Automated Attack Synthesis for Constant Product Market Makers

Sujin Han *KAIST* Daejeon, Korea sujinhan@kaist.ac.kr Jinseo Kim KAIST Daejeon, Korea jinseo@kaist.ac.kr Sung-Ju Lee *KAIST* Daejeon, Korea profsj@kaist.ac.kr Insu Yun *KAIST* Daejeon, Korea insuyun@kaist.ac.kr

Abstract—Decentralized Finance enables many novel applications that were impossible in traditional finances. However, it also introduces new types of vulnerabilities, such as composability bugs. The composability bugs refer to issues that lead to erroneous behaviors when multiple smart contracts operate together. One typical example of composability bugs is those between token contracts and Constant Product Market Makers (CPMM), the most widely used model for Decentralized Exchanges. Since 2022, 23 exploits of such kind have resulted in a total loss of 2.2M USD. BlockSec, a smart contract auditing company, once reported that 138 exploits of such kind occurred just in February 2023.

We propose CPMM-Exploiter, which automatically detects and generates end-to-end exploits for CPMM composability bugs. Generating such end-to-end exploits is challenging due to the large search space of multiple contracts and various fees involved with financial services. To tackle this, we investigated real-world exploits regarding these vulnerabilities and identified that they arise due to violating two safety invariants. Based on this observation, we implemented CPMM-Exploiter, a new grammar-based fuzzer targeting the detection of these bugs. CPMM-Exploiter uses fuzzing to find transactions that break the invariants. It then refines these transactions to make them profitable for the attacker. We evaluated CPMM-Exploiter on two real-world exploit datasets. CPMM-Exploiter obtained recalls of 0.91 and 0.89, respectively, while five baselines achieved maximum recalls of 0.36 and 0.58, respectively. We further evaluated CPMM-Exploiter by running it on the latest blocks of the Ethereum and Binance networks. It successfully generated 18 new exploits, which can result in 12.9K USD profit in total.

Index Terms-smart contract, security, composability, fuzzing

I. INTRODUCTION

Decentralized Finance (DeFi) provides new financial services using blockchain and smart contracts. These services use tokens, which are digital assets beyond native currencies on the blockchain. A key financial service smart contracts offer is Decentralized Exchanges (DEX). Unlike Centralized Exchanges (CEX), DEXes enable users to swap one asset for another without a central authority. Through DEXes, blockchain users can freely convert their assets, which determines the pricing of assets and provides fluidity in the blockchain economy. To enable swapping without an intermediary, most DEXes adopt the Constant Product Market Maker (CPMM) model to automatically determine appropriate exchange rates. As of June 2023, approximately 77% of DEXes adopt the Uniswap protocol, which implements the CPMM model [1]

This usefulness of DeFi is often threatened by new types of vulnerabilities, one of which is composability bugs. Composability bugs refer to issues that lead to erroneous behaviors when multiple smart contracts operate together. This concept has been formulated by Babel et al. [2], who defined economic composability of smart contracts in DeFi. Economic composability is a property of a system where adding a new contract, C_{new} , to the existing system does not result in negative economic impacts. If economic composability is violated, vulnerabilities in a specific contract can potentially affect the entire system. This paper examines vulnerabilities that cause such violations in the context of token contracts and DEXes that follow the CPMM model (i.e., CPMM composability bugs). Recently, this type of vulnerability has been frequently exploited. For instance, on January 20, 2023, an attacker leveraged BRA token's flawed tax mechanism to steal around 225K USD worth of digital assets from a DEX trading BRA tokens [3]. Moreover, BlockSec, which is a renowned security auditing company, reported 138 attacks of such kind in just the month of February 2023.

Several tools have been proposed to detect multi-contract bugs similar to CPMM composability bugs, but they are ineffective at detecting CPMM composability bugs. For example, Echidna [4], one of the most popular fuzzing tools, supports transactions involving multiple contracts. ItyFuzz [5] utilizes snapshots to efficiently explore the large search space composed of multiple contracts. These tools have the advantage of being able to detect various forms of vulnerabilities, but they are not suitable for targeting CPMM composability bugs on a large scale. Due to the diversity of financial models in DeFi, distinguishing signs of CPMM composability bugs and intended behaviors is challenging. For example, a decrease in DEX token balance is a common sign of a CPMM composability bug. At the same time, numerous tokens adopt the deflationary model that constantly decreases the total supply of tokens. This behavior may be extended to automatically decrease the DEX token balance. However, this does not imply that all deflationary tokens are exploitable because profitability depends on many other factors, such as the attacker's token balance relative to DEX's and taxes charged to token transfers. Hence, to precisely detect CPMM composability bugs, tangible evidence (i.e., profit-generating transaction) is necessary, but fuzzers running without knowledge of token flow are unlikely to generate such transactions given a tight budget.

To address these limitations, we propose a two-step approach to detect CPMM composability bugs. Analyzing existing vulnerabilities, we discovered that CPMM composability bugs may occur when two safety invariants are broken. If these invariants are violated, attackers have an opportunity to steal assets from CPMM DEXes. Hence, instead of directly detecting vulnerabilities, our approach identifies violations of invariants and finds ways to abuse these violations for profit, ultimately leading to vulnerability detection.

Based on this insight, we design and implement CPMM-Exploiter to detect CPMM composability bugs and automatically synthesize exploits for the detected vulnerabilities. CPMM-Exploiter works as follows. First, CPMM-Exploiter utilizes grammar-based fuzzing to find a transaction that breaks the safety invariants. Then, CPMM-Exploiter refines the transaction to make it profitable. Through this process, CPMM-Exploiter can eliminate false positives by avoiding intended violations that do not lead to token leaks. According to our evaluation, CPMM-Exploiter outperformed five baselines in detecting CPMM composability bugs. On two realworld exploit datasets, CPMM-Exploiter obtained recalls of 0.91 and 0.89, while baselines achieved recalls of 0.36 and 0.58. Furthermore, to demonstrate the effectiveness of CPMM-Exploiter in a large-scale setting, we ran CPMM-Exploiter on the latest blocks of the Ethereum and Binance networks and discovered 18 transactions that are profitable, which can result in 12.9K USD total profit.

To summarize, we make the following contributions:

- We formalize composability bugs between DEXes following the CPMM model and token smart contracts. We identify two safety invariants that, when broken, allow an attacker to steal funds from DEXes.
- We propose CPMM-Exploiter, a two-step framework that detects CPMM composability bugs and generates exploits based on detected bugs.
- We evaluate *CPMM-Exploiter* on public datasets and compare it with five baseline tools: Echidna [4], Ity-fuzz [5], DeFiTainter [6], Slither [7], and Mythril [8]. In addition, we demonstrate the applicability of our tool in the wild by running it on Ethereum and Binance chains.

II. BACKGROUND

A. ERC20 Tokens

Tokens are digital assets beyond native currencies on the blockchain. Among these tokens, the most commonly used, fungible tokens are referred to as ERC20 tokens as they are originally defined through the Ethereum Request for Comments 20 (ERC20).¹ Native currencies (e.g., ETH in Ethereum or BNB in Binance) can also utilize ERC20 services through compatible tokens, such as Wrapped Ethereum (WETH) or Wrapped Binance Coin (WBNB).



Fig. 1. Example swap operation in a CPMM with x * y = 1000.

The ERC20 standard requires a token smart contract to implement a set of Application Binary Interface (ABI) consisting of 9 functions and 2 events. These functions are necessary for basic operations of tokens, such as transfer(address, value) and balanceOf(address). Such a uniform interface allows developers to build financial services, such as DEXes, for a countless number of tokens without having to write custom code for each token. This design increases the flexibility of ERC20 token implementation; however, it also poses the risk of potentially violating various invariants within the service, leading to security vulnerabilities.

B. Constant Product Market Maker Model

The Constant Product Market Maker (CPMM) model is adopted by DEXes to automatically swap one ERC20 token for another ERC20 token at an appropriate exchange rate. The CPMM model states that, given a DEX holding x amount of X tokens and y amount of Y tokens, the product of x and y should remain the same (i.e., $x \times y = k$). When a user requests to swap Δx amount of X tokens for Y tokens, the amount of Y tokens the DEX returns, Δy , is calculated with the equation $(x + \Delta x) \times (y - \Delta y) = k$. Thus, any swap operation in a CPMM DEX can be represented as a movement along the curve $x \times y = k$ as shown in Figure 1.

