting the next fuzzing iteration.

6 EVALUATION

In this section, we evaluate multiple aspects of FUZZDELSOL on several datasets of Solana programs. We start by demonstrating the soundness and completeness precision of FUZZDELSOL with the *Neodyme Breakpoint Workshop* dataset [35] in Section 6.1. This dataset is a collection of prevalent Solana program vulnerabilities and is used in previous work [7]. Furthermore, we compare FUZZDELSOL with VRust [7], which is currently the only other approach addressing Solana program security. Second, we test FUZZDELSOL's vulnerability discovery effectiveness on real-world programs directly taken from the Solana mainnet blockchain. We present our findings and discuss newly discovered bugs in Section 6.2. Lastly, in Section 6.3, we demonstrate FUZZDELSOL's performance with bug bounty programs [24]. We focus on FUZZDELSOL's test case throughput and achieved code coverage.

Experimental Setup. We ran our evaluation on an AMD EPYC 7302P CPU with 16 cores clocked at 3 GHz with 256 GB RAM. The experiments are executed in parallel, keeping all physical CPU cores fully occupied. Each fuzzing experiment uses a single core and uses the same initial seed for all fuzzing runs.

6.1 Bug Detection Capabilities

To validate our design, we test FUZZDELSOL with the *Neodyme Breakpoint Workshop* dataset [35]. The dataset contains common Solana vulnerabilities (cf. Section 3.1). This dataset is organized into 5 different levels, where each level consists of a Solana program with a specific vulnerability. For this experiment, we fuzz each program with a timeout of 10 minutes. Our results are depicted in Table 1: FUZZDELSOL is able to find all the bugs in this dataset within less than 5 seconds. FUZZDELSOL does *not report any false alarms*, and can precisely detect each vulnerability.

We do not include the *Level 3* program, because the program has an *account confusion vulnerability*. Detecting account confusions requires knowledge of the underlying data layout that represents the expected data structure in memory. This information requires

| Program | Vulnerabilities | | | | | Time to first Bug (s) |
|-----------------------------------|-----------------|-------|-------|-------|---------|-----------------------|
| | IB | MOC | MSC | ACPI | MKC | |
| Level 0 | | ✓ | | | | 4 |
| Level 1 | ✓ | | | ✓ | | 2 |
| Level 2 | ✓ | | | | | 2 |
| Level 4 | | | | | ✓ | 1 |
| Wormhole* | | | | | ✓ | 37 |
| False Alarms (FUZZDELSOL : VRust) | 0 : 62 | 0 : 0 | 0 : 0 | 0 : 3 | 0 : 107 | N/A |

Table 1: Results of our validity measurement. We mark true bugs with ✓.

access to the source code of the program. However, FUZZDELSOL’s goal is to detect bugs in Solana programs, without relying on source code, and thus we skip this program.

Comparison with VRust. In contrast to VRust [7], FUZZDELSOL reliably detects bugs in smart contracts, without source code. Thus, a full comparison of every metric (e.g., performance) with VRust is impossible. While VRust can detect the same bugs as FUZZDELSOL, it reports false alarms regarding the integer bug in the *Level 2* program. Meanwhile, FUZZDELSOL does not report a single false alarm in this dataset. In addition, VRust only indicates a missing key check in the *Level 0* program and *Level 1* program. FUZZDELSOL, on the other hand, can trace the vulnerability to a missing owner check in the *Level 0* program and a missing signer check in the *Level 1* program, resulting in FUZZDELSOL being more precise compared to VRust.

The Infamous Wormhole Bug. In February 2022, wrapped Ether (wETH) with a value of 323 million USD has been stolen from the Wormhole program, which implements a bridge between Ethereum and Solana. The underlying bug is a missing key check in a Solana program. FUZZDELSOL implements a bug detection oracle that detects this bug. Furthermore, FUZZDELSOL works on a binary-only level and does not need to understand a program’s semantics to detect vulnerabilities. However, by design, the original Wormhole program and the underlying bug requires this level of context information. The context is provided by off-chain guardians that check and verify each transaction. However, we challenged FUZZDELSOL to detect this bug *without* any context information. Therefore, we created an emulation of this program which shares the same vulnerability as the Wormhole bug. Here, FUZZDELSOL was able to detect the bug in less than 40 seconds.

