Mitigating game theoretic vulnerabilities in Chainlink

HANS DEKKER , OĞUZHAN ERSOY , ZEKERIYA ERKIN

Cyber Security Group
Department of Intellegent Systems
Delft University of Technology

Abstract

Oracles are mechanisms that provide blockchain networks with data that only exists outside of the network, such as asset prices. Decentralized Finance (DeFi) protocols use this data, and therefore their usability depends on the reliability of oracles. One such oracle system, widely used by DeFi protocols for pricing feeds, is Chainlink. The Chainlink system mitigates the risk of oracle manipulation attacks that have occurred in various DeFi protocols with a decentralized data aggregation infrastructure. The participants of the Chainlink system are incentivized by a coordination game, which poses game theoretic risks. While some game theoretic analyses of blockchain based systems exist, no formal study has been done on the incentives securing the Chainlink system. In this paper, we present a formal incentive model of the participants in the Chainlink system. We show that users can not detect whether incentives are aligned such that honest node behaviour is a strictly dominant strategy, making it impossible for users to assess the security of the system. We propose a mitigation which enables users to assess the agent incentives of Chainlink nodes such that they can verify whether honest behaviour is a strictly dominant strategy for all participants.

1 Introduction

In the 2008 Bitcoin whitepaper, Satoshi Nakamoto proposed a peer-to-peer electronic payment system based on cryptographic proofs instead of trust in financial institutions [1]. Since Bitcoin has been implemented and launched in 2009, it has carried out over 645 million transactions without trusted intermediaries by leveraging the security properties of a blockchain [2].

The Ethereum network launched in 2015, and enabled decentralized applications (dApps) to be run on its blockchain [3]. This allowed developers to build alternatives for more complex financial services without depending on trusted intermediaries, such as borrowing or trading assets [4]. This new ecosystem of decentralized financial protocols, called Decentralized Finance (DeFi), has four defining properties: it is non-custodial, permissionless, openly auditable, and its services can be easily recombined into new ones [4]. DeFi is rapidly increasing in popularity: In January 2021 the total value locked (TVL) in DeFi protocols was approximately 25 billion USD [4], while as of June 2021, the TVL exceeds 65 billion USD [5].

These DeFi protocols need information from outside of the deterministic blockchain, such as asset prices. For example, Protocols for Loanable Funds (PLFs) need price information to calculate the required collateral of a loan. Such information is provided on-chain by oracles [4]. In order for the protocols to meet the strong security guarantees expected of decentralized applications, oracles need to serve them correct data.

As [6] showed, some DeFi protocols rely on centralized oracles, which are not held accountable for their performance. From a security standpoint this is not desirable, since the operator of the oracle can potentially gain more from dishonest behaviour than from honest behaviour. Furthermore, a single, centralized oracle is easier for adversaries to hack or manipulate than a decentralized oracle [7].

An oracle providing invalid data can have disastrous consequences for DeFi protocols, in one case of an oracle manipulation attack it was possible for adversaries to steal 1.1M USD worth of ETH tokens from bZx users [8]. bZx and various other protocols have integrated Chainlink price feeds to mitigate the risk of such oracle manipulation attacks [9]. Chainlink provides an infrastructure for Decentralized Oracle Networks (DONs) [10].

Fundamentally, Chainlink is a Schelling game based system, in which a participant is rewarded for voting for an answer which is voted for by a majority of participants. Such coordination games are vulnerable to game theoretic risks, since the incentives can become misaligned with the desired behaviour of the participants of a system. The incentives become misaligned when the expected profit of dishonest behaviour outweighs the expected profit of honest behaviour, which can happen through manipulation of the incentives by an external adversary [4], [7].

One way in which the incentives can be manipulated is by bribes offered by an adversary. As shown in [11], such bribes can be offered trustlessly through smart contracts. In one particular type of attack on Schelling game based systems, called $p + \varepsilon$ attacks, an attacker can theoretically bribe such a system without any financial cost [12]. The vulnerability to these type of bribe attacks of multiple differently structured systems has been examined in [13].