The majority of DEXes today charge a small percent fee for each exchange to provide profit for the liquidity providers who deposited the initial x amount of X tokens and y amount of Y tokens. For example, the Uniswap protocol, which is the most widely used DEX, charges a 0.3% fee for each exchange. As a result, most DEXes can be said to have adopted a modified version of the CPMM model, where the product of two assets slightly increases after each exchange (i.e., $x \times y \ge k$).

III. CPMM COMPOSABILITY BUGS

A. Terminology

Notation. Assume we have two ERC20 tokens, X token and Y token, and a DEX following the CPMM model for the two tokens, denoted by D. The quantity of X token and Y token in D are denoted with x and y, respectively.

Profitability. We define profitability in terms of gaining token X in one transaction. If an attacker begins with an initial balance of X tokens, a transaction that results in an increased balance of X tokens for the attacker without any additional financial input is a profitable transaction. Furthermore, we

¹Although each blockchain may refer to them differently according to their protocol (e.g., BEP20 or TRC20), we collectively call them as ERC20 Tokens in this paper.



Fig. 2. Example attack scenario where the attacker is able to decrease Y token balance of the DEX.

limit our scope to call sequences that can be executed in one transaction to exclude the impact of interest accumulation and other market players. In a typical attack scenario, X token would be a coin with relatively stable value, such as the native currency (e.g., WETH or WBNB) or stablecoins (e.g., USDT). CPMM composability bugs. We call a bug in the Y token contract that enables an attacker to construct a profitable transaction (i.e., call sequences that extract X tokens) from the system composed of X token, Y token, and D as a CPMM composability bug. For example, as demonstrated in Figure 1, if D has x = 20 and y = 50 ($x \times y = 1000$) and a user requests to swap 40 X tokens for Y tokens, the appropriate amount of Y tokens D should return according to the CPMM is 33. Disregarding other market players and fees, if the user immediately requests to swap the 33 Y tokens the user just obtained for X tokens, then D should return 40 X tokens. However, if the attacker is able to decrease y or increase one's own balance of Y tokens without cost, then the attacker is able to gain more than the expected 40 X tokens as depicted in Figure 2 and Figure 3. We can say that the additional X token is illegitimately obtained from the system because the attacker did not make any financial contributions to the system other than the initial 40 X tokens. Such vulnerabilities have been exploited to extract significant portions of DEX token balances and oftentimes drain the DEX entirely.

B. Type 1: DEX Token Balance Decrease

If an attacker can decrease y without making any payment (i.e., additional X tokens or liquidity tokens) to D, then the attacker can effectively alter the product, k, in the constant product function, $x \times y = k$ to a smaller value, effectively shifting the swap curve inward. An example scenario is shown in Figure 2. At point B, if the attacker can decrease the Y by around 8.3, the attacker can decrease k to 500 ($(16.7 - 8.3) \times$ $60 \approx 500$). Since the attacker only decreased y and X remains the same, the next swap will happen at a point straight below the previous point on the updated swap curve (i.e., point B'). The price of the Y token is greater at this point because the price of the Y token is determined by the ratio between x and y. Only decreasing y decreases the proportion of y, which increases the price of Y tokens in D and allows the attacker to swap the same amount of Y tokens for more X tokens (i.e., swapping to point A' instead of point A). Hence, when



Fig. 3. Example attack scenario where the attacker is able to increase one's own balance of Y tokens.

Invariant 1 is broken, attackers have an opportunity to extract X tokens owned by *D*.

Invariant 1 (DEX token balance decrease). Users should not be able to transfer or burn assets owned by a DEX without making any payment to the DEX.

Real-world example. For instance, on September 2, 2022, an attacker exploited a vulnerability in the ShadowFi token to steal around 1078 BNB, which was worth around 301K USD at the time of the exploit [9]. The vulnerability was that the ShaowFi token allowed any user to burn ShadowFi tokens in the Shadowfi DEX.

C. Type 2: Attacker Token Balance Increase

If an attacker can gain Y tokens without cost, then the attacker can also gain X tokens without cost through D. An example scenario is shown in Figure 3. At point B, if the attacker can increase its own balance of token Y without making any additional payments, then the attacker can gain more than the expected amount of X tokens (i.e., swapping to point A' instead of point A). Hence, when **Invariant 2** is broken, attackers have an opportunity to extract X tokens owned by D.

Invariant 2 (Attacker token balance increase). Users should not be able to obtain tokens traded in a DEX without cost.

Real-world example. For example, on January 10, 2023, an attacker leveraged BRA token's flawed tax mechanism to steal around 819 BNB, which was worth around 225K USD at the time of the exploit [3]. The BRA contract had a bug that transferred the same tax amount twice, allowing an attacker to accumulate a large sum of BRA tokens for free.

D. Prevalence of CPMM Composability Bugs

Recently, attackers have frequently exploited CPMM composability bugs to extract considerable amounts of tokens from DEXes. For example, BlockSec [10], a renowned smart contract auditing firm, reported that 138 exploits in February 2023 utilized **Invariant 1** violation. Furthermore, DeFiHack-Labs [11], a public exploit replication dataset, reported 23 realworld exploits that were consequences of either **Invariant 1** or **Invariant 2** breaking (Table I). The 23 exploits caused a cumulative loss of 2.2M USD.

 TABLE I

 CPMM composability bug found in DeFiHackLabs dataset.

| Vulnerable Token | Invariant Broken | Date of Exploit | Reported Loss | Reported Loss |
|---------------------|---------------------|--------------------|---------------|---------------|
| | Bronten | or Expron | | |
| Wdoge | 1 | 2022/04/24 | 78.6 BNB | 30.2K |
| LPC | 2 | 2022/07/25 | 45.1K B-USD | 45.1K |
| XST | 2 | 2022/08/10 | 27.4 ETH | 46.2K |
| Shadowfi | 1 | 2022/09/02 | 1.08K BNB | 300K |
| PLTD | 1 | 2022/10/18 | 24.5K B-USD | 24.5K |
| HEALTH | 1 | 2022/10/20 | 16.6 BNB | 4.54K |
| AES | 1 | 2022/12/07 | 61.6K B-USD | 61.6K |
| BGLD | 1 | 2022/12/12 | 8.80 BNB | 2.40K |
| BRA | 2 | 2023/01/10 | 228K B-USD | 228K |
| Upswing | 1 | 2023/01/18 | 22.6 ETH | 35.6K |
| ThoreumFi | 2 | 2023/01/19 | 2.26K BNB | 659K |
| SHEEP | 1 | 2023/02/10 | 9.54 BNB | 2.93K |
| Starlink | 1 | 2023/02/17 | 38.4 BNB | 11.8K |
| GPT | 1 | 2023/05/25 | 155K B-USD | 155K |
| ANCH | 2 | 2023/06/06 | 199K B-USD | 199K |
| Bamboo | 1 | 2023/07/04 | 235 BNB | 57.6K |
| ApeDAO | 1 | 2023/07/18 | 19.2K B-USD | 19.2K |
| BIGFI | 1 | 2023/09/07 | 30.3K B-USD | 30.3K |
| HCT | 1 | 2023/09/07 | 30.5 BNB | 6.58K |
| BFC | 1 | 2023/09/09 | 42.3K B-USD | 42.3K |
| pSeudoEth | 2 | 2023/10/08 | 1.44 ETH | 2.34K |
| TGBS | 1 | 2024/03/06 | 377 BNB | 154K |
| GHT | 1 | 2024/03/07 | 15.4 ETH | 58.6K |

Note. BSC-USD was denoted as B-USD.



Fig. 4. Overall workflow of CPMM-Exploiter.

IV. OVERVIEW

To automatically detect and exploit CPMM composability bugs, we propose *CPMM-Exploiter*. In this section, we describe our goals for *CPMM-Exploiter* design, the technical challenges in achieving those goals, and our approach to address the challenges.

