An End to End Analysis of Crypto Scams on Ethereum

JADYN KIMBER, Department of Computer Science, University of Saskatchewan, Canada
ENRICO BRANCA, Department of Computer Science, University of Saskatchewan, Canada
ANDREI NATADZE, Department of Computer Science, University of Saskatchewan, Canada
NATALIA STAKHANOVA, Department of Computer Science, University of Saskatchewan, Canada

The increasing number of Ethereum scams is causing significant concern within the blockchain community, costing users millions of dollars annually. Yet, our understanding of how these scams operate remains limited. In this study, we present the first end-to-end analysis of crypto scams using a large set of malicious Ethereum accounts as a case study. We examine the tactics these scams employ on social media platforms to deceive users and convince them to transfer funds to malicious accounts. Our analysis explores the full life cycle of these scams, considering both their distribution through social media and their activity on the Ethereum blockchain. We identify several unique aspects of Ethereum phishing scams that have not been documented in prior literature and find that these scams generally persist significantly longer and result in greater financial losses compared to traditional phishing scams studied in earlier research.

CCS Concepts: • **Security and privacy → Distributed systems security**; *Domain-specific security and privacy architectures*; Database and storage security; Network security.

Additional Key Words and Phrases: Blockchain, Phishing

ACM Reference Format:
Jadyn Kimber, Enrico Branca, Andrei Natadze, and Natalia Stakhanova. YYYY. An End to End Analysis of Crypto Scams on Ethereum. *ACM Trans. Internet Technol.* vv, nn, Article aaa (MMM YYYY), 28 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Ever since the emergence of blockchain technology in the digital ecosystem, it has gained immense popularity and has become predominantly associated with financial transactions. Unfortunately, these characteristics,

Authors' Contact Information: Jadyn Kimber, Department of Computer Science, University of Saskatchewan, Saskatoon, Canada; Enrico Branca, Department of Computer Science, University of Saskatchewan, Saskatoon, Canada; Andrei Natadze, Department of Computer Science, University of Saskatchewan, Saskatoon, Canada; Natalia Stakhanova, natalia@cs.usask.com, Department of Computer Science, University of Saskatchewan, Saskatoon, Canada.

Manuscript submitted to ACM

among others, have made blockchain platforms highly appealing targets for malicious actors worldwide, resulting in an unprecedented surge in blockchain scams. In July, 2022, users of Uniswap Labs fell victim to a significant crypto phishing attack, resulting in the theft of over $8 million [29]. The attacker distributed malicious tokens to platform users and directed them to a fraudulent interface where they were encouraged to exchange the tokens for the platform's currency. This incident serves as just one example among many others.

Currently, there are over 1,000 blockchain platforms available worldwide, many of which facilitate the decentralized creation and automatic execution of smart contracts. Ethereum was the first platform to support smart contracts, followed by others such as Tezos. A smart contract is self-executing code that, once compiled, has its bytecode hosted on a blockchain and can be triggered by transactions to perform specific functions, such as transferring cryptocurrency to another user. Once deployed, these contracts are immutable, meaning they cannot be altered or removed from the blockchain. Although immutability was originally intended to ensure transaction integrity in a trustless environment, it has become a significant obstacle in addressing blockchain attacks. Despite the rise in attacks, blockchain technology lacks mechanisms for automatic attack detection and instead relies heavily on manual analysis of account activity to identify suspicious behavior.

As blockchain attacks continue to increase, this work offers insights into the fundamental characteristics of blockchain scams within Ethereum, the second most popular blockchain after Bitcoin. Among fraudulent activities, phishing scams are especially prevalent in the Ethereum ecosystem [42]. Traditional phishing attacks attempt to trick users into providing sensitive information, e.g., login credentials, personal information or credit card information, by impersonating a trustworthy entity. This often involves setting up sophisticated phishing websites designed to appear highly legitimate to users [26]. In contrast, blockchain-based phishing attacks possess unique characteristics when compared to traditional phishing techniques. They typically target users' financial information and commonly aim to persuade users to execute phishing contracts or transfer money to fraudulent accounts. Some of these characteristics resemble scam activities. We thus venture to understand characteristics of this activity and compare its flow to traditional phishing attack.

