

Formal Analysis of Composable DeFi Protocols

Palina Tolmach^{1,2}, Yi Li², Shang-Wei Lin², and Yang Liu²

¹ Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore

² Nanyang Technological University, Singapore
{palina001,yi_li,shang-wei.lin,yangliu}@ntu.edu.sg

Abstract. Decentralized finance (DeFi) has become one of the most successful applications of blockchain and smart contracts. The DeFi ecosystem enables a wide range of crypto-financial activities, while the underlying smart contracts often contain bugs, with many vulnerabilities arising from the unforeseen consequences of composing DeFi protocols together. In this paper, we propose a formal process-algebraic technique that models DeFi protocols in a compositional manner to allow for efficient property verification. We also conduct a case study to demonstrate the proposed approach in analyzing the composition of two interacting DeFi protocols, namely, Curve and Compound. Finally, we discuss how the proposed modeling and verification approach can be used to analyze financial and security properties of interest.

1 Introduction

With more than \$12 billions currently locked inside, decentralized finance (DeFi) becomes one of the most prominent applications of the blockchain technology [10]. DeFi protocols implement various financial applications, including analogs of traditional-finance use cases, such as lending [21], exchange [12,4], investment [2], etc. These protocols give users access to digital assets, e.g., *tokens*, and expose them to the cryptocurrency market. As an example, stablecoins are cryptocurrencies providing minimum volatility by pegging their prices to fiat money, real-world commodity, or a more “stable” cryptocurrency, such as ETH [29].

At the same time, billions of dollars stored in DeFi stimulate the invention of new security attacks. Unlike other smart contracts applications, the security of DeFi protocols can be compromised by not only software vulnerabilities but also unforeseen movements in the cryptocurrency market or arbitrage and speculation opportunities. For example, an attacker drained \$2M of funds from the (twice audited) Akropolis DeFi platform [14] through a well-studied reentrancy vulnerability [15,27,35]. As another example, in March 2020, the network congestion caused by market instability led to major disruptions and losses in some of DeFi protocols during the events of so-called “Black Thursday” [31].

A distinctive feature of DeFi applications is their similarity to the pieces of so-called *Money Legos* [40]. In other words, the design of DeFi protocols often facilitates interoperability between them including the support of tokens issued

by different DeFi platforms. While the composability of DeFi applications enables the construction of a decentralized financial ecosystem, integrations between protocols contribute to the creation of new attack vectors. For example, a recent attack on the Harvest yield aggregation protocol [1] was made possible due to its dependence on the prices reported by the Curve decentralized exchange protocol [12]. By performing a \$17M trade in Curve, the attacker could indirectly manipulate the price of tokens in Harvest, obtaining \$24M of protocol funds [13]. An established way to rigorously verify correctness of safety-critical systems, including smart contracts, is to employ formal analysis [43]. In the field of DeFi, security audits often involve formal analysis, but usually focusing only on the verification of individual protocols. Yet, the “money-lego” structure of the DeFi ecosystem demands compositional analysis, which allows reasoning about the possible interplay between DeFi protocols and their impact on each other.

To model and analyze the behaviors of composable DeFi protocols, we formulate general formal models of components of DeFi protocols, particularly, *tokens* and *pools*. Based on their actual implementations, we develop process-algebraic models of two widely used DeFi protocols: a decentralized exchange—Curve Finance [12], and a lending protocol—Compound [21]. In addition, we formally model the behavior of the USDC stablecoin. Using the developed model, we formally verify some of the (already stated) relevant properties of the protocols under consideration. Finally, we formulate safety and correctness properties that are expected to hold throughout the interactions between the considered protocols.

2 Background

In this section, we provide necessary background for the rest of the paper.

2.1 DeFi Protocols

We consider two common types of DeFi protocols: *decentralized exchanges* (DEX) and *protocols for loanable funds* (PLF), a.k.a. *lending protocols*.

Decentralized Exchanges DEX is one of the first and most popular DeFi applications. While a centralized exchange has to match a seller with a specific buyer, a typical DEX uses smart contracts to execute trades asynchronously [11,5]. A *pool*, implemented using smart contracts, stores the reserves of two or more types of tokens and automatically determines the exchange rate between these tokens.