Chainlink proposed a unique security model in their whitepaper, as an improvement to their current security model, with a mechanism called super-linear staking [10]. However this model has been criticized for relying too much on trust in participants of the system who have previously shown good behaviour [14]. To the best of our knowledge, the game theoretic risks and possible mitigations of these risks in this security model have not been formally examined yet.

1.1 Contributions

In this paper we aim to answer the following research question: ”How can we mitigate the game theoretic security vulnerabilities in Chainlink?”.

- **Formal model of Chainlink agent incentives.** We present a formal model of the agent incentives in the Chainlink DONs as proposed in their 2021 whitepaper. We then use these models to show that the incentives of Chainlink nodes are not transparent for users relying on the system for data. As a consequence, users can not verify whether honest behaviour is a strictly dominant strategy for Chainlink nodes, and thus not assess the security risks of the system on which they rely.
- **A mitigation decreasing expected profits of dishonest behaviour.** The first mitigation we propose is a mechanism requiring Chainlink Tier 2 nodes to hold LINK tokens, the amount being specified in the service level agreement between the user and DON. We show that this decreases the expected profits nodes can obtain through dishonest behaviour.
- **A mitigation making incentives transparent.** The second mitigation we propose builds on the first mitigation, by enabling a majority of LINK token holders to fork the network and burn the tokens held by malicious Tier 2 nodes. We show how these two mechanisms combined result in improved transparency on the incentives of Chainlink nodes compared to the current Chainlink security model, enabling users to verify whether honest node behaviour is a strictly dominant strategy.

1.2 Paper structure

Section 2 provides an overview of the Chainlink system and related literature. Section 3 explains the methodology by which we compare different security models in the rest of the paper. Section 4 focuses on the problem analysis, showing how game theoretic security depends on incentives which are not transparent to Chainlink users. Section 5 provides the mechanisms we propose to mitigate these game theoretic risks. Section 6 discusses the ethical aspects of this work and Section 7 concludes.

2 Background

2.1 Chainlink

A Chainlink user who needs data on-chain can set up their own DON, or make data requests to an existing DON. The most popular use case of Chainlink DONs at the time of writing is pricing feeds [15].

Data is aggregated from multiple nodes, who are paid a specified [1] amount of LINK tokens as a reward, when their answer deviates less than a specified amount from the final answer. The way this final answer is calculated, after aggregating all responses, depends on the specifications of the service level agreement to which nodes and users requesting the data agreed.

Besides a reward for responding correctly, nodes are incentivized by an on-chain reputation system. The motivation for having a good reputation in the system is the potential future revenue that can be earned by being trusted to respond to future data requests honestly, as nodes with a reputation for being dishonest will likely not be included in DONs by users.

An addition to this security model, as proposed in [10], splits up a DON in two tiers. Tier 1 functions as the system described before, but additionally these nodes are now required to stake a security deposit. This is an extra financial incentive to vote honestly, as their deposit can be lost when one of the Tier 1 nodes alerts Tier 2, on grounds of a manipulated Tier 1 outcome. If Tier 2, which is a group of highly trusted nodes, decides that the majority vote of Tier 1 has been manipulated in some way, the aggregated security deposits from this malicious majority, is then rewarded to the single node who alerted the second tier. Under the assumption that Tier 2 is reliable, this mechanism called 'super-linear staking' increases the required capital for a successful bribe attack quadratically in the number of Tier 1 nodes.

2.2 Related Work

2.2.1 DeFi oracles

A primer on the DeFi ecosystem is presented in [4]. Klages-Mundt et al. [7] presents an analysis of stable coins, one of many categories of DeFi protocols relying on oracles. In their Appendix they include a brief analysis of different oracle designs and their challenges. In [6], the performance of DeFi oracles is measured and compared. Security risks are discussed, and more transparency, accountability, and operational robustness are recommended to mitigate these risks.