A. Goals

To apply *CPMM-Exploiter* in the real world, we set the following goals:

Fully automated. *CPMM-Exploiter* should automatically generate exploits for CPMM composability bugs without human intervention. This is crucial as the number of smart contracts is increasing rapidly, and manual analysis is not scalable. Many existing tools detect mere signs of bugs without providing an

end-to-end exploit (e.g., suspicious token balance changes [5]). Unfortunately, this design requires developers to manually analyze contracts to confirm vulnerabilities, making them unscalable and time-consuming.

Efficient. As the number of smart contracts is increasing rapidly, *CPMM-Exploiter* should be able to analyze a large number of contracts in a short time. We set up a goal to analyze one contract in a few minutes. By doing so, *CPMM-Exploiter* can be used to analyze many contracts in all blockchains within days.

Generic. We aim to make *CPMM-Exploiter* generic so that it can be applied to various tokens and DEXes. This is important because each token implements its own business logic, and each DEX has its own CPMM implementation. *CPMM-Exploiter* should be able to analyze any token and DEX without any modification.

B. Technical Challenges

Large search space. Several technical challenges exist in achieving the above goals. Since CPMM composability bugs are triggered when multiple contracts interact, one must explore the vast search space composed of them. Moreover, the stateful nature of smart contracts exacerbates the search space problem because the same functions may behave differently based on contract states. Moreover, conditional statements may depend on states, making argument generation difficult. At the same time, to generate exploits within a short time frame, we need to identify and prioritize areas that are highly likely to contain CPMM composability bugs. However, due to the diversity in token contract implementation, designing a suitable guidance strategy is a challenging task.

Complex financial ecosystem. Furthermore, even after identifying specific issues in the contracts, crafting an end-to-end exploit poses additional challenges. This is because the exploit should be profitable, even considering the various fees and costs attached to financial services. The most representative example would be the exchange fees imposed by DEXes. Most DEXes charge a percent fee for every exchange. To construct an exploit, we need to make its revenue high enough to offset the fees and costs involved with executing the transaction, including DEX exchange fees.

C. Scope

CPMM-Exploiter is implemented for smart contracts running on EVM-based blockchain (e.g., Ethereum and Binance Smart Chain). Moreover, *CPMM-Exploiter* aims to steal ERC20 tokens from Uniswap V2 DEXes, which is the most popular implementation of the CPMM model. However, we believe *CPMM-Exploiter* can be applied to other CPMMs. We support smart contracts with or without source code. However, our analysis can be more accurate (e.g., using ABIs) if the source code is available.

V. DESIGN

In this section, we describe the design of CPMM-Exploiter.

| CPMM-Exploiter | Testcase | Grammar |
|----------------|----------|---------|
|----------------|----------|---------|

| $\langle Transaction \rangle ::= \langle SwapXY \rangle \langle Payload \rangle \langle SwapYX \rangle$ |
|--|
| $\langle Payload \rangle ::= \langle Cycle \rangle$ $ \langle StateChange \rangle$ |
| $\langle Cycle \rangle \langle StateChange \rangle$ |
| $\langle Cycle \rangle \qquad ::= \langle CycleA \rangle$ |
| $\langle CycleB \rangle$ $\langle CycleC \rangle$ |
| $\langle CycleA \rangle ::= \langle CycleA \rangle \langle CycleA \rangle$ |
| $ \langle UserToUser \rangle$ |
| $\langle CycleB \rangle ::= \langle CycleB \rangle \langle CycleB \rangle$ |
| {UserIODEX} {DEXIOUSER} |
| $\langle CycleC \rangle$::= $\langle UserToDEX \rangle \langle CycleD \rangle \langle DEXToUser \rangle$ |
| $\begin{array}{ll} \langle CycleD \rangle & ::= & \langle CycleD \rangle \langle CycleD \rangle \\ & & & \langle DEXToDEX \rangle \end{array}$ |
| $\langle \textit{UserToUser} \rangle$::= Y.transfer(this, amount) |
| $\langle \textit{UserToDEX} \rangle ::=$ Y.transfer(DEX, amount) |
| <pre>(DEXToUser) ::= DEX.skim(this)</pre> |
| <pre>(DEXToDEX) ::= DEX.skim(DEX)</pre> |
| $\langle StateChange \rangle ::= \langle StateChange \rangle \langle StateChange \rangle$ |
| DEX.sync() |
| Y.burn(amount) |
| |

Fig. 5. *CPMM-Exploiter's* grammar represented in Backus-Naur Form. Repeat rules (marked with bold) are not used for testcase generation in the first step. They are used for building exploits in the second step.

A. Workflow

The overall workflow of CPMM-Exploiter is illustrated in Figure 4. CPMM-Exploiter adopts a two-step approach to generate an exploit based on CPMM composability bugs. In the first step, CPMM-Exploiter utilizes grammar-based fuzzing to find a transaction that breaks invariants in section III. For this, CPMM-Exploiter utilizes contract ABIs if available. If any transaction results in profit (i.e., exploits), CPMM-Exploiter terminates early. On the other hand, if CPMM-Exploiter can only find invariant-breaking transactions, CPMM-Exploiter proceeds to the second step. In particular, it refines the invariant-breaking transactions to build an exploit. This step is necessary to distinguish exploitable invariant violations from intended ones accompanied by safety measures to protect DEXes' assets (e.g., deflationary tokens that decrease balances of all token holders). To exacerbate the broken invariants, CPMM-Exploiter repeats segments of the transaction to either further decrease DEX token balance (breaking **Invariant 1**) or to further increase the attacker's one (breaking Invariant 2). Finally, if a profitable transaction is generated, CPMM-Exploiter returns the transaction and flags the set of contracts as vulnerable.

B. Finding Invariant Violations

As a first step, *CPMM-Exploiter* searches for invariant violations between given contracts and DEXes. Unfortunately,

it is difficult to find such violations by exploring all possible functions due to the complexity of DeFi contracts. To address this, we employ grammar-based fuzzing to explore states relevant to CPMM composability bugs efficiently. In particular, we focus on interactions between a target token (Y) and a DEX, which are the primary causes of CPMM bugs. Figure 5 illustrates the grammar used by CPMM-Exploiter to generate testcases. At the beginning, CPMM-Exploiter swaps X tokens (e.g., stablecoins) to the target token Y (<SwapXY>). Then, CPMM-Exploiter attempts to break invariants by injecting a payload that consists of cyclic token transfers and statechanging functions (<Payload>). Finally, CPMM-Exploiter swaps Y tokens back to X tokens to complete the transaction (<SwapYX>). In the following, we discuss our intuition behind the grammar and how we use it to find invariant-breaking transactions.

Cyclic transfers. One of the key components of our payload is cyclic token transfers. The cyclic token transfers are necessary to simulate diverse interactions between the attacker and the DEX, while ensuring that the attacker does not lose any tokens. In CPMM-Exploiter, we consider three types of cyclic transfers: <CycleA>, <CycleB>, and <CycleC>. <CycleA> represents a self-transfer in the attacker's account, <CycleB> represents an exchange between the attacker and the DEX, and <CycleC> is similar to <CycleB> but involves a self-transfer in the DEX's account. It is worth to note that DEX.skim(this) returns tokens if the DEX has more tokens than needed. In a normal scenario, such operations are meaningless as they circulate tokens between the attacker and the DEX. However, if the target token Y implements transferrelated features (e.g., deflationary mechanisms), this can lead to invariant violations.

State-changing functions. In addition to cyclic transfers, we consider state-changing functions that can alter the state of the DEX or Y. There could be various state-changing functions in the DEX and Y contracts. One of the most critical functions is DEX.sync(), which updates the k value in the CPMM curve (i.e., $x \times y = k$). Another example is burn-related functions in the Y contract, which can decrease the total supply of Y tokens. We also consider other state-changing functions in the Y contract that do not take arguments. Such functions are often used for various purposes, such as updating internal state variables or providing bonuses to token owners.