In this work, we present our findings from a large-scale effort to characterize scams on the Ethereum blockchain or *cryptoscams* in short. For this analysis, we collected 5,142,020,877 messages from the Telegram, Reddit, and Twitter platforms spanning a period of 1-6 years. Using a curated list of 7,915 known malicious Ethereum accounts, we find 724 messages that advertise 495 of these accounts. We examine the content of these messages to understand the tactics employed by attackers to lure victims to the malicious accounts. To construct a full life cycle of Ethereum scams, we retrieve full transaction history from the Ethereum blockchain for the accounts advertised on social media platforms. Finally, we use these messages along with complete transaction data to calculate losses for victims of these malicious accounts and characterize how accumulated profits are transferred out by the attackers.

There has yet to be a comprehensive study on the operation of modern blockchain-based phishing scams. Previous studies have focused extensively on traditional phishing attacks [26, 30] and spam [17, 20, 22, 31]. More recent research has examined specific types of malicious blockchain activity, such as token theft [15] and pump-and-dump schemes used to manipulate cryptocurrency prices [21]. Phishing detection within the Ethereum ecosystem has also been explored, primarily through the analysis of smart contract code [14, 24, 28, 40, 46] and historical transaction patterns [24, 27, 33, 42]. While these approaches provide promising avenues for scam detection, they are essentially reactive solutions aimed at identifying malicious behavior after it has already occurred. Designing proactive detection and mitigation measures requires a deeper understanding of the lifecycle of Ethereum phishing scams.

Our work presents a comprehensive longitudinal view of cryptoscams, incorporating off-chain activity that is typically hidden from traditional analysis and detection. By providing this broader perspective, we aim to enhance future scam detection methodologies, improving accuracy while reducing the need for extensive prior knowledge of the targeted cryptoscams.

Our extensive analysis of crypto phishing-related scams life cycle give insight into:

- *Lifecycle*: Crypto phishing attacks follow similar patterns to traditional phishing, although over a prolonged period of time with more even distribution of activity. The entire lifecycle of a malicious account can span over 5-7 years. Once a contract is deployed on the blockchain, it may take up to 28 days to see the first victim. However, 80% of the victim activity occurs within the first two weeks of the first transfer of funds for the majority of analyzed accounts (72%).
- *Social media distribution channels*: The first victim transaction is seen within 24 hours of message appearing on social media platform for 78% of malicious accounts. Attackers leverage multiple platforms to advertise the same set of malicious accounts. Yet, the patterns of advertisement is different for each platform.
- *Victims*: Most users become victims of crypto scams only once. However, a notable percentage of victims (15%) have been deceived by scams on multiple occasions, ranging from 2 to 5 separate instances. Furthermore, a small fraction (2%) of all victims have been repeatedly victimized more than 10 times.
- *Financial losses*: We estimate that the total amount of USD $1,978,415,380 (1,180,509 ETH) was transferred into the 495 malicious accounts throughout their entire lifetime. This leads to an average loss of $37K USD (22 ETH) per victim.

To summarize our contributions:

- We present a first end-to-end analysis of blockchain-based phishing related scams life cycle based on the analysis of malicious Ethereum accounts.
- We analyze crypto scam campaigns distributed through three social media sites: Twitter, Telegram, and Reddit, and connect these campaigns to malicious activity on Ethereum.
- We document the entire life cycle of crypto phishing-related scams and compare it with life span of traditional phishing attacks.

- We report our findings and release the dataset of malicious accounts collected in this study to facilitate further research in this area[1].

2 Background

Proposed in 2013, Ethereum is an public blockchain that enables developers to create, deploy, and execute applications, referred to as *smart contracts* in a distributed fashion. Users can interact with smart contracts on the Ethereum blockchain. Smart contracts are written in high-level programming language such as Solidity. To deploy a smart contract on the Ethereum chain, a user compiles a contract to the Ethereum Virtual Machine (EVM) bytecode and then deploys its compiled bytecode by issuing a transaction, i.e., cryptographically signed instruction from user's account. Once published on the blockchain, smart contracts are immutable and remain on the chain forever, although a few mechanisms have recently been implemented to upgrade or remove deployed contracts.