For a detailed overview of different oracle designs and an extensive analysis of their trade offs we recommend reading [16]. They examine 17 different oracles, and show that there is a wide variety of different oracle designs. They distinguish oracle designs based on differences in the the following modules: Data Sources, Data Feeders, Selection of Data Feeders, Aggregation, Dispute Phase. For each module they discuss possible attack vectors. In general, most oracles they examine, including Chainlink, try to incentivize their participants financially to provide correct data. One notable exception to this common paradigm are related technologies DECO [17] and Town Crier [18]. DECO allows data providers to prove that the data they provide has not been tampered with and comes from a certain source. This shifts the trust assumption for providing authentic data to the source. Town Crier relies on Trusted Execution Environments (TEE) to prove that an application has been executed without being tampered with, which shifts the trust assumption to the TEE. Vulnerabilities in the TEE remain a possible attack vector for systems like Town Crier [19], [20].

2.2.2 Game theoretic studies

For a game theoretic analysis of blockchain networks we refer the reader to [21]. This work does not cover oracle mechanisms.

[13] analyzes the bribe vulnerability of multiple differently structured Schelling game based systems. It is shown that Schelling game based systems with appeal mechanisms are less vulnerable to $p + \varepsilon$ attacks or other bribe attacks than simple Schelling games without appeal mechanisms. In [10], the Chainlink security model is described, which does feature an appeal mechanism. However, Chainlink's appeal mechanism functions differently than the ones analyzed in [13]. Chainlink appeals feature only one round, which is decided by a group of highly trusted nodes who are incentivized by their reputation. In [13] on the other hand, an appeal mechanism is analyzed in which all token holders can vote in multiple rounds of appeals, until no new appeal round is initiated. Furthermore, in each round, these voters risk losing a security deposit for voting incorrectly, whereas only Tier 1 Chainlink nodes have a security deposit at stake. Chainlink's appeal mechanism allows for faster response time compared

[1] This amount is specified in the service level agreement between the user and DON nodes

to the appeal mechanisms described in [13], since as mentioned in [10], a single round appeal mechanism already slows down confirmation time of data reports. Though the security model has been criticized by [14], to the best of our knowledge, the game theoretic security risks of this model have not yet been formally analyzed.

This research gap on Chainlink security is relevant since in the blockchain community, it is generally seen as the most established decentralized oracle solution [22], [23]. Further illustrating its importance in the Ethereum ecosystem, requests to Chainlink oracles were responsible for most Ethereum transactions of all oracles, though it should be noted that this metric is not a perfect measure for usage due to the possibility of requests being fulfilled off-chain or in a more transaction efficient manner [16].

3 Methodology

Because Chainlink provides a flexible architecture that can be used for multiple different types of data, the way data provided by oracle nodes is aggregated differs for each DON. While for binary data it is clear when a given answer deviates enough from the consensus to be considered wrong, for data like asset prices correctness is more ambiguous. The tolerance for imperfect data also depends on the use case, and therefore these parameters are all set by the user creating a DON.

Nevertheless, at the core every DON relies on the same principle of leveraging the Chainlink reputation system and financial rewards or punishments to increase the likelihood of Tier 1 nodes responding honestly to data requests. Though data requested might often not be of binary nature, to be able to reason about the security of these DONs in general, we simplify the choice of a Tier 1 node to either tell the truth or to respond dishonestly.

An underlying assumption in this analysis is that nodes know whether the data they provide is correct, and can not provide incorrect data by accident. Removing this assumption would add extra uncertainty to the payoff matrices, and unnecessarily complicate the strategies for the scope of this paper.

To compare different security mechanisms, we formalize the strategies of participants in the Chainlink system by means of payoff matrices, where agents have a binary choice of being honest or being dishonest.

4 Achieving incentive transparency

In [24], the desired outcome for a Chainlink DON is described as a Trusted Third Party, which you can rely on to carry out instructions honestly as requested. Because the security of Chainlink, like other decentralized oracle solutions, depends on participating oracle nodes providing ’true’ data, while the correctness is not objectively verifiable, it relies on incentives to stimulate desired participant behaviour.