A common way to determine the exchange rate between assets within a DEX pool is by maintaining a *constant-product* and/or *constant-sum* invariant between the values of the tokens contained in the pool. Essentially, the invariant implies that if a user trades t_1 for t_2 , the price of t_1 in the pool goes down, while the price of t_2 increases. This model, therefore, provides an arbitrage opportunity for the users of DEXes, encouraging them to deposit or sell tokens of type t_2 at a higher price, which thereby restores the balance between tokens.

Lending Protocols PLFs [17,30] rely on smart contracts to mediate token lending and borrowing. Different from DEXes, lending pools collect assets of (usually) one token type from liquidity providers. In return, the depositors are given pool tokens with the value constantly increasing from the interest fees paid by borrowers. The interest rate for borrowers depends on a chosen interest rate model and is usually decided by the utilization rate—the ratio between the supply and demand of the pool. To protect a protocol from the cryptocurrency volatility, the borrower is also supposed to supply a collateral (e.g., in ETH or a stablecoin) that is bigger than the amount of borrowed funds by at least a collateralization ratio.

2.2 Formal Modeling and Verification

Communicating Sequential Process (CSP) [18] is a formal language for describing patterns of interaction for concurrent systems [34]. A CSP model contains a set of synchronized or interleaving processes, each of which consists of a sequence of ordered events. For instance, a process P , with an event a followed by another event b , can be written as “ $P = a \rightarrow b$ ”. Multiple processes can be composed either sequentially or in parallel. Sequential composition of two processes P and Q (denoted by $P; Q$) acts as P first, and acts as Q upon the termination of P . The two processes can also be composed in parallel and synchronized on an event X ($P \mid [X] \mid Q$), or asynchronously ($P \parallel Q$). Finally, a process Q can interrupt another process P when event e happens ($P \nabla e \rightarrow Q$). The detailed syntax are summarized as follows.

$$\begin{array}{l}
 P := \quad STOP \quad | \quad SKIP \quad | \quad e \rightarrow P \quad | \quad P \square P \\
 \quad | \quad P \sqcap P \quad | \quad P \parallel P \quad | \quad P \mid [X] \mid P \quad | \quad P \setminus X \\
 \quad | \quad P; P \quad | \quad \text{if } b \text{ then } P \text{ else } P \quad | \quad P \nabla P
 \end{array}$$

CSP# [38,28] is an extension to CSP with embedding of data operations. CSP# combines high-level compositional operators from process algebra with program-like codes, which makes the language much more expressive. The models and properties specified in CSP# can be checked using Process Analysis Toolkit (PAT) [36,37,28], which is a framework for specification, simulation, and verification of concurrent and real-time systems. PAT supports event-based compositional models and efficient LTL model checking with various fairness assumptions. Model checking [9] is widely used to verify state-transition systems of one or several interacting smart contracts against a temporal logic specification [39]. In this work, we use the model checker of PAT to verify the properties of individual and interacting DeFi protocols, as described in Sect. 4.

One unique feature of PAT is that it allows users to define static functions and data types as C# libraries. These user-defined C# libraries are built as DLL files and are loaded during execution, which compensates for the common deficiencies of model checkers on complex data operations and data types. We utilize this capability and implement complex mathematical computations underlying the token price calculation in C#. Finally, the translation from high-level smart contract programming languages, such as Vyper and Solidity, to C# is straightforward.

3 Methodology

To reason about a system of interconnected protocols, we use a *process-algebraic* approach to model various components of the DeFi ecosystem. First, we formally define the main components of DeFi applications along with the environment models. Then, we model two widely used Ethereum DeFi protocols and their interactions using CSP#, by translating the major smart contract functions into CSP, in a similar fashion to some of the previous work [33,22].

3.1 Protocol Formal Modeling

In this section, we propose formal definitions for the two key constituents of lending and exchange DeFi protocols: *token* and *pool*.¹ The behaviors of the aforementioned objects can be formalized as state transition systems, and we focus on their states here.

We model the *states* of users, smart contracts and the environment variables (e.g., `block.number`) as global variables in the CSP# model. Functions, on the other hand, are translated into *processes*. Inspired by [6], we assume a set of blockchain *users* (\mathbb{U}) and a set of *tokens* (\mathbb{T}). Tokens are programmable assets managed by smart contracts [8]. The majority of tokens used in DeFi protocols, except the native platform cryptocurrency ETH, are implemented in the form of a contract conforming to the ERC20 standard [41]. ERC20 regulates the development of fungible tokens by specifying the interface of the corresponding smart contract, i.e., public functions and events that it should emit during executions. In accordance with the standard, we define tokens in Definition 1.