6.2 Discovering New Bugs

To evaluate FUZZDELSOL’s effectiveness in discovering unknown vulnerabilities, we assembled a dataset of 6049 real-world programs deployed on the Solana mainnet on March 27, 2023. We take the following steps to ensure that we have the most current and complete collection of Solana programs: First, we query an RPC node of the Solana network for all programs that belong to the most recent loader program³. Second, we use the Solana toolchain to dump each of the programs into an ELF file. VRust [7] cannot analyze any of these programs, which emphasizes the gap that FUZZDELSOL fills. Given the large data set of contracts, we set the timeout to

³At the time of writing, this is `BPFLoaderUpgradeab1e111111111111111111111111`.

| Program ID | Vulnerabilities | | | |
|-------------|-----------------|-------------|------|---------|
| | MSC | Integer Bug | ACPI | Lamport |
| 3nJ2...5erP | ✗ | | | |
| 3od3...jnsW | | | ✓ | |
| 3Vtj...4q8v | | ✓ | | |
| 3w57...obPW | | | ✓ | |
| 4hPk...JP5N | | | ✓ | |
| 4M2f...jStx | ✓ | ✓ | | |
| 6Lan...szqi | ✓ | | | |
| 7FWE...9p7p | | | ✓ | |
| 9a5d...jZP9 | | ✗ | | ✓ |
| 9tSW...11yy | ✓ | ✓ | | |
| GQ6q...1u6K | | ✓ | | |
| 9WoL...849B | | | ✓ | |
| H5rp...nPSG | | ✓ | | |

Table 2: Results of our bug discovery experiment. We mark true bugs with ✓ and false alarms with ✗.

5 minutes. In total, FUZZDELSOL reports 92 potential security vulnerabilities in 52 out of the 6049 programs, including 30 *missing signer checks*, 12 *arbitrary APIs*, and 30 *integer bugs*. Moreover, 20 reports indicate potential vulnerabilities to lamports theft without possessing the vulnerabilities listed before.

Confirming vulnerabilities is challenging due to the absence of source code. Hence, we opted for the following approach. We first generate instructions based on the payload information contained in the vulnerability report generated by FUZZDELSOL. Next, we analyze the program logs as well as the disassembled eBPF bytecode executed at the runtime of the instructions. Afterward, we craft transactions and observe if the transactions create an erroneous state in the blockchain. Note that this is a tedious validation process, but a common issue when developing smart contract fuzzers [19, 37, 42].

At the time of writing, we can validate the existence of 14 exploitable bugs, and 2 non-exploitable bugs. Accordingly, FUZZDELSOL currently has a false alarm rate of 12.5%. Table 2 shows the 16 discovered bugs and public key abbreviations of the vulnerable programs. We list the full public keys in the extended version of this paper [49]. In the following, we present five interesting vulnerabilities detected by our approach. To minimize potential damage, we ensured that none of these programs are actively managing assets and that no token accounts exist to which the programs are assigned as authority.

Responsible Disclosure and Ethical Concerns. Since we conduct this experiment on all Solana programs present on the Solana blockchain, this includes many programs of unknown origin, i.e., the authors are anonymous. We tried our best to reach out to the program authors of this experiment. Due to the lack of contact information, we could not disclose our findings directly to the authors of vulnerable Solana programs. Hence, we decided to disclose *all* of our findings to the Solana foundation⁴ and offered collaboration to fix the vulnerabilities.

⁴<https://solana.org/>

Finding 1: Integer bugs in 3Vtj...4q8v. This program contains an *integer bug* that an attacker can use to steal lamports from a program-controlled account. The vulnerable instruction requires four accounts, and subtracts one SOL from the fourth account while adding it to the first account. However, when reducing and crediting, the program does not check whether an overflow or underflow of lamports has occurred. Given that the first account owns too much SOL and the fourth account owns too little, an attacker could exploit the integer bug to add lamports to the fourth account and subtract lamports from the first account.

Finding 2: Arbitrary CPI in 3w57...obPW. FUZZDELSOL found an *arbitrary CPI* vulnerability. We confirmed the bug by sending a malicious instruction. The instruction accepts five accounts, where the last account signs the transaction. The program then invokes the first account supplied without checking the account.

Finding 3: Arbitrary CPI in 9WoL...849B. FUZZDELSOL detected an *arbitrary CPI* vulnerability which grants additional privileges to the invoked program by signing a PDA in the CPI instruction. The program 9WoL...849B manages accounts whose public key follows a PDA seed structure consisting of one seed corresponding to a passable public key (e.g., a wallet account). Similarly to 3w57...obPW, we have also created an instruction for 9WoL...849B to confirm the bug: The instruction expects 13 accounts, where account 13 is the program an attacker can invoke arbitrarily, and accounts 3, 8, and 11 are accounts that have the same public key matching the PDA seed structure of 9WoL...849B, i.e., a PDA associated to 9WoL...849B. The instruction results in 9WoL...849B invoking the arbitrary invocable program and signing the PDA in the CPI instruction. This leads to the invoked program having additional privileges than originally included in the transaction. Since the program states in the program log that it generates an instruction for calling the function *borrow_obligation_liquidity* of the Serum Swap program before performing CPI, we assume that it is supposed to manage accounts of the Serum Swap program as an authority.