In game theory, a strategy is called strictly dominant if that strategy always yields a player the largest expected profit compared to other strategies, regardless of the strategy chosen by other players. We assume that agents in the Chainlink system are economically rational agents, and therefore act to maximize their own profit. For this reason, if providing correct data strictly dominates providing incorrect data for Tier 1 nodes, we have a strong economic security guarantee that the Chainlink system will behave as intended, and users receive correct data.

In subsection 4.1 we show that given a hypothetical, completely reliable Tier 2, for a Tier 1 node providing correct data strictly dominates providing incorrect data, unless the profit by corruption exceeds a certain lower bound. However, as we also show, when Tier 2 is completely dishonest, dishonest behaviour becomes the dominant strategy for Tier 1 as well. We show that this dynamic represents a security vulnerability since Tier 2 is only incentivized by future revenue, while users relying on Chainlink data have no insight into the size of this future revenue as perceived by Tier 2 nodes.

Due to this lack of transparency on incentives of Tier 2 nodes, users can not verify whether Tier 1 nodes are incentivized to provide correct data. This is why in this work, we aim to find a mitigation to the Chainlink security model that enables users to quantify the game theoretic incentives of Tier 2

Chainlink nodes, such that users can verify whether honest behaviour is a strictly dominant strategy for Tier 1 nodes.

4.1 Showing dependence on Tier 2

A payoff matrix with expected profits of a Tier 1 node N within the Chainlink system is given in Table 2, while the variables are defined in Table 1.

Table 1: Symbols and notations

Variable	Definition
r	reward for voting according to consensus
d	security deposit, which is lost if Tier 2 decides a node has voted dishonestly
P	probability of Tier 2 reversing the outcome of the vote, punishing the majority and rewarding the whistleblower, if alerted by a Tier 1 node
n	size of Tier 1 majority (number of nodes)
PoC	Profit of Corruption: the expected profits a node gains through dishonest voting. This can be either a bribe or a share of profit by collusion and successfully corrupting data.

Table 2: Expected profit of N given Tier 1 consensus

		Tier 1 consensus	
		Honest	*Dishonest*
N	*Honest*	$r-(P\cdot d)$	$-d+(P\cdot d\cdot n)$
	Dishonest	$-d+(P\cdot d\cdot n)$ $+PoC$	$r-(P\cdot d)$ $+PoC$

Let us assume a completely reliable Tier 2. Then, $P=1$ if Tier 1 consensus is dishonest, and $P=0$ if Tier 1 consensus is honest. This is how the system is assumed to work in the Chainlink whitepaper, ensuring that dishonest behaviour is only the strictly dominant strategy when PoC is larger than $d\cdot(n-1)$. However, if we assume instead that Tier 2 is completely dishonest, meaning $P=0$ if Tier 1 consensus is dishonest, and $P=1$ if Tier 1 consensus is honest, r needs to be greater than $d\cdot(n-1)+PoC$ for honest behaviour to be a dominant strategy. This shows that the desired outcome of honest Tier 1 behaviour being a strictly dominant strategy, depends completely on a reliable Tier 2.

4.2 Current Tier 2 incentives

As we have shown, what strategy is strictly dominant for Tier 1 nodes depends on an honest Tier 2. In Table 4 we show the incentives of Tier 2 node $N2$ as proposed in [10], while the variables are defined in Table 3.

Table 3: Symbols and notations

Variable	Definition
F	Expected future revenue
ΔF	Change in expected future revenue. Assumed to be greater than 0 on grounds of dishonest majority losing credibility and larger total revenue going to honest nodes
M	Expected profit gained by exploiting corrupted data
B	Expected profit gained by bribes
S	Expected profit gained through short selling
n	size of Tier 2 majority (number of nodes)
st	Value of LINK tokens held

Table 4: Expected profit of N2 given Tier 2 consensus

		Tier 2 consensus	
		Honest	*Dishonest*
N2	*Honest*	F	$F+\Delta F$
	Dishonest	0	$M+B+S$

Since the value of F is not known by users, and depends on multiple ambiguous factors such as the desire of nodes to keep participating in the Chainlink system in the future, users can not verify whether or not $F + \Delta F > M + B + S$ and thus whether Tier 2 (and consequently Tier 1) is well incentivized to act honestly. It is not clear either how to estimate the value of ΔF. This value may be positive on grounds of a share of future Chainlink data requests going from dishonest to honest nodes, but the value may also be negative, if this dishonest consensus of Tier 2 would lead to a decrease in future Chainlink usage.