Definition 1 (Token). *A token $t \in \mathbb{T}$ is a tuple $(\mathbb{U}, TS, B, A, \mathbb{F})$, where \mathbb{U} is a set of users, $TS \in \mathbb{Z}_{\geq 0}$ is the total supply, $B : \mathbb{U} \mapsto \mathbb{Z}_{\geq 0}$ is the mapping from users to their token balances, $A : \mathbb{U} \times \mathbb{U} \mapsto \mathbb{Z}_{\geq 0}$ specifies the allowances, i.e., amounts of token that a user is allowed to spend from another user’s balance, and \mathbb{F} is the set of state-changing functions modifying the state of the token.*

Given a token $t \in \mathbb{T}$, we use $t.TS$ to denote its total supply and $t.A$ to denote its allowances, and so on. The balance invariant of t satisfies the formula: $t.TS = \sum_{u \in \mathbb{U}} t.B(u)$. \mathbb{F} includes functions changing the values of A , B , and TS , e.g., `approve()`, `transfer()`, `transferFrom()`, `mint()`, `burn()`, etc. Figure 1 demonstrates a partial implementation of the state of the USDC token in a model with N participants. Formally, the state of a user account u is the set of balances of tokens in the user’s possession, i.e., $\{t.B(u) \mid u \in \mathbb{U} \text{ and } t \in \mathbb{T}\}$.

Definition 2 specifies pools, which are smart contracts used to aggregate a number of tokens.

Definition 2 (Pool). *A pool P is a tuple $(\mathbb{T}_P, \mathbb{T}_P^R, \mathbb{F}_P)$, where $\mathbb{T}_P \subset \mathbb{T}$ is a set of pool tokens of P ,² $\mathbb{T}_P^R \subset \mathbb{T}$ is a set of liquidity tokens supported by P , and \mathbb{F}_P is a set of functions $\{(\mathbb{T}_P \times \mathbb{T}_P^R) \mapsto (\mathbb{T}_P \times \mathbb{T}_P^R)\}_i$ changing the state of P .*

¹ Depending on the application, pools are also referred to as *markets*, *vaults*, or *pairs*.

² Most of the pools in DeFi support a single pool token.

```

var USDC_balances[N]:{0..} = [b1, b2, ..., bN];
var USDC_allowed[N][N]:{0..};
var USDC_totalSupply = ts;
...

```

Fig. 1: USDC token state implementation in CSP#.

```

Curve_addLiquidity(uamounts, min_mint_amount, sender) = atomic {
...
  USDC_transferFrom(user, curveDeposit, uamounts, ...);
...
  cUSDC_mint(uamounts, curveDeposit);
...
  cUSDC_approve(curveSwap, cAmounts[USDC], curveDeposit);
  Curve_swap_addLiquidity(cAmounts, min_mint_amount, sender);
  cCrv_transfer(user, cCrv_mintAmount, curveDeposit);}

```

Fig. 2: CSP# process for the add_liquidity() function of a Curve pool.

Depending on the protocol application, liquidity tokens $\mathbb{T}_P^R \subset \mathbb{T}$ are used to facilitate decentralized token exchange, lending, investments, or other DeFi use cases. Liquidity pools in DEX usually hold liquidity in two or more types of tokens [4,12,3], while lending protocol [21] or yield aggregator [2] pools typically accept a single type of token as input. In both cases, the users depositing tokens (a.k.a. *liquidity providers*) receive a certain amount of pool tokens (\mathbb{T}_P), which represent user’s share and can be used to redeem the deposit with the earned interests from the pool. \mathbb{F}_P is a set of functions that can change the state of a pool. Figure 2 illustrates the process-algebraic encoding of a state-changing function that implements adding liquidity to a pool from the Curve protocol. To mimic the atomic transaction execution model in Ethereum, we mark state-changing processes as *atomic*, so that their executions cannot be interrupted by an interleaving.

3.2 Protocol Composition

Now, we illustrate how interactions between users and DeFi protocols (*user-protocol*) as well as interactions among different protocols (*protocol-protocol*) can be modeled formally. In both cases, the initiator of a transaction sends a certain amount of tokens to a receiving DeFi protocol and/or receives some tokens from it in return.