Generation. Given Y contract ABI, *CPMM-Exploiter* derives all possible testcases by following the grammar in Figure 5. It is worth noting that *CPMM-Exploiter* does not perform repetition at this point to first survey the large space. For amount, *CPMM-Exploiter* randomly fills them with various values, including zero, the attacker's balance, and the DEX's balance. To start, the initial X token of the attacker is set to an amount that can swap out a random value between 1% and 99% of the DEX's Y token balance.

C. Refining to Build Exploits

Even if we can find invariant-breaking transactions, it does not always imply a CPMM composability bug. Since

TABLE IILOC TO IMPLEMENT CPMM-Exploiter.

| Components | Lines of Code |
|--------------------------------|--|
| Grammar-based fuzzing | 2,613 Loc of Rust |
| Solidity Execution Environment | 893 LoC of Rust 1,064 LoC of Solidity |

each token adopts a unique economic model, it may have intentionally broken invariants with mitigation measures to protect assets in DEX. To verify whether the detected violation is exploitable, *CPMM-Exploiter* attempts to synthesize an exploit (i.e., profitable transaction) with the invariant-breaking transaction found earlier.

To build an exploit, *CPMM-Exploiter* aims to increase the exploit revenue by repeating call sequences. Such repeatable segments are <CycleA>, <CycleB>, <CycleD>, and <StateChange> in Figure 5. For **Invariant 1**, we need to decrease the DEX token balance of Y to increase its price in DEX. For **Invariant 2**, we need to earn more Y tokens without cost to gain more X tokens. Repeating invariant-breaking calls can exacerbate the broken invariant, increasing exploit revenue and chances to build a profitable transaction.

D. State Tracking

CPMM-Exploiter executes the generated testcases on a simulated on-chain environment. We modified the EVM backend to keep track of important state variables, such as the current token balances of the attacker and the DEX. These values are used for two purposes. First, they replace arguments in the testcases as such variables can change over time (e.g., token balances). Second, they are also used to determine whether safety invariants from section III have been broken (i.e., whether DEX Y token balance decreased or attacker Y token balance increased).

VI. IMPLEMENTATION

CPMM-Exploiter was built on top of Foundry [12] and relies on Foundry to set up and fetch on-chain data necessary for simulated on-chain testing. The LoC to implement *CPMM-Exploiter* is shown in Table II.

Most token exchanges (i.e., <SwapXY> and <SwapYX>) are handled through the relevant Uniswap Router, instead of directly exchanging through the DEX, because direct exchanges through DEX require calculation of appropriate input and output amounts before the exchange. However, tokens that take exclusive fees (i.e., fees outside of the transfer amount) are incompatible with Uniswap Routers. Thus, when token exchange through Uniswap V2 Router fails, *CPMM-Exploiter* runs another simulated environment to calculate the fee percent and utilizes the fee information to directly swap tokens by calculating the correct input and output quantities.

VII. EVALUATION

To evaluate *CPMM-Exploiter*, we answer the following research questions:

- **RQ1:** How effective is *CPMM-Exploiter* at detecting CPMM composability bugs compared to existing tools?
- **RQ2:** How efficient is *CPMM-Exploiter* at detecting CPMM composability bugs compared to existing tools?
- **RQ3:** How significant are the techniques applied to *CPMM-Exploiter*?
- **RQ4:** How effective is *CPMM-Exploiter* at detecting undiscovered CPMM composability bugs in the real-world?

A. Experimental Setup

1) Baseline Selection: Among many existing tools for smart contract analysis, we selected five tools as baselines for the evaluation: ItyFuzz [5], Echidna [4], DeFiTainter [6], Slither [7] and Mythril [8]. We selected ItyFuzz, Echidna, and DeFiTainter as they support multi-contract analysis and can detect (a subset of) CPMM composability bugs. We also included Slither and Mythril, which do not support multicontract analysis, to demonstrate that tools designed for single contracts are ineffective at detecting CPMM composability bugs. We also attempted to include EF/CF [13] and Clockwork Finance [2] as baselines, but they were not selected because we could not run them for our datasets. EF/CF does not support on-chain fuzzing if contracts require large on-chain storage, which is true for many token contracts. Meanwhile, Clockwork Finance requires manual modeling for each contract, which is not feasible for our large-scale evaluation.

In the following, we describe the configurations for each tool used in the evaluation. We tried to configure each tool to provide the fairest possible comparison.

ItyFuzz. We run ItyFuzz with only the bug oracle that detects ERC20 token leaks. ItyFuzz also has a bug oracle for detecting token imbalances in DEXes (similar to **Invariant 1**), but these issues do not always lead to vulnerabilities. Here, we expect ItyFuzz to detect end-to-end exploits that result in profit, similar to how *CPMM-Exploiter* operates.

Echidna. Echidna requires custom oracles to detect vulnerabilities. Thus, we implemented an oracle that checks whether the attacker contract can get more native tokens after exchanging all ERC20 tokens for native ones. We also set up enough initial native currency balance for each exploit in Echnida (i.e., 10,000 ETH or 10,000 BNB).

DeFiTainter. DeFiTainter determines whether a given function contains a price manipulation vulnerability. Thus, we ran DeFiTainter for all public and external functions of a contract and flagged the contract as vulnerable if any of the functions outputted a positive result. Since DeFiTainter requires source code analysis, we could not run it for close-sourced contracts. **Mythril**. Mythril has no detector for ERC20 token or ether leaks. However, other detectors may have detected the programmatic error, leading to broken safety invariants for CP-MMs. Thus, we manually validate each result to check if Mythril can find the root cause of each exploit.

Slither. Since Slither includes many non-critical detectors, we ran Slither with only detectors that could be a potential

TABLE III

CPMM COMPOSABILITY BUG DATASETS USED FOR EVALUATION. DEFIHACKALABS CONTAINS PAST EXPLOITS THAT BREAK EITHER INVARIANT 1 OR 2. BLOCKSEC CONTAINS PAST EXPLOITS THAT BREAK INVARIANT 1. REALWORLD-ETH AND REALWORLD-BSC CONTAIN ON-CHAIN UNISWAP/PANCAKE SWAP CONTRACTS WITH ASSETS WORTH MORE THAN 1,000 USD.

| Dataset | Number of Contracts |
|---------------|---------------------|
| DeFiHackLabs | 23 |
| BlockSec | 124 |
| RealWorld-ETH | 23,701 |
| RealWorld-BSC | 20,607 |

root cause for CPMM composability bugs (i.e. arbitrary-senderc20, protected-vars, arbitrary-send-erc20-permit, arbitrarysend-eth, unchecked-transfer). Then, similar to Mythril, we manually validate its result to check if it can discover the root cause of each exploit. Since Slither requires source code analysis, we cannot run it for close-sourced contracts.

2) Datasets: Datasets used for evaluation are shown in Table III. We used two datasets for evaluation. First, we use DeFiHackLabs [11], a public dataset containing exploit replications for reported DeFi hacking incidents. This dataset has been widely used for evaluating smart contract analysis tools [14], [15]. Among these, we selected 23 exploits that utilize CPMM composability bugs. Second, for a more thorough evaluation, we also used BlockSec [16], a public dataset containing 138 real-world exploits that involve breaking Invariant 1. Out of the 138 exploits, we use 124 exploits for this evaluation, as 14 are duplicate exploits of the same DEX. For duplicates, the exploit with the earliest block number was kept. Only one exploit, the SHEEP token exploit, is included in both the DeFiHackLabs and BlockSec datasets. Third, we also evaluate CPMM-Exploiter on Uniswap and PancakeSwap contracts on the Ethereum and Binance networks that contain more than 1,000 USD worth of native tokens or stablecoins. These datasets were used to evaluate how effective CPMM-Exploiter is at detecting undiscovered CPMM composability bugs in the real world.

B. Effectiveness in Detecting CPMM Composability Bugs

To compare the effectiveness of *CPMM-Exploiter* in detecting CPMM composability bugs with existing tools, we measured the recall of *CPMM-Exploiter* and each of the baselines for the DeFiHackLabs and BlockSec datasets ² To avoid non-deterministic results from fuzzing, we ran fuzzing-based approaches (i.e., *CPMM-Exploiter*, ItyFuzz, and Echidna) three times for each exploit and reported the average recall. We use 20 minutes as the timeout for each contract, which is reasonably long enough if we consider the number of contracts to analyze in the real world (e.g., tens of thousands of DEX contracts).