To support pseudo-anonymity, users on the blockchain are identified by unique accounts. There are two types of accounts on Ethereum: *contract accounts*, which are associated with smart contracts deployed on the chain, and *Externally Owned Accounts (EOAs)*, which are controlled by users interacting with the chain, typically through transactions. The addresses of these accounts are essential for facilitating transactions within the Ethereum blockchain. An EOA address is generated from the user's public key, while a contract account address is generated deterministically using the address of the account deploying the contract (usually an EOA) and a unique identifier known as the nonce.

Although smart contracts hosted by contract accounts are intended to be permanent and immutable [1], Solidity includes a 'self-destruct' function that allows a contract to be removed from future states of the chain [2]. This can result in a contract account without an associated contract.

A user represented by an EOA can initiate a transaction that, in turn, can execute a smart contract published on the chain. These transactions are known as *normal transactions*. Contract accounts cannot initiate new transactions on their own but can issue transactions in response to transactions they have received. For example, smart contracts may interact with each other by invoking functions in other contracts. These types of transactions are called *internal transactions*. The Ethereum blockchain runs based on the *Ether* cryptocurrency. Contract accounts can store an Ether balance which they receive and transfer through transactions.

Ethereum is a powerful platform that is leveraged for a variety of services. Among these services is the exchange of tokens that can represent various functionality. To help facilitate the interactions between different tokens, several standards, e.g., ERC20, ERC721, were developed. Any transactions related to the transfer of tokens under these standards can be labelled as *token transactions*.

[1]https://github.com/YNclusk/scamsonethereum

3 Related work

In recent years, blockchain-based cryptocurrency technologies have grown rapidly and attracted significant research interest. Many studies exploring blockchain security and scams have focused heavily on the Bitcoin chain. Notably, Vasek et al. provided the first major overview of Bitcoin-based cryptoscams in 2015 [36], presenting the first large-scale detection of cryptoscams and introducing an initial taxonomy of blockchain-specific scams. Vasek et al. further refined this work in 2019 with the first detailed analysis of previously identified Ponzi schemes on the Bitcoin chain, even including high-level timeline analysis of these scams [37]. Another significant study focusing on Bitcoin was conducted by Atondo et al. in 2022, which analyzed cryptoscam advertisements within Bitcointalk, a forum dedicated to discussing Bitcoin [6].

In recent years, the Ethereum blockchain [1] received significant attention due to its widespread use of smart contracts. The earliest major platform review of Ethereum was published by its creator Vitalik Buterin in 2016 in order to spur prospective research on the topic [10]. This was followed by a 2018 work by Tikhomirov et al. [32] which summarized the state of knowledge of academic studies of Ethereum. In 2020 Chen et al. published a survey on Ethereum security with a focus on distributed applications (dApps) [12]. Later in 2021, Wang et al. provided another survey of research opportunities in Ethereum [41].

These emerging studies for the Ethereum blockchain were primarily aimed at detection of vulnerabilities in smart contracts through static analysis of code in an offline setting, i.e., analysis of code without its execution (e.g., [9, 18, 34, 39]). Mostly based on the patterns of known insecure behavior in code, these studies viewed contracts in isolation, hence, failing to take into account dynamic contract interactions occurring on the chain after the contract's deployment.