In the case of *user-protocol* interaction, we model the behavior of a user by a *sequential composition* (denoted by ‘;’) of one or more processes. These processes correspond to the public state-changing functions of DeFi protocols and tokens invoked by the user. For example, the behavior of a depositor in Curve (i.e., *Curve_Depositor*) is demonstrated in Fig. 3.

```

Curve_Depositor() = USDC_approve(curveDeposit, suppliedTokens, user);
Curve_addLiquidity(suppliedTokens, minMintTokens, user);
...
cCrv_approve(curveDeposit, add, user);
Curve_remove_liquidity_one_coin(add, 0, user, true);

```

Fig. 3: The implementation of Curve depositor behavior in CSP#.

The subject system is then modeled by an *interleaving* (denoted by ‘|||’) of such user processes. For instance, Fig. 4 shows the depositor, exchanger, and borrower processes composed asynchronously, which simulates possible state changes in interacting protocols caused by concurrently acting users. The processes simulating state-changing functions are atomic, i.e., executing without interruption so that the interleaving between user processes can only happen after a state-changing process is finished. We simulate the block mining using a process that increases the value of the block number variable.

```

System() = Curve_Depositor() ||| Curve_Exchange() |||
Compound_Depositor() ||| Compound_Borrower() ||| IncreaseBlockNum();

```

Fig. 4: The analyzed user composition.

The *protocol-protocol* interactions in DeFi smart contracts are external calls to a function of another protocol. Following a similar approach, we model smart contract functions with external calls to other DeFi applications and token contracts as an *atomic sequential composition* of corresponding processes. The sequential composition of two processes ensures that the former process has to finish before the latter can start, so that the model operates similarly as the execution of internal transactions in blockchain. The CSP# representation of a function that implements adding liquidity to a pool of the Curve DeFi protocol is shown in Fig. 2. The communication among users, tokens, and different protocols is simulated via shared global variables, such as token balances shown in Fig. 1.

4 Evaluation

In this section, we evaluate our modeling approach by checking a set of relevant properties on Compound pool of the Curve DEX³ using PAT and report on the results of property verification. We performed the evaluation on a virtual machine with Windows 10, 8GB RAM and 1 CPU core, using PAT version 3.5.0. The virtual machine is running on MacOS Catalina v.10.15.7, 32GB RAM and 2 GHz quad-core Intel Core i5 processor.

The Curve Compound pool allows trading between a pair of stablecoins: USDC and DAI. Under the hood, the Curve pool transfers its USDC and DAI to a lending platform Compound, in exchange for the corresponding Compound’s pool tokens—cUSDC and cDAI. cUSDC and cDAI are, therefore, used for all

³ <https://www.curve.fi/compound>

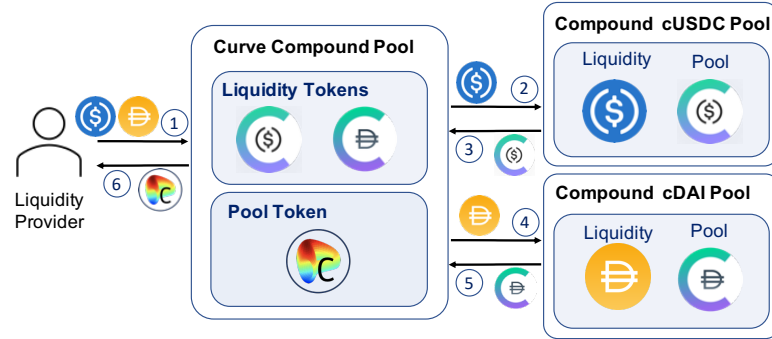


Fig. 5: A scheme of token transfers between Curve Compound pool participants.

the operations within the Curve Compound pool. Figure 5 outlines the process of adding liquidity to the Curve Compound pool: ① a liquidity provider sends USDC and/or DAI to the pool; ② Curve supplies the received USDC to the USDC Compound pool and ③ receives an appropriate amount of cUSDC in return; ④-⑤ the same process is repeated for DAI/cDAI; ⑥ the user receives a certain amount of cCrv—a pool token of the Curve Compound pool.