TABLE IV CPMM COMPOSABILITY BUG DETECTION RATE OF CPMM-Exploiter and baselines on the DeFiHackLabs dataset.

| Token | Ours | ItyFuzz | Echidna | DeFiTainter | Slither | Mythril |
|-----------------|---------------|-----------------|--------------|--------------|--------------|--------------|
| AES | 1 | 1 | 0 | 0 | 0 | 0 |
| ANCH | 1 | 0 | 0 | 0 | 0 | 0 |
| ApeDAO | 0 | 0 | 0 | 1 | 0 | 0 |
| Bamboo | 1 | 1 | 0 | 0 | 0 | 0 |
| BFC | 1 | 0.33 | 0 | 0 | 0 | 0 |
| BGLD | 1 | 0 | 0 | 0 | 0 | 0 |
| BIGFI | 1 | 0.33 | 0 | 0 | 0 | 0 |
| BRA | 1 | 0 | 0 | 0 | 0 | 0 |
| GHT | 0 | 0 | 0 | - | - | 0 |
| GPT | 1 | 0 | 0 | - | - | 0 |
| HCT | 1 | 0.33 | 0 | 0 | 0 | 0 |
| HEALTH | 1 | 0 | 0 | 0 | 0 | 0 |
| LPC | 1 | 1 | 0 | 0 | 0 | 0 |
| PLTD | 1 | 0 | 0 | 0 | 0 | 0 |
| pSeudoEth | 1 | 1 | 0 | - | - | 0 |
| Shadowfi | 1 | 0 | 0 | 0 | 0 | 0 |
| SHEEP | 1 | 0.33 | 0 | 0 | 0 | 0 |
| Starlink | 1 | 0 | 0 | 0 | 0 | 0 |
| TGBS | 1 | 1 | 0 | 0 | 0 | 0 |
| ThoreumFi | 1 | 0 | 0 | - | - | 0 |
| Upswing | 1 | 1 | 0 | 0 | 0 | 0 |
| Wdoge | 1 | 1 | 0 | 0 | 0 | 0 |
| XST | 1 | 0 | 0 | 0 | 0 | 0 |
| Total Recall | 21/23 0.91 | 8.33/23 0.36 | 0/23 0.00 | 1/19 0.05 | 0/19 0.00 | 0/23 0.00 |

TABLE V CPMM composability bug detection rate of *CPMM-Exploiter* and baselines on the BlockSec dataset.

| | Ours | ItyFuzz | Echidna | DeFiTainter | Slither | Mythril |
|--------|---------|---------|---------|-------------|---------|---------|
| Total | 110/124 | 72/124 | 11/124 | 1/123 | 0/123 | 0/124 |
| Recall | 0.89 | 0.58 | 0.09 | 0.01 | 0.00 | 0.00 |

Table IV and Table V show the results of running *CPMM*-*Exploiter* and baselines on the DeFiHackLabs and BlockSec datasets, respectively. In summary, *CPMM*-*Exploiter* outperformed other tools in detecting CPMM composability bugs. In DeFiHackLabs dataset (Table IV), *CPMM*-*Exploiter* detected 21 out of 23 exploits, while ItyFuzz detected 8.33 exploits on average. Moreover, *CPMM*-*Exploiter* achieved the highest recall value of 0.91, while ItyFuzz had the second-highest recall value of 0.36. DeFiTainter detected only one vulnerability out of 19 contracts it could analyze, thus having a recall value of 0.05. Other tools failed to detect any vulnerabilities. In the BlockSec dataset (Table V), *CPMM*-*Exploiter* also achieved the highest recall of 0.89, while ItyFuzz had the second-highest recall of 0.58.

CPMM-Exploiter obtained significantly higher recalls than other tools because it efficiently explores various contract states focusing on token transfers. ItyFuzz and Echidna were suboptimal in generating exploits because they generated testcases without guidance on the flow of tokens. Furthermore, Slither and Mythril could not identify erroneous behaviors in token contracts because they do not know how those contracts operate with DEXes. The results indicate the need for a

²We did not measure the precision of *CPMM-Exploiter* as it generates endto-end exploits in on-chain environment, always yielding 100% precision.

TABLE VI Average time taken by *CPMM-Exploiter*, ItyFuzz, and Echidna to detect bugs in the DeFiHackLabs dataset (seconds).

| Token | Ours | ItyFuzz |
|-----------|------|---------|
| AES | 22 | 508 |
| ANCH | 11 | - |
| ApeDAO | - | - |
| Bamboo | 8 | 17 |
| BFC | 12 | 86 |
| BGLD | 133 | - |
| BIGFI | 10 | 927 |
| BRA | 12 | - |
| GHT | - | - |
| GPT | 12 | - |
| HCT | 9 | 177 |
| HEALTH | 11 | - |
| LPC | 12 | 102 |
| PLTD | 32 | - |
| pSeudoEth | 7 | 26 |
| Shadowfi | 106 | - |
| SHEEP | 7 | 795 |
| Starlink | 11 | - |
| TGBS | 15 | 15 |
| ThoreumFi | 24 | - |
| Upswing | 9 | 33 |
| Wdoge | 669 | 27 |
| XST | 11 | - |
| Average | 54 | 246 |



Fig. 6. Heatmap of time taken to detect CPMM composability bugs in BlockSec dataset (in seconds).

targeted approach to detect CPMM composability bugs.

| Answer | to RQ1: CPMM-Exploiter outperforms | |
|----------|---|-----|
| existing | tools in detecting CPMM composability buy | gs. |

C. Efficiency in Detecting CPMM Composability Bugs

To compare the efficiency of *CPMM-Exploiter* in detecting CPMM composability bugs with existing tools, we measured the time taken by *CPMM-Exploiter*, ItyFuzz, and Echidna to detect CPMM composability bugs in DeFiHackLabs and BlockSec datasets. We report the average time taken, excluding trials when bugs were undetected. Static analysis tools (i.e., DeFiTainter, Slither, and Mythril) are not included in this comparison, as they could not detect nearly all CPMM composability bugs.

Table VI shows the time taken to detect each vulnerability in the DeFiHackLabs dataset. The last row contains the average time taken to detect vulnerabilities overall. On average, *CPMM-Exploiter* took 54 seconds to detect vulnerabilities, around 4.56 times faster than the average time taken by ItyFuzz, which is 246 seconds. However, ItyFuzz was able to detect one vulnerability faster than *CPMM-Exploiter* for one exploit (i.e., Wdoge). This is because the Wdoge charges an exclusive fee for token transfers. To calculate and take account of such fees, *CPMM-Exploiter* deploys another simulated environment and tests transfers, which incurs nonnegligible latency overhead.

Figure 6 is a visual representation of how fast each tool was at finding each vulnerability in the BlockSec dataset. Each cell in the heatmap contains the result for one vulnerability, thus a total of 124 cells per tool. The red color indicates that the tool took a relatively long time (close to 1200 seconds or 20 minutes) to detect the vulnerability, while the blue color indicates that the tool took a relatively short time (close to 0 seconds) to detect the vulnerability. White cells indicate that the tool was not able to detect vulnerability for all three trials.

CPMM-Exploiter was able to detect most vulnerabilities in a short period, while ItyFuzz detected vulnerabilities in varying time frames. On average, ItyFuzz and Echidna took 410 seconds and 151 seconds to detect vulnerabilities, respectively. Meanwhile, *CPMM-Exploiter* took only 11 seconds to detect vulnerabilities, around 37 times faster than the average time taken by ItyFuzz and 13.7 times faster than the average time taken by Echidna.

Answer to RQ2: On average, *CPMM-Exploiter* detects CPMM composability bugs 4.56 to 37 times faster than existing tools.