5 Improving Tier 2 Incentives

5.1 Requiring Tier 2 to hold LINK tokens

The first mitigation we propose is a mechanism that requires Tier 2 nodes to lock a certain amount of LINK tokens in a smart contract for a certain amount of time. The amount and time being specified by a user setting up the DON. Under the assumption that the market value of LINK tokens will substantially decrease when the market finds out the system has not functioned correctly, the payoff matrix of a Tier 2 node $N2$ is given in Table 5. Compared to the current security model, $N2$'s expected profits for voting dishonestly are decreased by st. The variables are defined in Table 3.

Table 5: Expected profit of N2 given Tier 2 consensus

		Tier 2 consensus	
		Honest	*Dishonest*
N2	*Honest*	F	$F + \Delta F$
	Dishonest	0	$M + B + S - st$

5.2 Short selling by malicious nodes

In this subsection, we consider a method by which malicious Tier 2 nodes could work around our proposed mitigation by short selling LINK tokens.

The proposed mitigation depends on the market value of locked LINK owned by Tier 2 declining after dishonest behaviour by a majority of Tier 2. A possible workaround to this mitigation for a dishonest majority would be hedging against this exposure to LINK by short selling through PLFs, or DeFi derivative markets.

Centralized exchanges are subject to financial regulation and in the past have shown willingness to freeze funds associated with hacks or other malicious activity [25]. Therefore, the feasibility and scalability of short selling depends on available liquidity on DeFi protocols, which is not in the scope of this work. However, in order to compare security models under the strongest possible threat model, we assume that an adversary is able to obtain $S > st$.

Note that the value of S is neither positively or negatively impacted by the proposed mitigation. Therefore the mitigation decreases the total expected profit of a dishonest strategy by a Tier 2 node. Nevertheless, given sufficient available shorting liquidity, S can be equal to or larger than st, such that users still can not verify whether honest behaviour is a strictly dominant strategy.

5.3 Achieving $S = 0$

In this section we propose a mitigation that requires Tier 2 nodes to have financial assets at stake, without being able to hedge against the loss of these assets in case of an attack on the system. This is achieved with a mechanism inspired by the Ethereum Proof of Stake security model [26].

The idea is that after a manipulation by a dishonest Tier 2 majority, honest Chainlink participants initiate a hard fork of the network. Such hard forks are used if a substantial part of a community participating in a blockchain network, no longer accept a previous state of the network [27]. For the Chainlink system it means, that stakeholders can create a new copy of the network in which the malicious participants lose their tokens st. The assumption is that the market recognizes that an attack has been attempted, and recognizes the new fork as the legitimate network. As a consequence, the tokens in the original network would become worthless, and since the tokens of adversaries have been removed in the new network, effectively, they lost st.

Because of the threat of such a fork, the expected value of S is approximately 0, under the assumption that LINK price will not decrease as the market perceives such a fork as a successful counterattack. Additionally, the supply of the token is reduced, which may stimulate LINK price to go up. Hence, malicious Tier 2 nodes can not expect to profit from shorting after corrupting the system.

This mitigation makes Tier 2 incentives transparent. Users can verify whether or not Tier 2 nodes of a DON have an amount of LINK tokens at stake, larger than $M + B$, and thus whether honest behaviour is a strictly dominant strategy. Here we assume that at least the order of magnitude of M is known to the user, since they know what they will use the data for, and no economically rational adversary will pay a bribe B exceeding M.

As can be seen in Table 6, honest voting is a strictly dominant strategy if $M + B <= st$. The variables are defined in Table 3.