State-changing actions of interest include providing and removing liquidity in both Curve and Compound, exchanging tokens in Curve, and taking/repaying a loan in Compound. In this paper, we mostly concentrate on the operations that involve USDC: in our model, a liquidity provider on Curve adds and withdraws liquidity in USDC, while Compound depositors and borrowers also perform the corresponding actions with the USDC Compound pool. To model slippage and front-running that can occur in the pool of a DEX, the token exchanges between cUSDC and cDAI in Curve can happen in both directions. We assume that the modeled trading activity reflects the possible changes in the USDC/DAI exchange rate, which we do not explicitly consider otherwise. In addition, since we focus on the operations involving the USDC stablecoin, we simplify the implementation of DAI to basic ERC20 functionality and do not consider the underlying stabilizing mechanism implemented by MakerDAO. We modeled pools and tokens by manually translating their source code written in Solidity or Vyper to CSP# and C# languages supported by PAT. While the translation between high-level languages (e.g., Solidity/C#) is straightforward, data operations and programming constructs supported by CSP# also facilitate translation to a modeling language. The source code of the model can be found in a repository: <https://github.com/polinatolmach/DeFi-csp-models/>.

Based on the defined model, we formulated and verified properties for tokens, individual DEX and lending DeFi applications as well as their composition. LTL formulae and verification results for the properties are demonstrated in Table 1. The first property in Table 1 is the *Balance Invariant* [39]—an important property related to tokens, which we verify for all the tokens involved in the modeled composition: stablecoins (USDC and DAI) and pool tokens (cCrv, cUSDC, etc.). Property (2) in Table 1 is a token-related requirement for a composition of protocols stating that *the positively-valued tokens should never produce zero*

Table 1: A summary of verified properties.

#	Properties	LTL Formulae	Protocols	Results	Stats
(1)	Balance Invariants	$\square((\text{sum}(\text{cCrv_balances}) == \text{cCrv_totalSupply}) \ \&\& \ \text{sum}(\text{cDAI_accountTokens}) == \text{cDAI_totalSupply}) \dots)$	All Tokens	Valid	Time (s): 275.5 s #State: 127337 #Transition: 133763
(2)	Proportional Token Exchange	$\square((\text{suppliedTokens} > 0) \rightarrow \diamond((\text{mintedCTokens} > 0) \ \&\& \ (\text{mintedCCrvTokens} > 0)))$	Curve Compound	Valid	Time (s): 277.7 s #State: 127367 #Transition: 133821
(3)	Non-decreasing Exchange Rate	$\square(\text{prevExchangeRate} \leq \text{newExchangeRate})$	Compound	Valid	Time (s): 277.9s #State: 127337 #Transition: 133763
(4)	Non-negative Profit	$\square(\text{Mint.cUSDC} \rightarrow \square(\text{depositorProfit} \geq 0))$	Compound	Invalid	Time (s): 1.0 s #State: 430 #Transition: 453
(5)	Bounded Loss	$\square(\text{AddLiquidity} \rightarrow \square(\text{depositorLoss} \leq \text{ADMISSIBLE_LOSS}))$	Curve	Invalid	Time (s): 0.5s #State: 177 #Transition: 196

tokens (Proportional Token Exchange) [7]. We verified that this requirement holds for all pairs of tokens involved in the process of adding liquidity to the Curve Compound pool (Fig. 5).

Among the properties of individual protocols, our model allows verification of *the exchange rate of the pool token in Compound being non-decreasing*, meaning that a liquidity provider always receives a guaranteed interest on her deposit (Property (3) in Table 1). For a liquidity provider on Compound, we additionally checked whether her *profit from providing and then redeeming liquidity can only be non-negative* (Property (4) in Table 1). While this requirement holds under normal conditions, it does fail in the event of *overutilization*, i.e., if the pool does not have enough liquidity to repay the depositor. To model *overutilization*, we defined a user who borrows all the available liquidity from the Compound pool. For simplicity, we omitted the collateralization requirements in our model—each loan is assumed to be collateralized using the token that is not considered in the current model (e.g., ETH). Although the simplifications assumed in our model allow reaching overutilization easier than it is in reality, it remains one of the main risks associated with lending protocols [6].

Overutilization in a Compound pool causes a violation of an analogous property defined for a Curve liquidity provider, showing the potentially harmful effects of composability. In other words, the users of both Compound and Curve are not always able to redeem their original deposit back. Considering that a liquidity provider in a DEX can legitimately suffer losses from the *impermanent loss*, the property (5) in Table 1 requires *the loss to be bounded by a certain value*, which we set to 20% of the original deposit. This requirement can also be violated in an anticipated way due to *front-running* and *slippage* caused by massive trades made by other users. The violations of both properties are identified by PAT in sub-second time. Being an on-the-fly model checker, PAT stops constructing and exploring the state space after detecting the violation, which explains the time discrepancy between the verification of properties (1)–(3) and (4), (5). For both violated properties, the reachability analysis in PAT also helps identify the maximum possible losses and profits for both Curve and Compound depositors.