More recently, the focus of research on Ethereum has shifted to a detection of malicious activity involving contracts already deployed on the blockchain (e.g., SODA [14], Sereum [28], TXSpector [46], EtherProv [24], ContractGuard [40], and EtherShield [27]), or replaying historic transactions (e.g., HORUS [33], EthScope [42]). As blockchain transactions naturally form a graph, many graph-based approaches were proposed for detection of crypto scams [13, 38, 42, 44], ponzi scams [25]. Some studies focused on services not directly connected to the chain but used extensively such as Uniswap [43]. Our study however takes a broader view, we aim to understand the lifespan of a phishing scam on Ethereum. Phishing attacks traditionally refer to a type of internet based scam that involves impersonating a trusted source and aiming to steal private information from victims connected to that source. This can include private information such as usernames and passwords, personally identifying information, and even financial information [4]. A traditional phishing attack usually commences by the malicious actor sending some sort of message or alert to a victim (in many classic cases, an email) and prompting the victim to willingly give up personal information under the assumption that they are interacting with a trusted source [4]. These phishing attacks generally follow a fairly consistent lifespan, quickly appearing and netting victims and then disappearing equally as quickly [26]. These scams usually have a low success rate, but target many victims, leading to significant losses. Other more targeted phishing scams do occur, but tend to target smaller groups with a more convincing attack. These attacks are commonly known as spear-phishing attacks [26]. Blockchain

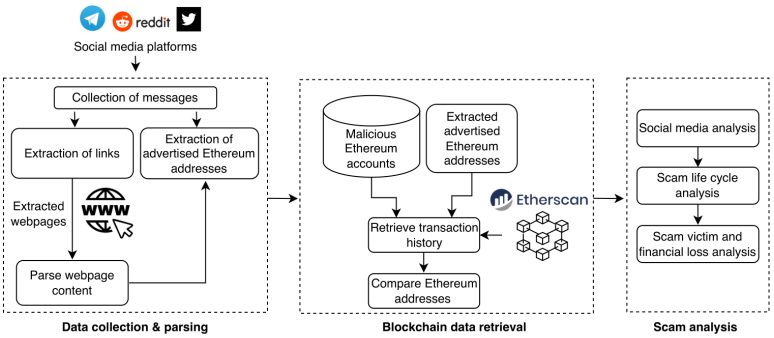

Fig. 1. The overview of the analysis pipeline.

environments often deal with financial assets in an anonymous fashion. This makes them ideal targets for phishing scams [42]. The unique architecture of the blockchain creates a distinctly different environment for the lifespan of phishing attacks, significantly altering the anatomy of such attacks in this context [42].

This traditional understanding of phishing attacks and scam-related activities are the product of over a decade of analysis, characterization and measurement [22, 30]. However, as we show technical niches provide exemptions to this pattern. While large scale characterizations of phishing attacks, e.g., Oest et al [26], have led to a much more in-depth understanding of the end-to-end behavior of traditional phishing attacks, there is very little effort in understanding the anatomy of Ethereum cryptoscams. Bartolletti [8] gave a high-level classification of cryptocurrency scams. Recent studies looked at different elements of malicious activity on blockchain, mostly exploring specific attacks schemes, e.g., theft of tokens [15], pump-and-dump schemes to manipulate cryptocurrency prices [21], giveaway scams [23], 1-day rug pull token schemes [11] and most recently, Txphishing attacks [19]. These studies primarily aim to understand the tactics of specific attacks, and in some cases showcase detection tools for them. Our work seeks to take a much larger view of the phishing scams, measure their impact on victims and characterize their life cycle starting from social media distribution to transaction activity on the blockchain. Our view is different from an analysis conducted by Li et al. [23] that has a narrow focus on giveaway scam websites only.

4 Analysis methodology

To provide a comprehensive longitudinal view of cryptoscams that incorporates both on- and off-chain activity, we designed a three-phased approach, presented in Figure 1. This approach includes: (1) data collection and parsing, (2) blockchain data retrieval, and (3) scam analysis. Given a set of social media sources, we collect publicly available messages posted by users. Our analysis focuses on social network sites that provide an open platform for users to exchange messages, specifically targeting attacks that use social media as a distribution platform.

The collected messages are then parsed to identify Ethereum account addresses, if present, as well as links to external webpages. For each link, we retrieve the corresponding website content and similarly parse it for Ethereum account addresses.