Table 6: Expected profit of N2 given Tier 2 consensus

		Tier 2 consensus	
		Honest	*Dishonest*
N2	*Honest*	F	$F + \Delta F$
	Dishonest	0	$M + B - st$

6 Responsible Research

Because DeFi and oracles are a relatively new area of research, relevant literature is scarce and non peer-reviewed sources such as blog articles were used as well. Although an honest attempt was made to identify the most relevant security risk in the system, this work does not present an exhaustive study of vulnerabilities. Therefore we emphasize that this work should not be interpreted as evidence for absence of any other possible vulnerabilities. We tried to make assumptions, by which different models were compared, explicit and free of bias. We encourage the scientific community to critically examine the assumptions by which models were compared, as well as underlying assumptions we may have forgotten to mention.

Since this work mainly covers vulnerabilities in a version of the Chainlink system that is not yet in use, and the vulnerabilities described are not exploitable by most people, publishing this work does not introduce substantial risk to the system. We emphasize that our goal with this research is to contribute to the improvement of the security of the DeFi ecosystem, and not to help any adversaries attacking protocols.

7 Conclusion and Discussion

Decentralized Finance (DeFi) protocols rely on oracles to provide data from outside of the blockchain network they are run on. In this work, we have shown how the security of Chainlink, one of the most established and widely used oracles, ultimately depends on the reliability of a group of nodes called Tier 2. The only incentive Tier 2 has to behave as intended is based on future revenue and reputation.

Since it is unfeasible to measure the perceived value of this future revenue from the perspective of Tier 2 nodes for users relying on Chainlink oracles, the security is ultimately not quantifiable. This represents a game theoretic security vulnerability, as the profits of corruption could outweigh the cost of corruption, and then rational participants of the system might not behave as intended.

The research question we aimed to answer in this work is: "How can we mitigate the game theoretic security vulnerabilities in Chainlink?". We have proposed two additions to the Chainlink security model. The first mitigation decreases the expected profit of Tier 2 misbehaviour by requiring Tier 2 nodes to hold a certain amount of LINK and thus forcing some financial exposure to the success of the Chainlink system as perceived by the market. Because malicious Tier 2 nodes could decrease the effectiveness of this mitigation by short selling LINK tokens prior to an attack, this mitigation alone does not result in honest node behaviour being a strictly dominant strategy over dishonest behaviour.

We therefore have proposed a second, stronger mitigation inspired by the Ethereum Proof of Stake security model. By adding to the first mitigation, the threat of the stake of malicious Tier 2 nodes being burned through a hard fork of the network, these nodes could be effectively punished for misbehaviour without the possibility of making a profit through short selling. We have shown that this second mitigation allows oracle users to measure the cost of corruption by Tier 2 nodes, and thus enables users to verify whether or not honest behaviour is game theoretically a strictly dominant strategy.

It should be mentioned that in our comparison of security models, only economical incentives were considered, and other external factors, such as legal consequences malicious operators of oracles might face, have not been considered. We focused on economical incentives because relying on strong identities is not in line with the DeFi design goal of permissionless protocols. In practice however, users might be willing to sacrifice some decentralization in favor of ease of use.

As mentioned in Section 3, when comparing node strategies, we assumed that a node has a choice between being honest and dishonest, and thus can only provide incorrect data on purpose. How uncertainty will impact node strategies remains an open question for future work. Another interesting question that remains is the scalability of Chainlink's security model. It is clear that when more value is exchanged through DeFi protocols, more profit can potentially be gained by corrupting the used oracles such as Chainlink. The mechanisms we proposed allow users to verify whether the value put at stake by nodes is in proportion to the potential profit by corruption. However, when the potential profit of corruption amounts to billions of dollars, it is unclear whether it is still feasible for Chainlink participants to put up such large amounts of LINK tokens as a security deposit. Such possibly prohibitive costs of participating in the network might harm decentralization. Finally, even though the second mitigation serves as a method of last resort, and hard forks are not expected to happen often, the possibility of such a fork might require changes in the functionality of DeFi protocols relying on Chainlink. The feasibility of these changes is outside of the scope of this paper and remains an open question.