D. Ablation Study

To demonstrate the effectiveness of the two-step approach, we conducted evaluations with two modified versions of *CPMM-Exploiter: CPMM-Exploiter-NoRepeat* and *CPMM-Exploiter-NoIC. CPMM-Exploiter-NoRepeat* only runs generated testcases and does not utilize repetitions to generate exploits. *CPMM-Exploiter-NoIC* generates testcases with a random number of repetitions and directly checks for profit generation. Similar to previous evaluations, *CPMM-Exploiter-NoRepeat* and *CPMM-Exploiter-NoIC* were each run three times.

On average, CPMM-Exploiter-NoRepeat detected 11 out of 23 vulnerabilities, and CPMM-Exploiter-NoIC detected 16 out of 23 vulnerabilities. Meanwhile, CPMM-Exploiter detected 21 out of 23 vulnerabilities. Such an outcome is expected. CPMM-Exploiter-NoRepeat cannot detect vulnerabilities that require repetition for profit. CPMM-Exploiter-NoIC cannot efficiently allocate resources to function calls more likely to lead to exploits.

Answer to RQ3: *CPMM-Exploiter's* two-step approach outperforms grammar-based fuzzing without repetition and grammar-based fuzzing without invariant checks.

TABLE VII REAL-WORLD EXPLOITS GENERATED BY CPMM-Exploiter. As THESE VULNERABILITIES HAVE NOT BEEN PATCHED, WE DENOTE THEM WITH NUMBERS TO AVOID PROVIDING DETAILS FOR EXPLOITABLE VULNERABILITIES.

| Exploit Number | Invariant Broken | Nework | Maximum Achievable Profit in USD | % Pair Asset |
|-------------------|---------------------|--------|--|--------------|
| 1 | 1 | BSC | 0.65 | 107.87 |
| 2 | 1 | BSC | 186.41 | 1.68 |
| 3 | 2 | BSC | 398.00 | 0.74 |
| 4 | 1 | BSC | 125.00 | 2.86 |
| 5 | 1 | BSC | 0.37 | 0.01 |
| 6 | 1 | BSC | 4796.00 | 189.66 |
| 7 | 1 | BSC | 282.76 | 2.40 |
| 8 | 1 | BSC | 76.64 | 3.54 |
| 9 | 1 | BSC | 1.87 | 0.05 |
| 10 | 1 | BSC | 1.55 | 0.19 |
| 11 | 2 | ETH | 191.06 | 0.24 |
| 12 | 1 | ETH | 338.54 | 1.66 |
| 13 | 2 | ETH | 30.17 | 0.28 |
| 14 | 1 | ETH | 26.81 | 0.39 |
| 15 | 1 | ETH | 5597.62 | 99.85 |
| 16 | 1 | ETH | 16.76 | 1.90 |
| 17 | 1 | ETH | 0.03 | 0.06 |
| 18 | 2 | ETH | 811.15 | 56.93 |

E. Effectiveness in the Real World

To demonstrate the effectiveness of *CPMM-Exploiter* in detecting undiscovered CPMM composability bugs in the real world, we ran *CPMM-Exploiter* on the latest blocks of Ethereum and Binance networks from 22nd January 2024 to 9th February 2024. For each network, we extracted a list of DEXes following the Uniswap V2 protocol with more than 1,000 USD worth of native tokens or stablecoins. 23,701 contracts from Ethereum and 20,607 contracts from Binance satisfied the criteria. We ran *CPMM-Exploiter* with a timeout of 20 minutes per DEX.

Table VII contains the summary of exploits generated by CPMM-Exploiter. Please note that we represent them with exploit numbers instead of token names or addresses because these vulnerabilities have not yet been patched. In summary, CPMM-Exploiter could generate 18 exploits by exploiting CPMM composability bugs in the real world, resulting in a total 12.9K USD profit. To demonstrate the impact of each vulnerability, we report the maximum achievable profit (column 4) and the proportion of maximum achievable profit to the pair stablecoin balance before the exploit (column 5). As CPMM-Exploiter halts when it finds a profit-generating transaction and does not proceed to maximize profit, we manually adjusted some parameters of the exploit (e.g., initial token balance or the number of repetitions) to maximize the profit. Profit maximization was straightforward for all exploits.

Interestingly, *CPMM-Exploiter* was able to generate two exploits that gained more profit than the entire token balance of the DEX (exploits 1 and 6). Such exploits were possible because the exploit also withdraws tokens owned by the vulnerable token contract. Moreover, *CPMM-Exploiter* generated one exploit that almost drains the DEX (exploit 15) and one

| 1 | <pre>function transfer(address addr, uint amount) external {</pre> |
|---|--|
| 2 | if (addr == DEX_ADDR) { |
| 3 | // missing interval check |
| 4 | <pre>maintainPrice();</pre> |
| 5 | } |
| 6 | // transfer tokens |
| 7 | <pre>balances[msg.sender] -= amount;</pre> |
| 8 | <pre>balances[addr] += amount;</pre> |
| 9 | } |
| 0 | <pre>function maintainPrice() internal {</pre> |
| 1 | <pre>// decrease pair token balance by 10%</pre> |
| 2 | balances[DEX_ADDR] = |
| 3 | balances[DEX_ADDR] * 9 / 10; |
| 4 | 1 |

Fig. 7. Vulnerable code snippet from Exploit 15.

exploit that withdraws more than half of the tokens owned by the DEX (exploit 18).

Other exploits withdraw only a small portion of tokens owned by the DEX. Oftentimes, the profits were restricted because *CPMM-Exploiter* could only execute invariant-breaking call sequences once due to interval checking mechanisms that prevent functions from executing multiple times in one transaction. If we remove the constraint that the exploit has to be completed in one transaction, *CPMM-Exploiter* can repeatedly withdraw small portions in intervals, which may ultimately drain the DEX. Thus, the remaining vulnerabilities may also be critical.

Answer to RQ4: *CPMM-Exploiter* can generate impactful real-world exploits.

VIII. CASE STUDY

This section reports case studies for two real-world CPMM composability bug that CPMM-Exploiter detected.

A. Exploit 15: Breaking Invariant 1

Exploit 15 is a real-world bug that breaks **Invariant 1**. This bug is caused due to a missing interval check in the token contract. When a transfer is made to the DEX, the vulnerable token, henceforth Token 15, removes a portion of the DEX token balance. Unfortunately, the token developers did not include an interval check for this behavior, allowing an attacker to drain the DEX's stablecoin asset by repeatedly triggering the vulnerability. Consequently, an attacker can almost drain the DEX's stablecoin balance, which is worth 5597.62 USD.

The simplified version of vulnerable code is shown in Figure 7. Whenever users sell Token 15 to the DEX, the price of Token 15 falls, which may pressure users to sell Token 15. Thus, Token 15 has a function that periodically burns a share of Token 15 in the DEX to maintain its price (i.e., maintainPrice() function in lines 10-14). However, the token developer did not implement an interval check before executing the maintainPrice() function. This bug enables an attacker to burn an arbitrary portion of the DEX token balance, breaking **Invariant 1**, and leverage the vulnerability to drain stablecoins from the DEX.

B. Exploit 18: Breaking Invariant 2

Exploit 18 is a real-world bug that breaks **Invariant 2**. This token, henceforth Token 18, rewards users whenever

```
function getRate() public {
2
       return totalTokenSupply / totalShareSupply;
3
4
   function transfer(address addr, uint amount) external {
5
       // transfer tokens
6
       uint shareAmount = amount / getRate();
       shareBalances[msg.sender] -= shareAmount;
8
       shareBalances[addr] += shareAmount;
9
10
   function maintainToken() external {
11
       // missing interval check
12
       // check that caller is a token owner
13
       require(shareBalances[msg.sender] > minAmount);
14
       // proportionally decrease variables
totalTokenSupply = totalTokenSupply * 9 / 10;
15
       totalShareSupply = totalShareSupply * 9 / 10;
16
17
          award caller
       shareBalances[msg.sender] += awareAmount;
18
19
   function balanceOf(address addr) external {
20
       return shareBalances[addr] * getRate();
21
22 }
```

Fig. 8. Vulnerable code snippet from Exploit 18.

they call a maintenance function. Unfortunately, this reward can be repeatedly reaped, allowing an attacker to accumulate a significant amount of Token 18. We concluded that this vulnerability can be leveraged to drain around 57% of relevant DEX's stablecoin balance, which is worth 811.15 USD.