In the subsequent blockchain data retrieval phase, we retrieve the complete transaction history of the collected Ethereum account addresses from the Ethereum blockchain. Since we collect all publicly available social media messages, it is essential to filter for those that contain malicious Ethereum addresses. In the absence of ground truth, we leverage known malicious Ethereum account sets. We match Ethereum accounts collected from social media with a curated set of known malicious addresses.

In the final step of our analysis, we focus on scams. We analyze the collected social media messages and their distribution patterns. We reconstruct a timeline of malicious Ethereum accounts, starting from the deployment of the malicious smart contract to the appearance of the first messages on social platforms. Finally, we explore financial losses and scam victims.

The following sections describe our data collection and analysis phases in more details.

4.1 Data collection & parsing

4.1.1 Collection of messages. For our collection we considered three social media platforms with varied levels of moderation: Telegram, Reddit, and Twitter. To select relevant Telegram channels, we collected channels from public Telegram metadata provider tgstat.ru and then selected 330 channels related to cryptocurrency topics using keywords 'crypto', 'ethereum', 'ether', and 'eth'. For the Reddit platform, we extracted the list of all subreddits using a Python script[2] and then selected 745 subreddits related to cryptocurrency topics using the same set of keywords. We have gathered all available messages from the start of the channel until the date of data collection for each of these platforms. We have also extracted 37 months of archived Twitter data from archive.org. This data is less focused than the data collected from Reddit and Telegram as Twitter does not provide content feeds in the same capacity as Reddit and Telegram.

4.1.2 Extraction of Ethereum Addresses. All obtained from these platforms messages were parsed to extract Ethereum addresses and if present any website links potentially leading to cryprocurrency scams and Ethereum addresses. An Ethereum address has a unique format. For both contract addresses and EOAs, the address is generated using the Keccak-256 hashing algorithm. For EOAs, the address is derived by hashing the public key, while for contract addresses, it is created by hashing the creator's address and nonce. The resulting hash is 256 bits (32 bytes) long, but only the last 20 bytes (40 hexadecimal characters) are used for the address. The final Ethereum address is represented in hexadecimal format with a '0x' prefix, resulting in 42 characters (including the prefix). This unique format, distinct from other cryptocurrencies, allows us to use regular expressions to extract potential Ethereum account addresses.

[2]https://gist.github.com/andrew/1362791

4.1.3 Extraction of Links and Webpages. The obtained lists of website links were de-duplicated and tested for availability. Although we retained the original affiliation, given a large amount of duplication of links among messages, we were bound to reduce the number of times each website was crawled.

Given a list of domains, our data collection module picked a domain URL from the list, and made a first connection to verify if it was alive. If the domain was operational and the website was available (response code: HTTP 200 OK), we loaded the contents of the website directly. If our request was redirected, which happened in many cases, we then also collected the first and the last hop information. For each link, we downloaded the corresponding page with all associated content. Subsequently, our module extracted possible Ethereum account addresses present in loaded content using regular expressions. This analysis primarily aimed to retrieve the malicious addresses explicitly mentioned on the website, so while our web driver loaded the webpage in its entirety (including Javascript elements), we missed content which required explicit user interaction to display.

4.2 Blockchain data retrieval

4.2.1 Malicious Ethereum Accounts. Once extracted, Ethereum accounts are checked to verify their involvement in malicious activities. One of the primary challenges in analysis of malicious blockchain activity is the lack of ground truth. In our analysis, we leveraged a set of known malicious accounts.

Over the years, a significant amount of research was dedicated to detection of smart contracts' vulnerabilities and accounts exhibiting illicit activity on the chain. We have collected a large collection of Ethereum accounts labeled by the previous studies as malicious. We accumulated 36,600 Ethereum accounts associated with malicious activity, mostly related to phishing, from 14 different sources covering a period from 2017 to 2023. This set includes CryptoScamDB[3], PhishingDB[4], EtherScamDB[5], Ethereum repository of malicious accounts[6], Ethereum phishing accounts [7], Forta labelled malicious smart contracts[8] and phishing contracts[9], Etherscan set of phish-hack addresses[10], datasets used in papers [5, 7, 13, 16, 45, 47]. The majority of these sources contained duplicate accounts. After deduplication, we were left with 7,915 unique accounts.