References

[1] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” p. 9, Oct. 31, 2008.

[2] “N-transactions-total,” Blockchain.com, [Online]. Available: https://www.blockchain.com/charts/n-transactions-total (visited on 06/13/2021).

[3] S. Tual. (Jul. 30, 2015). “Ethereum launches,” [Online]. Available: https://blog.ethereum.org/2015/07/30/ethereum-launches/ (visited on 06/13/2021).

[4] S. M. Werner, D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. J. Knottenbelt, “SoK: Decentralized finance (DeFi),” *arXiv:2101.08778 [cs, econ, q-fin]*, Mar. 2, 2021. arXiv: 2101.08778. [Online]. Available: http://arxiv.org/abs/2101.08778 (visited on 04/25/2021).

[5] “DeFi - the decentralized finance leaderboard at DeFi pulse,” [Online]. Available: https://defipulse.com (visited on 04/25/2021).

[6] B. Liu, P. Szalachowski, and J. Zhou, “A first look into DeFi oracles,” *arXiv:2005.04377 [cs]*, Dec. 11, 2020. arXiv: 2005.04377. [Online]. Available: http://arxiv.org/abs/2005.04377 (visited on 06/07/2021).

[7] A. Klages-Mundt, D. Harz, L. Gudgeon, J.-Y. Liu, and A. Minca, “Stablecoins 2.0: Economic foundations and risk-based models,” *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 59–79, Oct. 21, 2020. DOI: 10.1145/3419614.3423261. arXiv: 2006.12388. [Online]. Available: http://arxiv.org/abs/2006.12388 (visited on 05/26/2021).

[8] K. Qin, L. Zhou, B. Livshits, and A. Gervais, “Attacking the DeFi ecosystem with flash loans for fun and profit,” *arXiv:2003.03810 [cs]*, Mar. 20, 2021. arXiv: 2003.03810. [Online]. Available: http://arxiv.org/abs/2003.03810 (visited on 04/25/2021).

[9] (Mar. 9, 2020). “bZx integrates with chainlink to prevent future price oracle exploits,” [Online]. Available: https://bzx.network/blog/chainlink-oracles (visited on 04/25/2021).

[10] L. Breidenbach, C. Cachin, A. Coventry, S. Ellis, A. Juels, A. Miller, B. Magauran, S. Nazarov, A. Topliceanu, F. Zhang, B. Chan, F. Koushanfar, D. Moroz, and F. Tramer, “Chainlink 2.0: Next steps in the evolution of decentralized oracle networks,” p. 136, 2021.

[11] P. McCorry, A. Hicks, and S. Meiklejohn. (Feb. 10, 2019). “Smart contracts for bribing miners,” In: Zohar, A, (ed.) Proceedings of International Conference on Financial Cryptography and Data Security, FC 2018. Lecture Notes in Computer Science, vol 10958. (pp. pp. 3-18). Springer: Berlin, Heidelberg. (2019). Conference Name: International Conference on Financial Cryptography and Data Security, FC 2018. Lecture Notes in Computer Science Meeting Name: International Conference on Financial Cryptography and Data Security, FC 2018. Lecture Notes in Computer Science Num Pages: 16 Pages: 3-18 Place: Berlin, Heidelberg Publisher: Springer Volume: 10958, [Online]. Available: https://doi.org/10.1007/978-3-662-58820-8_1 (visited on 06/25/2021).

[12] V. Buterin. (Jan. 28, 2015). “The p + epsilon attack,” [Online]. Available: https://blog.ethereum.org/2015/01/28/p-epsilon-attack/ (visited on 06/07/2021).

[13] W. George and C. Lesaege, “An analysis of p + epsilon attacks on various models of schelling game based systems,” presented at the Cryptoeconomic Systems ’20 Conference, MIT, Mar. 8, 2020.