Finding 4: Multiple Vulnerabilities in 4M2f...jStx. Here, FUZZDELSOL discovered both a *missing signer check* and an *integer bug* in this program. The program expects an instruction containing four accounts and allows playing a gambling game in which a user can win or lose. The number of lamports of the second account multiplied by an odd of 1.9934 determines the total payout of the game. In case the user wins the game, the program credits the payout to the first account and subtracts 1) from the fourth account the lamports worth 0.9334 multiplied by the lamports of the second account, and 2) from the second account, its total lamport balance. When crediting and subtracting the lamports, the program does not check whether overflows or underflows have occurred. Hence, an attacker can exploit the integer bug to credit the fourth account with lamports instead of subtracting lamports in the case that the user won the game. In addition, the program does not check which account signed the transaction. Thus, an attacker can submit arbitrary program-controlled accounts and start gambling without the program ever checking whether the attacker is authorized to gamble with the submitted accounts. We note that the programs 6Lan...szqi and 9tSW...11yy also allow gambling similar to 4M2f...jStx, and also do not verify that gambling with the submitted accounts is authorized.

Finding 5: Integer Bug in 9a5d...jZP9. Besides suffering from an integer bug, this program also enables an attacker to transfer lamports from a program-controlled account to an arbitrary, attacker-controlled account. The program expects two accounts and subtracts all lamports from the first account and credits them to the second account without checking for overflow or underflow of lamports. Thus, the integer bug is not exploitable, as only the lamports field of the second account can overflow but not of the first account. However, the program allows passing an arbitrary program-controlled account as a first account, transferring its lamports to the second account. In general, such a behavior is undesirable, as it allows an attacker to drain funds of all accounts belonging to this program. This is a clear indication of an access control bug: the intended program behavior would surely only allow an *authorized* account to transfer lamports from program-controlled accounts to carefully selected accounts. To address this bug, the program must verify that the authorized account is included in the instruction and that it signed the transaction to prevent exploitation.

6.3 Performance Analysis

Given that FUZZDELSOL is the first fuzzer for Solana programs, there exists no qualitative baseline or dataset to measure common fuzzing metrics like coverage and execution speed. Thus, we also aim at establishing a baseline allowing the community to compare future fuzzers with FUZZDELSOL.

We assembled a dataset from the Immunefi bug bounty list [24]. This dataset includes a diverse set of Solana programs used in production and offers a higher grade of code quality and complexity, compared to the average mainnet programs. Hence, we assess the execution speed and code coverage of FUZZDELSOL based on this dataset. For this experiment, we use a representative timeout of 24 hours and collect metrics on execution speed and code coverage.

Binary-only Approach Baseline. Unlike the experiments in Section 6.2, source code for the bug-bounty dataset is available. The performance and coverage of FUZZDELSOL could potentially be optimized by analyzing the source code to extract context information about authority or configuration accounts, solving assertions, and uncovering new code paths. However, we refrain from doing so because (1) we aim to provide representative measurements and (2) for the large majority of Solana programs no source code is available.

Challenge: Measuring Code Complexity of Solana programs. There is currently no tool support to measure the complexity of Solana smart contracts. Previous work, like VRust [7], relied on lines of code (LOC) to estimate the complexity of a program. However, this metric is insufficient since it includes unreachable code. Furthermore, FUZZDELSOL's coverage works on traversed edges in a program's control flow graph (CFG) and is therefore incomparable to LOC. To tackle this challenge, we developed a static analysis approach which measures the complexity based on traversing the eBPF code and counting every control flow instruction that eBPF supports. We use this number to over-approximate the number of edges in the CFG.