Finally, we confirmed the violation of properties on a locally deployed Ethereum network, assuming the same set of simplifications to smart contracts as in the model.

The performed evaluation demonstrates the suitability of applying process algebra CSP for modeling concurrently acting users and DeFi protocols on blockchain. The results also confirm that model checking can automatically reveal undesirable conditions in the operation of a single DeFi protocol or a composition of those. However, with expanding the composition of modeled users and protocols, the number of states grows exponentially. To combat the state explosion problem, we consider utilizing techniques from the area of compositional verification, such as assume-guarantee reasoning [23,24,25], which we leave for future work.

5 Related Work

The analysis of DeFi protocols is a relatively new field. The existing works often focus on specific types of DeFi protocols or investigate abnormal behaviors observed in the wild. For example, Liu and Szalachowski explored the usage of *oracles* in four major DeFi platforms [26], revealing the operational issues inherent in oracles and common deviations between the real and reported prices.

A number of articles analyze the attack vectors that involve a *flash loan* [32,16], while Wang et al. [42] proposed a framework that allows the identification and classification of flash loan transactions. Their technique is able to detect speculative usage of flash loans and other potentially harmful behaviors.

Several studies explore the operation and properties of DeFi *lending protocols* [17,19,30,6]. Kao et al. [19] utilized agent-based simulations to analyze the market risks faced by the Compound lending protocol users. Stress-tests were performed to demonstrate the scalability of the protocol on a larger borrow size under reasonably volatile conditions. Formal models of lending protocols and their pools were formulated in two recent publications [30,6]. Bartoletti et al. [6] also formulated the fundamental properties of lending pools and typical ways of their interaction with other DeFi protocols. Meanwhile, Perez et al. [30] utilized the abstract formal model of Compound to explore the possibility of liquidations of undercollateralized positions. Different from the discussed publications, this paper formulates a more general formal model of a pool, which can be used to formalize both lending and DEX protocols.

In addition, Klages-Mundt et al. [20] proposed a framework for modeling and classifying stablecoins. The authors also formulated and examined the risks associated with stablecoins and their use in the DeFi ecosystem. The formal model of a token considered in this paper is of a higher level and does not cover its underlying economical mechanism.

Finally, a recent publication by Bernardi et al. [7] proposed a set of invariants that are relevant for individual DeFi protocols, including DEXes and lending platforms. While our study involves verification of some of the invariants proposed in this article, we further extend them to the system of interacting DeFi protocols.

6 Conclusion and Future Work

In this paper, we proposed formal definitions for the main components of DeFi protocols and an approach to model their implementations and interactions in a process-algebraic modeling language. We demonstrated how model checking can automatically verify correctness properties for a composition of DeFi protocols and tokens. The proposed technique successfully identifies the DeFi-specific conditions that cause the violations of these properties.

As future work, we plan to enrich the models to account for functionality related to liquidity-mining and governance mechanisms in the considered DeFi protocols. We would also like to extend the set of properties to cover both security vulnerabilities and the cryptoeconomical aspects of DeFi executions. Finally, to address the state explosion problem, we plan to integrate techniques from the area of compositional verification.

Acknowledgements. This research is partially supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (Award No. 2018-T1-002-069) and Tier 2 (Award No. MOE2018-T2-1-068), and by the National Research Foundation, Singapore, and the Energy Market Authority, under its Energy Programme (EP Award No. NRF2017EWT-EP003-023). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation, Singapore and the Energy Market Authority.