Figure 8 shows the simplified version of Token 18. This token manages its balances using two variables, totalTokenSupply and totalShareSupply. As more users join the market for Token 18, the two variables will increase and may result in integer overflow. To prevent such a situation, Token 18 has to decrease the two variables periodically. Such maintenance function is implemented in lines 10 to 19 in Figure 8. Unlike other tokens that commonly embed these functions in a commonly called function, such as transfer, Token 18 adopts a different approach where it incentivizes users to directly call the function with rewards. However, the developers did not limit the number of times this function call can be called. Thus, an attacker can repeatedly call the maintainToken () function to accumulate a significant amount of Token 18 without cost, breaking Invariant 2. This vulnerability can be leveraged to extract a sizable amount of stablecoins from DEXes trading Token 18.

IX. DISCUSSION

A. Responsible Disclosure for Smart Contract Vulnerabilities

We attempted to notify the token maintainers about bugs found by *CPMM-Exploiter*. However, these tokens were no longer maintained or the maintainers were unreachable. As a result, these vulnerabilities remain unpatched, so we could not open-source our code. Currently, we are in the process of reporting our findings to CISA (Cybersecurity & Infrastructure Security Agency). We also discussed this issue with SEAL 911, which is a group of security researchers who focus on blockchain security. We hope for an ethical way to manage vulnerabilities in projects with no active maintainers.

B. Limitations

CPMM-Exploiter's main limitations are twofold. First, the current implementation of CPMM-Exploiter only supports

Uniswap V2 DEXes. We believe *CPMM-Exploiter* can be extended to support other CPMM implementations with more development effort. Second, *CPMM-Exploiter* does not utilize all the functions available from token contracts when generating testcases. Currently, we only support burn-related functions and zero-argument functions. To support more functions, additional analysis is necessary to infer the purpose of those functions and generate suitable arguments.

C. Threats to Validity

One threat to the validity is evaluation datasets. Since we suggest a new category of vulnerability, we could not evaluate our system on datasets used in previous works [4]–[6]. Instead, we used a subset of a popular dataset, DeFiHackLabs [11], as one of our evaluation datasets. Since we manually selected vulnerabilities that fall into the category of CPMM composability bugs, some bias and subjectivity may have been introduced. Furthermore, there are significantly fewer cases that break **Invariant 2** than cases that break **Invariant 1** in both DeFiHackLabs and BlockSec datasets. Such imbalance may have skewed the detection results for some tools.

X. RELATED WORK

Numerous tools detect smart contract vulnerabilities. Some utilize static analysis techniques, such as model checking [2], [17]–[19] and symbolic execution [20]–[26]. While others utilize dynamic analysis techniques, most notably fuzzing [4], [5], [13], [27]–[36]. Recent works also utilize machine learning [37]–[44], including Large Language Models [14].

Multi-contract vulnerability detection. Several works were proposed to detect multi-contract vulnerabilities. Some focus on detecting commonly appearing ones, such as reentrancy and delegatecall-related vulnerabilities [45]–[47], while some aim to detect a wide variety of vulnerabilities [4], [5], [13]. In particular, ItyFuzz explores various combinations of contract states through fuzzing with snapshots, and Echidna uses a static analyzer, Slither, to extract useful information before fuzzing. Although CPMM composability bugs can theoretically be detected with such methods, our evaluation indicates that a generic approach is ineffective. Since CPMM composability bugs are closely tied to the business logic of contracts and sometimes require a long sequence of function calls for exploitation, a targeted approach is more suitable, as demonstrated by *CPMM-Exploiter*.

Automatic exploit generation for smart contracts. Some works propose systems that automatically generate exploits. EthPloit [48] generates exploits for single contracts based on fuzzing. FlashSyn [49] utilizes counterexample driven approximation to generate flashloan attacks. Recently, Gritti et al. [47] designed a system that analyzes multiple contracts to automatically detect and exploit confused deputy vulnerabilities. *CPMM-Exploiter* pursues the same goal of exploit generation, but it targets a vulnerability that the aforementioned tools cannot detect.

XI. CONCLUSION

Smart contracts enable many novel and innovative applications in DeFi. At the same time, they introduce new vulnerabilities, such as CPMM composability bugs. To mitigate this issue, we propose *CPMM-Exploiter*, a novel grammar-based fuzzing tool that automatically detects and generates an endto-end exploit for CPMM composability bugs within minutes. *CPMM-Exploiter* obtained recalls of 0.91 and 0.89 on realworld exploits of CPMM composability bugs, while the best recalls from five baselines were 0.36 and 0.58. Furthermore, *CPMM-Exploiter* produced 18 exploits with a total profit of 12.9K USD when run on Ethereum and Binance networks.

REFERENCES

- S. P. Lee, "Market share of decentralized crypto exchanges, by trading volume." https://www.coingecko.com/research/publications/ decentralized-crypto-exchanges-market-share, 2023.
- [2] K. Babel, P. Daian, M. Kelkar, and A. Juels, "Clockwork finance: Automated analysis of economic security in smart contracts," pp. 2499– 2516, IEEE, 2023.
- [3] N. Mutual, "How was bra token exploited?." https://medium.com/ neptune-mutual/how-was-bra-token-exploited-24ff323249d, 2023.
- [4] G. Grieco, W. Song, A. Cygan, J. Feist, and A. Groce, "Echidna: effective, usable, and fast fuzzing for smart contracts," in *Proceedings* of the 29th ACM SIGSOFT international symposium on software testing and analysis, pp. 557–560, 2020.
- [5] C. Shou, S. Tan, and K. Sen, "Ityfuzz: Snapshot-based fuzzer for smart contract," pp. 322–333, 2023.
- [6] Q. Kong, J. Chen, Y. Wang, Z. Jiang, and Z. Zheng, "Defitainter: Detecting price manipulation vulnerabilities in defi protocols," pp. 1144– 1156, Association for Computing Machinery, 2023.
- [7] J. Feist, G. Grieco, and A. Groce, "Slither: A static analysis framework for smart contracts," in 2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB), IEEE, May 2019.
- [8] "Mythril." https://github.com/Consensys/mythril, 2017. GitHub repository.
- Q. W. Security, "Shadowfi \$301k burn function exploit analysis—quilaudits." https://medium.com/quillhash/ shadowfi-301k-burn-function-exploit-analysis-quillaudits-45a17ce04193, 2022.
- [10] "Blocksec." https://blocksec.com/, 2024.
- [11] "Defihacklabs." https://github.com/SunWeb3Sec/DeFiHackLabs, 2020. GitHub repository.
- [12] "Foundry." https://github.com/foundry-rs/foundry, 2024. GitHub repository.
- [13] M. Rodler, D. Paaßen, W. Li, L. Bernhard, T. Holz, G. Karame, and L. Davi, "Ef/cf: High performance smart contract fuzzing for exploit generation," arXiv preprint arXiv:2304.06341, 2023.
- [14] Y. Sun, D. Wu, Y. Xue, H. Liu, H. Wang, Z. Xu, X. Xie, and Y. Liu, "Gptscan: Detecting logic vulnerabilities in smart contracts by combining gpt with program analysis," *Proc. IEEE/ACM ICSE*, 2024.
- [15] Z. Zhang, Z. Lin, M. Morales, X. Zhang, and K. Zhang, "Your exploit is mine: Instantly synthesizing counterattack smart contract," in *32nd* USENIX Security Symposium (USENIX Security 23), pp. 1757–1774, 2023.
- [16] "Blocksec twitter." https://twitter.com/BlockSecTeam/status/ 1624077078852210691, 2023.
- [17] Y. Mo, J. Chen, Y. Wang, and Z. Zheng, "Toward automated detecting unanticipated price feed in smart contract," pp. 1257–1268, Association for Computing Machinery, 2023.
- [18] F. Ma, M. Ren, L. Ouyang, Y. Chen, J. Zhu, T. Chen, Y. Zheng, X. Dai, Y. Jiang, and J. Sun, "Pied-piper: Revealing the backdoor threats in ethereum erc token contracts," ACM Transactions on Software Engineering and Methodology, vol. 32, no. 3, pp. 1–24, 2023.
- [19] J. Ye, M. Ma, Y. Lin, Y. Sui, and Y. Xue, "Clairvoyance: Crosscontract static analysis for detecting practical reentrancy vulnerabilities in smart contracts," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Companion Proceedings*, pp. 274– 275, 2020.