This malicious set also allowed us to account for possible false positives in our collected set. By selecting known Ethereum addresses, we were able to eliminate any non-Ethereum hashes that coincidentally matched the Ethereum format.

Our initial examination of this set showed several inconsistencies. It appeared that some accounts had multiple labels (e.g., malicious and phish-hack, or heist and phishing). The labeling of accounts between different sources were not always consistent, and labels of some accounts varied throughout the time period. We performed three scans of Etherscan, a popular analytics platform for Ethereum blockchain for

[3]https://github.com/CryptoScamDB/blacklist
[4]https://github.com/brianleect/etherscan-labels
[5]https://github.com/MrLuit/EtherScamDB
[6]https://github.com/MyEtherWallet/ethereum-lists
[7]https://github.com/yuanqi7/Phishing-Detection-on-Ethereum
[8]https://github.com/forta-network/labelled-datasets
[9]https://raw.githubusercontent.com/forta-network/labelled-datasets/main/labels/1/phishing_scams.csv
[10]https://github.com/dawsbot/evm-labels

these accounts over a period of three months. We confirmed that these 7,915 accounts were labeled by Etherscan, as malicious or dangerous. However, we noticed Etherscan does not return consistent results over a period of time. The vast majority of accounts had at least one label 'phish' or 'phish-hack' on at least one scan of Etherscan. We also verified that all these accounts were labeled as 'phishing' or 'phish-hack' by CryptoScamDB and by Forta, a commercial platform for security monitoring of blockchain activity. Since the exact meaning of 'phish-hack' label is not defined by Etherscan, we broadly view this set of malicious accounts as phishing-related scams.

4.2.2 Retrieval of transactional history. For each of the extracted Ethereum account addresses, our module retrieves the corresponding transaction history from Ethereum mainnet through the Etherscan platform. For each transaction, we retrieve timestamp, sender and receiver addresses, amount of Ether transferred, data field, and if present, the Etherscan detection label. If an account is associated with a contract, we also retrieve contract's bytecode, source code (when available), all associated internal, normal and token transactions. Of the 7,915 malicious accounts, 1,349 had no incoming transactions and were therefore considered inactive. While it is possible that these accounts were intended for scams that failed to attract attention, their lack of activity excluded them from our temporal analysis. Similarly, 1,709 of malicious accounts had only 1-2 incoming transactions, we suspect these are similarly associated with the least successful fraudulent schemes. The resulting set contained 4,857 accounts.

4.2.3 Comparison. At this step, we validate the Ethereum addresses collected on social media platforms. To ensure that our longitudinal analysis accurately reflects cryptoscams life cycle, we match the collected addresses against a set of verified Ethereum account addresses known to be actively involved in malicious activities.

4.3 Scam analysis

We analyze both on- and off-chain activity associated with cryptoscams, focusing on three key aspects:

- Social Media as a Distribution Vehicle: We examine how social media platforms are leveraged by attackers to distribute cryptoscams. Our analysis reveals various tactics used to lure victims into engaging with malicious content.
- Lifecycle Analysis: We study the timeline of scam-related messages, particularly in relation to the timing of the first and last transactions by victims associated with malicious Ethereum accounts.
- Financial Impact: We assess the financial losses incurred by victims, providing insights into the economic effects of these cryptoscams.

5 How modern crypto scams work

To facilitate an understanding of crypto scam, we start by presenting a high-level overview of the cryptoscam's life cycle of based on our data.