References

1. Harvest Finance. Retrieved October 12, 2020 from <https://harvest.finance/> (2020)
2. Introduction to Yearn - [yearn.finance](https://docs.yearn.finance/). Retrieved November 12, 2020 from <https://docs.yearn.finance/> (2020)
3. Whitepaper - Balancer. Retrieved October 12, 2020 from <https://balancer.finance/whitepaper/> (2020)
4. Adams, H., Zinsmeister, N., Robinson, D.: Uniswap v2 Core Whitepaper. Retrieved October 12, 2020 from <https://uniswap.org/whitepaper.pdf> (2020)
5. Angeris, G., Kao, H.T., Chiang, R., Noyes, C., Chitra, T.: An analysis of uniswap markets (2020), <https://arxiv.org/abs/1911.03380>
6. Bartoletti, M., Chiang, J.H., Lluch-Lafuente, A.: Sok: Lending pools in decentralized finance (2020), <https://arxiv.org/abs/2012.13230>
7. Bernardi, T., Dor, N., Fedotov, A., Grossman, S., Immerman, N., Jackson, D., Nutz, A., Oppenheim, L., Pistiner, O., Rinetzky, N., Sagiv, M., Taube, M., Toman, J.A., Wilcox, J.R.: WIP: Finding bugs automatically in smart contracts with parameterized invariants. Retrieved July 14, 2020 from <https://www.certora.com/pubs/sbc2020.pdf> (2020)
8. Chen, W., Zhang, T., Chen, Z., Zheng, Z., Lu, Y.: Traveling the token world: A graph analysis of Ethereum ERC20 token ecosystem. In: Proceedings of The Web Conference 2020. p. 1411–1421. WWW '20, ACM (2020). <https://doi.org/10.1145/3366423.3380215>

9. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press, Cambridge, MA, USA (2000)
10. Coingape: DeFi success story continues as TVL peaks yet again; hits \$12 billion. Retrieved November 18, 2020 from <https://coinmarketcap.com/ru/headlines/news/defi-success-story-continues-as-tvl-peaks-yet-again-hits-12-billion> (2020)
11. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges (2019), <https://arxiv.org/abs/1904.05234>
12. Egorov, M.: StableSwap - efficient mechanism for Stablecoin liquidity — Curve. fi Whitepaper. Retrieved November 18, 2020 from <https://www.curve.fi/stableswap-paper.pdf> (2020)
13. Finance, H.: Harvest Flashloan Economic Attack Post-Mortem. Retrieved November 18, 2020 from <https://medium.com/harvest-finance/harvest-flashloan-economic-attack-post-mortem-3cf900d65217> (2020)
14. Foxley, W.: DeFi project Akropolis drained of \$2m in DAI. Retrieved November 14, 2020 from <https://www.coindesk.com/defi-project-akropolis-token-pool-drained> (2020)
15. Grossman, S., Abraham, I., Golan-Gueta, G., Michalevsky, Y., Rinetzky, N., Sagiv, M., Zohar, Y.: Online detection of effectively callback free objects with applications to smart contracts. Proc. ACM Program. Lang. pp. 1–28 (dec 2017). <https://doi.org/10.1145/3158136>
16. Gudgeon, L., Perez, D., Harz, D., Livshits, B., Gervais, A.: The decentralized financial crisis (2020), <https://arxiv.org/abs/2002.08099>
17. Gudgeon, L., Werner, S.M., Perez, D., Knottenbelt, W.J.: DeFi protocols for loanable funds: Interest rates, liquidity and market efficiency (2020), <https://arxiv.org/abs/2006.13922>
18. Hoare, C.A.R.: Communicating Sequential Processes. International Series on Computer Science, Prentice-Hall (1985)
19. Kao, H.T., Chitra, T., Chiang, R., Morrow, J.: An Analysis of the Market Risk to Participants in the Compound Protocol. Retrieved November 18, 2020 from https://scfab.github.io/2020/FAB2020_p5.pdf (2020)
20. Klages-Mundt, A., Harz, D., Gudgeon, L., Liu, J.Y., Minca, A.: Stablecoins 2.0. Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (Oct 2020). <https://doi.org/10.1145/3419614.3423261>
21. Leshner, R., Hayes, G.: Compound: The Money Market Protocol — Whitepaper. Retrieved November 18, 2020 from <https://compound.finance/documents/Compound.Whitepaper.pdf> (2020)
22. Li, X., Su, C., Xiong, Y., Huang, W., Wang, W.: Formal verification of BNB smart contract. In: Proceedings of the BIGCOM. pp. 74–78 (Aug 2019). <https://doi.org/10.1109/BIGCOM.2019.00021>
23. Lin, S., Andre, E., Liu, Y., Sun, J., Dong, J.: Learning assumptions for compositional verification of timed systems. IEEE Transactions on Software Engineering (02), 137–153 (feb 2014). <https://doi.org/10.1109/TSE.2013.57>
24. Lin, S.W., Liu, Y., Sun, J., Dong, J.S., André, É.: Automatic compositional verification of timed systems. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012: Formal Methods. pp. 272–276. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
25. Lin, S.W., Sun, J., Nguyen, T.K., Liu, Y., Dong, J.S.: Interpolation guided compositional verification. In: Proceedings of the 30th IEEE/ACM International Confer-