Note that this ensures that we include every eBPF JMP, CALL, and RET instruction. This includes any edges that are not reachable, since these may represent dead code or Solana-specific error

| Program Name | Bounty (\$) | #CFG Edges | Covered Edges | Mean Tx/s |
|--------------------|-------------|-------------------|------------------|------------------|
| Drift Protocol | 500 000 | 67 552 | 2336 | 5211 |
| Jet Airspace | 100 000 | 7559 | 1398 | 1951 |
| Jet Control | 100 000 | 9506 | 2716 | 1500 |
| ✗ Jet Fixed Term | 100 000 | 27 246 | 2375 | 1256 |
| Jet Margin | 100 000 | 18 216 | 1332 | 1771 |
| Jet Margin Swap | 100 000 | 12 843 | 2626 | 1332 |
| Jet Metadata | 100 000 | 5164 | 1310 | 786 |
| ✓ Jet Test Service | 100 000 | 17 972 | 2706 | 2894 |
| Lido | 2 000 000 | 8731 | 1305 | 274 |
| Marinade Finance | 250 000 | 22 424 | 2648 | 1023 |
| Port Finance VRL | 500 000 | 10 704 | 1643 | 1066 |
| Pyth | 500 000 | 4438 | 2984 | 576 |
| Solend Program | 1 000 000 | 10 818 | 1681 | 1129 |
| Sundial | 500 000 | 16 792 | 3171 | 1896 |
| Token Faucet | 500 000 | 6500 | 1299 | 1217 |
| Whirlpool | 500 000 | 20 593 | 2726 | 1229 |
| $n = 16$ | | $\bar{x} = 16691$ | $\bar{x} = 2141$ | $\bar{x} = 1569$ |

Table 3: Results of our performance experiment.

handling routines. For example, one routine is the handling of an incorrect serialized program id at the program input, which is never executed due to the valid instructions generated by the transaction generator.

We argue that this is sufficient to estimate the complexity of Solana programs because it provides a better insight into the complexity of a program than LOC. We analyze the target contracts with our CFG-based approach, and compare the results with the covered code paths by FUZZDELSOL.

Coverage. Table 3 shows the results of this experiment. First, we observe that the estimated complexity of the dataset varies widely, ranging from 4438 edges in the control flow graph to 67 552. This confirms that the size and complexity of the programs in the dataset is diverse. The number of covered edges by FUZZDELSOL ranges from 1299 to 3171. This provides two important insights: First, FUZZDELSOL can consistently generate meaningful transactions to uncover new program paths. Second, the binary-only analysis approach leads to several programs, where the number of covered edges does not increase over time. By further investigation, we learned that certain barriers or assertions prevent FUZZDELSOL from reaching deeper nested code. Hence, there is room for optimization for future work. For example, extending FUZZDELSOL with symbolic execution [23, 33], using Redqueen [3, 11], or, as mentioned before, incorporating the available source code (Section 6.2), to overcome these roadblocks.

On Fuzzing Throughput. Another interesting insight is that FUZZDELSOL is capable to generate on average more than 1569 transactions per second for every Solana program. Furthermore, FUZZDELSOL has an average of over 1000 transactions per second for 13 out of 16 programs. However, even in these three outliers, FUZZDELSOL can generate at minimum 274 transactions per second. We understand that FUZZDELSOL regularly extracts new runtime semantics for these programs, causing the blockchain emulator to initiate

updating the ledger snapshot as well as deriving new PDAs. We measured that initializing the snapshot takes about 60ms, which results in fewer transactions per second being generated from the input bytes.

Vulnerability Reports. FUZZDELSOL reported 2 bugs in the bug dataset both belonging to the Jet Protocol. We investigate the bugs while also consulting the source code of each program. As we will see, one of the bugs is a false positive while we believe the other one is a true positive which is currently under review by the developers.

Finding 6: False Alarm in Jet Fixed Term. FUZZDELSOL reports a missing signer check in the Jet Fixed Term program, which is a program for fixed-term lending and borrowing. The missing signer check exists in a function that cancels orders, which requires two accounts as its input. While the first account strictly belongs to a user of this program, the second account is a public order book containing the orders. By signing the transaction, a user is granted authority to remove an order from the marketplace. As a reminder, the oracle of FUZZDELSOL (c.f. Section 5.3.3) considers transitive signer checks if accounts are linked in some way, i.e., an account a may refer to another account b if a 's data field contains the public key of b . Additionally, the public key stored in a must be compared to the public key of b . However, in this particular case, the order book's account data may contain the public key of the authority account, but the public key of the authority account—which signed the transaction—is never compared to it. As a result, FUZZDELSOL reports a missing signer check for this program. However, we consider this a false alarm, as the bug is not exploitable because the program checks whether the user account owns the canceled order.

Finding 7: True Positive Bug in Jet Test Service. FUZZDELSOL reports an arbitrary cross-program invocation for the Jet Test Service program from the Jet Protocol. The arbitrary cross-program invocation exists in a function that accepts an arbitrary number of accounts as long as a minimum of two accounts is provided: the first account is a potentially uninitialized account, and the second account can be any other program of the Solana blockchain.