[14] E. Wall. (May 1, 2021). “What’s wrong with the chainlink 2.0 whitepaper? (for simpletons),” Medium, [Online]. Available: https://ercwl.medium.com/whats-wrong-with-the-chainlink-2-0-whitepaper-for-simpletons-d50f27049464 (visited on 06/11/2021).

[15] LinkPool. “Metrics - chainlink market,” [Online]. Available: https://market.link/ (visited on 06/20/2021).

[16] S. Eskandari, M. Salehi, W. C. Gu, and J. Clark, “SoK: Oracles from the ground truth to market manipulation,” *arXiv:2106.00667 [cs, eess]*, Jun. 1, 2021. arXiv: 2106.00667. [Online]. Available: http://arxiv.org/abs/2106.00667 (visited on 06/14/2021).

[17] F. Zhang, D. Maram, H. Malvai, S. Goldfeder, and A. Juels, “DECO: Liberating web data using decentralized oracles for TLS,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’20, New York, NY, USA: Association for Computing Machinery, Oct. 30, 2020, pp. 1919–1938, ISBN: 978-1-4503-7089-9. DOI: 10.1145/3372297.3417239. [Online]. Available: https://doi.org/10.1145/3372297.3417239 (visited on 06/24/2021).

[18] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi, “Town crier: An authenticated data feed for smart contracts,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’16, New York, NY, USA: Association for Computing Machinery, Oct. 24, 2016, pp. 270–282, ISBN: 978-1-4503-4139-4. DOI: 10.1145/2976749.2978326. [Online]. Available: https://doi.org/10.1145/2976749.2978326 (visited on 06/24/2021).

[19] F. Brasser, U. Müller, A. Dmitrienko, K. Kostiainen, S. Capkun, and A.-R. Sadeghi, “Software grand exposure: SGX cache attacks are practical,” in *Proceedings of the 11th USENIX Conference on Offensive Technologies*, ser. WOOT’17, USA: USENIX Association, Aug. 14, 2017, p. 11. (visited on 06/24/2021).

[20] J. V. Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx, “Foreshadow: Extracting the keys to the intel {SGX} kingdom with transient out-of-order execution,” presented at the USENIX Security Symposium (USENIX Security 18), 2018, pp. 991–1008, ISBN: 978-1-939133-04-5. [Online]. Available: https://www.usenix.org/conference/usenixsecurity18/presentation/bulck (visited on 06/24/2021).

[21] S. Azouvi and A. Hicks, “SoK: Tools for game theoretic models of security for cryptocurrencies,” *arXiv:1905.08595 [cs]*, Feb. 21, 2020. arXiv: 1905.08595. [Online]. Available: http://arxiv.org/abs/1905.08595 (visited on 06/13/2021).

[22] J. Finneseth. (Feb. 14, 2021). “Oracle-focused tokens rally as DeFi searches for trusted data providers,” Cointelegraph, [Online]. Available: https://cointelegraph.com/news/oracle-focused-tokens-rally-as-defi-searches-for-trusted-data-providers (visited on 06/14/2021).

[23] (Nov. 10, 2020). “Oracles: The all-seeing eyes that guide crypto networks | CoinMarketCap,” [Online]. Available: https://coinmarketcap.com/alexandria/article/oracles-the-all-seeing-eye-that-guides-crypto-networks (visited on 06/14/2021).

[24] S. Ellis, A. Juels, and S. Nazarov, “ChainLink: A decentralized oracle network,” Sep. 4, 2017.

[25] H. Partz. (May 13, 2020). “Binance freezes funds stolen from upbit in late 2019,” Cointelegraph, [Online]. Available: https://cointelegraph.com/news/binance-freezes-funds-stolen-from-upbit-in-late-2019 (visited on 06/25/2021).

[26] (2020). “Proof-of-stake-faqs,” Ethereum Wiki, [Online]. Available: https://eth.wiki/en/concepts/proof-of-stake-faqs (visited on 06/27/2021).

[27] J. Frankenfield. (Jun. 24, 2021). “Hard fork (blockchain) definition,” Investopedia, [Online]. Available: https://www.investopedia.com/terms/h/hard-fork.asp (visited on 06/27/2021).