- ence on Automated Software Engineering. p. 65–74. ASE '15, IEEE Press (2015). <https://doi.org/10.1109/ASE.2015.33>
26. Liu, B., Szalachowski, P.: A first look into DeFi oracles (2020), <https://arxiv.org/abs/2005.04377>
 27. Liu, C., Liu, H., Cao, Z., Chen, Z., Chen, B., Roscoe, B.: ReGuard: Finding reentrancy bugs in smart contracts. In: Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings. pp. 65–68. ACM (2018). <https://doi.org/10.1145/3183440.3183495>
 28. Liu, Y., Sun, J., Dong, J.S.: Pat 3: An extensible architecture for building multi-domain model checkers. In: 2011 IEEE 22nd International Symposium on Software Reliability Engineering. pp. 190–199 (2011). <https://doi.org/10.1109/ISSRE.2011.19>
 29. Moin, A., Sierer, E.G., Sekniqi, K.: A classification framework for stablecoin designs (2019), <https://arxiv.org/abs/1910.10098>
 30. Perez, D., Werner, S.M., Xu, J., Livshits, B.: Liquidations: DeFi on a knife-edge (2020), <https://arxiv.org/abs/2009.13235>
 31. Pulse, D.: DeFi status report post-black thursday. Retrieved November 18, 2020 from <https://defipulse.com/blog/defi-status-report-black-thursday> (2020)
 32. Qin, K., Zhou, L., Livshits, B., Gervais, A.: Attacking the DeFi ecosystem with flash loans for fun and profit (2020), <https://arxiv.org/abs/2003.03810>
 33. Qu, M., Huang, X., Chen, X., Wang, Y., Ma, X., Liu, D.: Formal verification of smart contracts from the perspective of concurrency. In: Proc. of the SmartBlock. vol. 11373 LNCS, pp. 32–43. Springer Verlag (2018)
 34. Roscoe, A.W.: The Theory and Practice of Concurrency. Prentice Hall (1997), ISBN 0-13-674409-5
 35. Samreen, N., Alalfi, M.H.: Reentrancy vulnerability identification in Ethereum smart contracts. In: 2020 IEEE International Workshop on Blockchain Oriented Software Engineering (IWBOSE). pp. 22–29 (2020). <https://doi.org/10.1109/IWBOSE50093.2020.9050260>
 36. Sun, J., Liu, Y., Dong, J.S., Pang, J.: Pat: Towards flexible verification under fairness. In: Proc. of the CAV. vol. 5643, pp. 709–714. Springer (2009)
 37. Sun, J., Liu, Y., Dong, J.S.: Model checking csp revisited: Introducing a process analysis toolkit. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation. pp. 307–322. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
 38. Sun, J., Liu, Y., Dong, J.S., Chen, C.: Integrating specification and programs for system modeling and verification. In: Proceedings of the 2009 Third IEEE International Symposium on Theoretical Aspects of Software Engineering. pp. 127–135 (2009). <https://doi.org/10.1109/TASE.2009.32>
 39. Tolmach, P., Li, Y., Lin, S.W., Liu, Y., Li, Z.: A survey of smart contract formal specification and verification (2020), <https://arxiv.org/abs/2008.02712>
 40. Totle: Building with Money Legos. Retrieved November 17, 2020 from <https://medium.com/totle/building-with-money-legos-ab63a58ae764>, accessed: November 17, 2020
 41. Vogelsteller, F., Buterin, V.: ERC-20 token standard. Retrieved July 14, 2020 from <https://eips.ethereum.org/EIPS/eip-20> (2015)
 42. Wang, D., Wu, S., Lin, Z., Wu, L., Yuan, X., Zhou, Y., Wang, H., Ren, K.: Towards understanding flash loan and its applications in DeFi ecosystem (2020), <https://arxiv.org/abs/2010.12252>
 43. Woodcock, J., Larsen, P.G., Bicarregui, J., Fitzgerald, J.: Formal methods: Practice and experience. ACM computing surveys **41**(4), 1–36 (2009). <https://doi.org/10.1145/1592434.1592436>