The function then checks whether the first account provided is initialized on the Solana blockchain, and if this is *not* the case, the function invokes the second program using CPI. Since there are no restrictions on the choice of account to invoke, we consider this a true bug in the Jet Test Service program. This bug can be overcome by having the Jet Test Service program verify the public key of the second account before invoking the second account using CPI.

We are now in contact with the vendor to resolve these issues and confirm the bugs. In conclusion, we can confirm the ability to fuzz complex targets with high transaction throughput and coverage.

7 RELATED WORK

In this section, we survey additional recent research work in the area of detecting bugs in smart contracts and fuzzing.

Solana Bug Detection Approaches. To the best of our knowledge, VRust [7] is currently the only static analysis tool for Solana programs. VRust finds missing owner, signer, and key checks, integer, account confusion, cross program invocation, numerical precision error, and bump seed bugs. It detects these bugs by analyzing source code and checking for vulnerable patterns. However, VRust relies

on the availability of source code which is unfortunately not available for the majority of Solana programs. Moreover, in contrast to FUZZDELSOL, VRust suffers from a high false alarm rate of 89.58%. Hence, due to the overwhelming number of alarms, it is very likely that developers will miss the true positives.

Library Fuzzing. Fuzzing [11, 12, 28, 48] is a popular technique for evaluating the security of software and hardware components [10, 31] and finding critical vulnerabilities. Different approaches are needed to fuzz complex targets that require well-structured data types [20, 22, 56]. For example, Fuzzil [20] is a grammar-based fuzzer for JavaScript engines. Recent research has focused on fuzzing targets on different platforms [6, 40, 46]. Fuzzware [46], SGXFuzz [6], and USBFuzz [40] fuzz their targets in an emulated environment. FUZZDELSOL follows a similar fuzzing strategy as it generates valid Solana transactions from randomly mutated input bytes and executes them in an emulated Solana blockchain.

Taint propagation [3, 21, 29, 41, 56] is widely used by fuzzers to identify which part of an input should be changed. Dowser [21] applies taints by identifying the input data bytes used in security-relevant operations. Vuzzer [41], Redqueen [3], and PATA [29] concentrate on guiding the fuzzer to pass barriers, e.g., passing a magic value validation. A similar technique could be incorporated in FUZZDELSOL to further improve the performance. In this paper, we use Libafl [12] as a fuzzer. Libafl is a high-performance fuzzer with state-of-the-art fuzzing techniques, e.g., including a *persistent mode* which avoids the bottleneck of forking a new process for each fuzzing iteration. Furthermore, the extensible design of Libafl enables us to integrate FUZZDELSOL's coverage information component into the fuzzing loop. Therefore, Libafl is the best available fuzzer for this work.

Smart Contract Fuzzing. Fuzzing has also been applied to Ethereum smart contracts [8, 25, 42, 51, 55]. Harvey [55] is a coverage-guided fuzzer for smart contracts. It uses program instrumentation to create a feedback mechanism for input prediction. In addition to that, it creates transaction sequences to detect smart contract bugs, like reentrancy and integer bugs. ConFuzzius [51] is a hybrid fuzzer. It leverages data dependency analysis and symbolic taint analysis to solve input constraints with the goal of reaching deeper nested paths, thereby increasing coverage. EFCF [42] is a coverage-guided binary-only fuzzer that tightly integrates with well-known fuzzing frameworks for native and legacy programs. EFCF transpiles the EVM bytecode of a smart contract to native C++ programs and uses state-of-the-art fuzzing optimization techniques [11].

However, all these smart contract fuzzing approaches only cover Ethereum smart contracts. Ethereum and Solana vulnerabilities are fundamentally different and require different bug detection mechanisms. In addition, the aforementioned approaches benefit from Ethereum's more advanced security tool landscape, which Solana lacks due to its immaturity, and transferring these techniques require extensive effort. As a result, these techniques are neither applicable to detect Solana-specific vulnerabilities, nor do they support Solana programs at all.

Smart Contract Analysis. In addition to fuzzing, other approaches to detect bugs and securing smart have been proposed [4, 33, 43, 44, 47, 50, 52]. Securify [52] deploys formal verification to detect vulnerabilities in Ethereum smart contracts. It discovers a variety

of vulnerabilities, including reentrancy, race conditions, and timestamp dependency. Oyente [4] and Manticore [33] are symbolic execution tools which analyze contracts to discover reentrancy and integer overflow issues. Sereum [44] deploys dynamic analysis to detect a variety of reentrancy attack patterns. EVMPatch [43] instruments the bytecode of smart contracts to enable instant patching of smart contracts. Osiris [50] introduces a symbolic execution approach for detecting integer-related bugs in Ethereum smart contracts. While these approaches are successful in detecting bugs in Ethereum smart contracts, they are heavily reliant on Solidity or EVM bytecode, which are closely tied to the Ethereum ecosystem. Furthermore, many works focus on bug classes that do not exist in Solana, like reentrancy. To conclude, these approaches cannot analyze Solana programs, and also lack detection capabilities for Solana-specific bugs.