- [20] P. Bose, D. Das, Y. Chen, Y. Feng, C. Kruegel, and G. Vigna, "Sailfish: Vetting smart contract state-inconsistency bugs in seconds," pp. 161– 178, IEEE, 2022.
- [21] A. Ghaleb, J. Rubin, and K. Pattabiraman, "Achecker: Statically detecting smart contract access control vulnerabilities," in 2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE), pp. 945– 956, IEEE, 2023.
- [22] M. Mossberg, F. Manzano, E. Hennenfent, A. Groce, G. Grieco, J. Feist, T. Brunson, and A. Dinaburg, "Manticore: A user-friendly symbolic execution framework for binaries and smart contracts," in 2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 1186–1189, IEEE, 2019.
- [23] S. So, S. Hong, and H. Oh, "{SmarTest}: Effectively hunting vulnerable transaction sequences in smart contracts through language {Model-Guided} symbolic execution," in 30th USENIX Security Symposium (USENIX Security 21), pp. 1361–1378, 2021.
- [24] J. Chen, X. Xia, D. Lo, J. Grundy, X. Luo, and T. Chen, "Defectchecker: Automated smart contract defect detection by analyzing evm bytecode," *IEEE Transactions on Software Engineering*, vol. 48, no. 7, pp. 2189– 2207, 2021.
- [25] H. Wang, Y. Liu, Y. Li, S.-W. Lin, C. Artho, L. Ma, and Y. Liu, "Oracle-supported dynamic exploit generation for smart contracts," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 3, pp. 1795–1809, 2020.
- [26] S.-W. Lin, P. Tolmach, Y. Liu, and Y. Li, "Solsee: a source-level symbolic execution engine for solidity," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 1687–1691, 2022.
- [27] "Detecting state inconsistency bugs in dapps via on-chain transaction replay and fuzzing," ISSTA 2023 - Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis, pp. 298–309, 7 2023.
- [28] C. F. Torres, A. K. Iannillo, A. Gervais, and R. State, "Confuzzius: A data dependency-aware hybrid fuzzer for smart contracts," pp. 103–119, IEEE, 2021.
- [29] J. Choi, D. Kim, S. Kim, G. Grieco, A. Groce, and S. K. Cha, "Smartian: Enhancing smart contract fuzzing with static and dynamic data-flow analyses," *Proceedings - 2021 36th IEEE/ACM International Conference* on Automated Software Engineering, ASE 2021, pp. 227–239, 2021.
- [30] M. Olsthoorn, D. Stallenberg, A. Van Deursen, and A. Panichella, "Syntest-solidity: Automated test case generation and fuzzing for smart contracts," in *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings*, pp. 202–206, 2022.
- [31] W. Chen, Z. Sun, H. Wang, X. Luo, H. Cai, and L. Wu, "Wasai: Uncovering vulnerabilities in wasm smart contracts," in *Proceedings of* the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis, pp. 703–715, 2022.
- [32] N. Parasaram, E. T. Barr, S. Mechtaev, and M. Böhme, "Precise datadriven approximation for program analysis via fuzzing," in 2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE), pp. 611–623, IEEE, 2023.
- [33] B. Jiang, Y. Liu, and W. K. Chan, "Contractfuzzer: Fuzzing smart contracts for vulnerability detection," in *Proceedings of the 33rd ACM/IEEE international conference on automated software engineering*, pp. 259– 269, 2018.
- [34] I. Ashraf, X. Ma, B. Jiang, and W. K. Chan, "Gasfuzzer: Fuzzing ethereum smart contract binaries to expose gas-oriented exception security vulnerabilities," *IEEE Access*, vol. 8, pp. 99552–99564, 2020.
- [35] Y. Huang, B. Jiang, and W. K. Chan, "Eosfuzzer: Fuzzing eosio smart contracts for vulnerability detection," in *Proceedings of the 12th Asia-Pacific Symposium on Internetware*, pp. 99–109, 2020.
- [36] T. D. Nguyen, L. H. Pham, J. Sun, Y. Lin, and Q. T. Minh, "sfuzz: An efficient adaptive fuzzer for solidity smart contracts," in *Proceedings of* the ACM/IEEE 42nd International Conference on Software Engineering, pp. 778–788, 2020.
- [37] Y. Chen, Z. Sun, Z. Gong, and D. Hao, "Improving smart contract security with contrastive learning-based vulnerability detection," in 2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE), pp. 940–940, IEEE Computer Society, 2024.
- [38] Z. Yang, J. Keung, M. Zhang, Y. Xiao, Y. Huang, and T. Hui, "Smart contracts vulnerability auditing with multi-semantics," in 2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC), pp. 892–901, IEEE, 2020.

- [39] H. H. Nguyen, N.-M. Nguyen, C. Xie, Z. Ahmadi, D. Kudendo, T.-N. Doan, and L. Jiang, "Mando: Multi-level heterogeneous graph embeddings for fine-grained detection of smart contract vulnerabilities," in 2022 IEEE 9th International Conference on Data Science and Advanced Analytics (DSAA), pp. 1–10, IEEE, 2022.
- [40] Z. Zhang, Y. Lei, M. Yan, Y. Yu, J. Chen, S. Wang, and X. Mao, "Reentrancy vulnerability detection and localization: A deep learning based two-phase approach," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pp. 1– 13, 2022.
- [41] H. Wu, Z. Zhang, S. Wang, Y. Lei, B. Lin, Y. Qin, H. Zhang, and X. Mao, "Peculiar: Smart contract vulnerability detection based on crucial data flow graph and pre-training techniques," in 2021 IEEE 32nd International Symposium on Software Reliability Engineering (ISSRE), pp. 378–389, IEEE, 2021.
- [42] M. Li, X. Ren, H. Fu, Z. Li, and J. Sun, "Convmhsa-scvd: Enhancing smart contract vulnerability detection through a knowledge-driven and data-driven framework," in 2023 IEEE 34th International Symposium on Software Reliability Engineering (ISSRE), pp. 578–589, IEEE, 2023.
- [43] F. Luo, R. Luo, T. Chen, A. Qiao, Z. He, S. Song, Y. Jiang, and S. Li, "Scvhunter: Smart contract vulnerability detection based on heterogeneous graph attention network," in 2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE), pp. 954–954, IEEE Computer Society, 2024.
- [44] H. Liu, C. Liu, W. Zhao, Y. Jiang, and J. Sun, "S-gram: Towards semantic-aware security auditing for ethereum smart contracts," in *Proceedings of the 33rd ACM/IEEE international conference on automated software engineering*, pp. 814–819, 2018.
- [45] Y. Xue, J. Ye, W. Zhang, J. Sun, L. Ma, H. Wang, and J. Zhao, "xfuzz: Machine learning guided cross-contract fuzzing," *IEEE Transactions on Dependable and Secure Computing*, 2022.
- [46] Z. Liao, Z. Zheng, X. Chen, and Y. Nan, "Smartdagger: a bytecodebased static analysis approach for detecting cross-contract vulnerability," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 752–764, 2022.
- [47] F. Gritti, N. Ruaro, R. McLaughlin, P. Bose, D. Das, I. Grishchenko, C. Kruegel, and G. Vigna, "Confusum contractum: Confused deputy vulnerabilities in ethereum smart contracts," in *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 1793–1810, 2023.
- [48] Q. Zhang, Y. Wang, J. Li, and S. Ma, "Ethploit: From fuzzing to efficient exploit generation against smart contracts," pp. 116–126, IEEE, 2020.
- [49] Z. Chen, S. M. Beillahi, and F. Long, "Flashsyn: Flash loan attack synthesis via counter example driven approximation," arXiv preprint arXiv:2206.10708, 2022.