8 CONCLUSION

In this paper, we propose FUZZDELSOL, the first coverage-guided fuzzer for Solana programs. FUZZDELSOL implements five oracles to detect common vulnerabilities in Solana programs, namely missing signer checks, missing owner checks, missing key checks, arbitrary APIs, and integer bugs. Moreover, FUZZDELSOL also implements a generic oracle based on lampport gains. To evaluate the capability of FUZZDELSOL to detect these vulnerabilities, we evaluated FUZZDELSOL against known vulnerabilities [35], including the infamous Wormhole bug [18], and compared its performance with VRust [7]. The evaluation has shown that FUZZDELSOL is able to detect the above-mentioned Solana program vulnerabilities and outperforms the state-of-the-art Solana bug detection tool. We also performed an extensive evaluation in which we fuzzed 6049 programs from the Solana mainnet blockchain. We showed that FUZZDELSOL is the only existing approach to effectively and precisely detect bugs in Solana smart contracts, without relying on source code. FUZZDELSOL reports 92 bugs in 52 programs, 14 of which we have verified as exploitable bugs at the time of writing the paper. Moreover, we evaluated FUZZDELSOL's performance against a set of bug bounty programs consisting of complex and productively used Solana programs. The evaluation has shown that FUZZDELSOL is capable to fuzz complex programs with up to 5211 transactions per second. Our large-scale evaluation is the most extensive security analysis of the Solana blockchain known to date.

ACKNOWLEDGMENT

Part of this research was conducted within a student project group at the University of Duisburg-Essen. We thank the project participants, Michael Mboni and Yelle Lieder, for their contribution. This work has been partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—SFB 1119 (CROSSING) 236615297 within project T1, EXC 2092 (CASA) 39078197—and the European Union (ERC, CONSEC, No. 101042266, and Horizon 2020 R&I, DYNABIC, No. 101070455).

REFERENCES

- [1] Vector 35. *Binary Ninja*. 2016. URL: <https://binary.ninja/>.
- [2] National Security Agency. *Ghidra*. 2019. URL: <https://ghidra-sre.org/>.
- [3] Aschermann et al. "REDQUEEN: Fuzzing with Input-to-State Correspondence". In: *NDSS Symp.* (2019).

- [4] Syed Badruddoja et al. "Making Smart Contracts Smarter". In: *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. May 2021, pp. 1–3.
- [5] Budweiser. *Budverse NFT*. <https://nft.budweiser.com/>. Accessed: 2023-4-22. 2023.
- [6] Tobias Cloosters et al. "SGXFuzz: Efficiently Synthesizing Nested Structures for SGX Enclave Fuzzing". In: *USENIX Security*. Aug. 2022.
- [7] Siwei Cui et al. "VRust: Automated Vulnerability Detection for Solana Smart Contracts". In: *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. 2022.
- [8] Mengjie Ding et al. "HFContractFuzzer: Fuzzing Hyperledger Fabric Smart Contracts for Vulnerability Detection". In: *Evaluation and Assessment in Software Engineering*. EASE 2021. Trondheim, Norway: Association for Computing Machinery, June 2021, pp. 321–328.
- [9] William Enck et al. "A study of android application security." In: *USENIX security symposium*. Vol. 2. 2. 2011.
- [10] Bo Feng, Alejandro Mera, and Long Lu. *P 2 IM: Scalable and hardware-independent firmware testing via automatic peripheral interface modeling*. <https://www.usenix.org/system/files/sec20-feng.pdf>. 2020.
- [11] Andrea Fioraldi et al. "AFL++: Combining Incremental Steps of Fuzzing Research". In: *14th USENIX Workshop on Offensive Technologies (WOOT 20)*. 2020.
- [12] Andrea Fioraldi et al. "LibAFL: A Framework to Build Modular and Reusable Fuzzers". In: *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. CCS '22. 2022.
- [13] Ethereum Foundation. *Ethereum*. <https://ethereum.org/>. Accessed: 2023-4-12. 2023.
- [14] Solana Foundation. *Solana*. <https://solana.com/>. Accessed: 2023-4-12. 2023.
- [15] Solana Foundation. *Solana Documentaion*. en. <https://docs.solana.com/>. Accessed: 2023-4-16. 2023.
- [16] Solana Foundation. *solana_rbpf*. en. <https://github.com/solana-labs/rbpf>. Accessed: 2023-4-16. 2023.
- [17] Joel Frank, Cornelius Aschermann, and Thorsten Holz. "ETHBMC: A Bounded Model Checker for Smart Contracts". In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 2757–2774. ISBN: 978-1-939133-17-5. URL: <https://www.usenix.org/conference/usenixsecurity20/presentation/frank>.
- [18] Dan Goodin. *How \$323M in crypto was stolen from a blockchain bridge called Wormhole*. en. <https://arstechnica.com/information-technology/2022/02/how-323-million-in-crypto-was-stolen-from-a-blockchain-bridge-called-wormhole/>. Accessed: 2023-4-12. Feb. 2022.
- [19] Gustavo Grieco et al. "Echidna: effective, usable, and fast fuzzing for smart contracts". In: *Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*. ISSTA 2020. Virtual Event, USA: Association for Computing Machinery, July 2020, pp. 557–560.
- [20] Samuel Groß. "Fuzzil: Coverage guided fuzzing for javascript engines". In: *Department of Informatics, Karlsruhe Institute of Technology* (2018).
- [21] Istvan Haller et al. "Dowsing for overflows: a guided fuzzer to find buffer boundary violations". In: *USENIX Security Symposium*. 2013, pp. 49–64.
- [22] Hyungseok Han, Donghyeon Oh, and Sang Kil Cha. "CodeAlchemist: Semantics-aware code generation to find vulnerabilities in JavaScript engines". In: *Proceedings 2019 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2019.
- [23] Jingxuan He et al. "Learning to Fuzz from Symbolic Execution with Application to Smart Contracts". In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. CCS. ACM, 2019. doi: [10.1145/3319535.3363230](https://doi.org/10.1145/3319535.3363230).
- [24] IMMUNI SOFTWARE PTE. LTD. *ImmuneFi Bug Bounties*. <https://immunefi.com/>. Accessed: 2023-3-30. 2020.
- [25] Bo Jiang, Ye Liu, and W. K. Chan. "ContractFuzzer: fuzzing smart contracts for vulnerability detection". In: *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. Ed. by Marianne Huchard, Christian Kästner, and Gordon Fraser. ASE. ACM, 2018. doi: [10.1145/3238147.3238177](https://doi.org/10.1145/3238147.3238177).
- [26] Lacoste. *Lacoste NFT*. <https://www.lacoste.com/en/undw3.html>. Accessed: 2023-4-22. 2023.
- [27] Martin Lee. *Solana: Scalability through speed*. en. <https://www.nansen.ai/research/solana-scalability-through-speed>. Accessed: 2023-4-12. Apr. 2022.
- [28] Hongliang Liang et al. "Fuzzing: State of the Art". In: *IEEE Trans. Reliab.* 67.3 (Sept. 2018), pp. 1199–1218.
- [29] Jie Liang et al. "PATA: Fuzzing with Path Aware Taint Analysis". In: *2022 IEEE Symposium on Security and Privacy (SP)*. May 2022, pp. 1–17.
- [30] OtterSec LLC. *BN-eBPF-Solana*. Accessed: 2023-8-9. 2022. URL: <https://github.com/otter-sec/bn-ebpf-solana>.
- [31] Dominik Maier, Lukas Seidel, and Shinjo Park. "BaseSAFE: baseband sanitized fuzzing through emulation". In: *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. WiSec '20. Linz, Austria: Association for Computing Machinery, July 2020, pp. 122–132.
- [32] Niko Matsakis. *Rust RFC 1211: MIR*. en. <https://rust-lang.github.io/rfcs/1211-mir.html>. Accessed: 2023-8-9. 2015.
- [33] Mark Mossberg et al. "Manticore: A User-Friendly Symbolic Execution Framework for Binaries and Smart Contracts". In: *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. Nov. 2019, pp. 1186–1189.
- [34] NBA. *NBA NFT*. <https://nbatopshot.com/>. Accessed: 2023-4-20. 2023.
- [35] Neodyme. *Introduction - Solana Security Workshop*. en. <https://workshop.neodyme.io/>. Accessed: 2023-4-16. 2021.
- [36] Neodyme. *Solana security.txt*. en. <https://github.com/neodyme-labs/solana-security-txt>. Accessed: 2023-8-9. 2022.
- [37] Tai D Nguyen et al. "sFuzz: an efficient adaptive fuzzer for solidity smart contracts". In: *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. ICSE '20. Seoul, South Korea: Association for Computing Machinery, Oct. 2020, pp. 778–788.
- [38] Nike. *Nike NFT*. <https://www.swoosh.nike/>. Accessed: 2023-4-22. 2023.
- [39] Richard Patel. *ghidra-eBPF*. Accessed: 2023-8-9. 2022. URL: <https://github.com/terorie/ghidra-ebpf>.
- [40] Hui Peng and Mathias Payer. *USBfuzz: A framework for fuzzing USB drivers by device emulation*. https://www.usenix.org/system/files/sec20-peng_0.pdf. Accessed: 2023-2-7. 2020.
- [41] Sanjay Rawat et al. "VUZZer: Application-aware evolutionary fuzzing". In: *Proceedings 2017 Network and Distributed System Security Symposium*. Vol. 17. San Diego, CA: Internet Society, 2017, pp. 1–14.
- [42] Michael Rodler et al. *EF/CF: High Performance Smart Contract Fuzzing for Exploit Generation*. 2023.
- [43] Michael Rodler et al. "EVMPatch: Timely and Automated Patching of Ethereum Smart Contracts". In: *USENIX Security Symposium*. 2021, pp. 1289–1306.
- [44] Michael Rodler et al. "Sereum: Protecting existing smart contracts against reentrancy attacks". In: *Proceedings 2019 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2019.
- [45] Sara Rouhani and Ralph Deters. "Performance analysis of ethereum transactions in private blockchain". In: *2017 8th IEEE International Conference on Software Engineering and Service Science (ICSESS)*. Nov. 2017, pp. 70–74.
- [46] Tobias Scharnowski et al. "Fuzzware: Using precise {MMIO} modeling for effective firmware fuzzing". In: *31st USENIX Security Symposium (USENIX Security 22)*. 2022, pp. 1239–1256.
- [47] Clara Schneidewind et al. "eThor: Practical and Provably Sound Static Analysis of Ethereum Smart Contracts". In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. CCS '20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 621–640.
- [48] Sergej Schumilo et al. "Nyx: Greybox hypervisor fuzzing using fast snapshots and affine types". In: *30th USENIX Security Symposium (USENIX Security 21)*. 2021, pp. 2597–2614.
- [49] Sven Smolka et al. *Fuzz on the Beach: Fuzzing Solana Smart Contracts*. 2023. arXiv: [2309.03006](https://arxiv.org/abs/2309.03006) [cs.CR].
- [50] 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*. ACSAC '18. San Juan, PR, USA: Association for Computing Machinery, Dec. 2018, pp. 664–676.
- [51] Christof Ferreira Torres et al. "Confuzzius: A Data Dependency-Aware Hybrid Fuzzer for Smart Contracts". In: *IEEE European Symposium on Security and Privacy*. EuroS&P. IEEE, 2021. doi: [10.1109/EuroSP51992.2021.00018](https://doi.org/10.1109/EuroSP51992.2021.00018).
- [52] Petar Tsankov et al. "Securityfy: Practical Security Analysis of Smart Contracts". In: *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. CCS '18. Toronto, Canada: Association for Computing Machinery, Oct. 2018, pp. 67–82.
- [53] Molly White. *Mango Markets exploiter arrested despite claiming all his actions were legal*. <https://web3isgoinggreat.com/?blockchain=solana&id=mango-markets-exploiter-arrested-despite-claiming-all-his-actions-were-legal>. Accessed: 2023-4-12. Dec. 2022.
- [54] Molly White. *Oracle attack on Solend costs the project \$1.26 million*. <https://web3isgoinggreat.com/?blockchain=solana&id=oracle-attack-on-solend-costs-the-project-1-26-million>. Accessed: 2023-4-12. Nov. 2022.
- [55] Valentin Wüstholtz and Maria Christakis. "Harvey: a greybox fuzzer for smart contracts". In: *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ESEC/FSE 2020. ACM, Nov. 2020. doi: [10.1145/3368089.3417064](https://doi.org/10.1145/3368089.3417064).
- [56] Wei You et al. "ProFuzzer: On-the-fly Input Type Probing for Better Zero-Day Vulnerability Discovery". In: *2019 IEEE Symposium on Security and Privacy (SP)*. ieeexplore.ieee.org, May 2019, pp. 769–786.
- [57] Michal Zalewski. *American Fuzzy Lop*. URL: <https://lcamtuf.coredump.cx/afll/>.