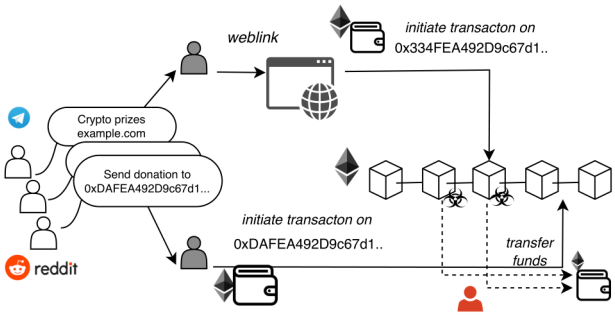

Fig. 2. Modern blockchain-based scam flow

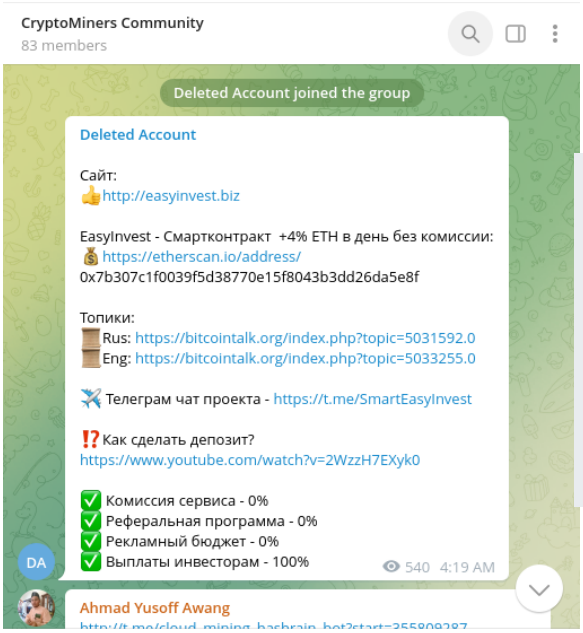

Fig. 3. An example of a Telegram message advertising an EasyInvest account.

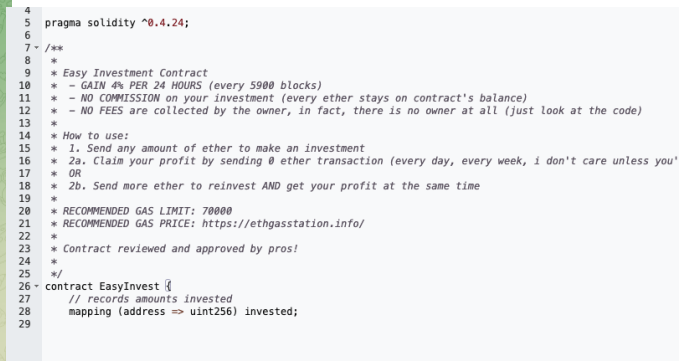

```
pragma solidity ^0.4.24;

/**
 *
 * Easy Investment Contract
 *  - GAIN 4% PER 24 HOURS (every 5900 blocks)
 *  - NO COMMISSION on your investment (every ether stays on contract's balance)
 *  - NO FEES are collected by the owner, in fact, there is no owner at all (just look at the code)
 *
 * How to use:
 *  1. Send any amount of ether to make an investment
 *  2a. Claim your profit by sending 0 ether transaction (every day, every week, i don't care unless you'
 *  OR
 *  2b. Send more ether to reinvest AND get your profit at the same time
 *
 * RECOMMENDED GAS LIMIT: 70000
 * RECOMMENDED GAS PRICE: https://ethgasstation.info/
 *
 * Contract reviewed and approved by pros!
 *
 */
contract EasyInvest {
    // records amounts invested
    mapping (address => uint256) invested;
```

Fig. 4. An EasyInvest smart contract.

The flow of generic traditional phishing attacks has been described by Oest et al. [26]. These attacks follow a relatively consistent life-cycle that spans 22 hours, from their first emergence to the last victim interactions before going offline. The lifespan of a traditional attack begins with what is believed to be testing visits to the phishing URL by the attackers. Within 30 minutes of this initial activity, the first fraudulent messages via an email, or social media site luring victims to the phishing site are sent out. The first victims of a phishing scam often show up within the next 30 minutes, i.e., within one hour of the site's first activity.

By the time the first crawler reaches the site, it is then can take up to six hours for the site to be listed as malicious in the ecosystem. By the time of detection the site has received more than 50% of its visitor activity.

Once detected, activity quickly begins to subside as mitigation efforts start recognizing the site. Mitigation efforts, however, are slow in general and the site receives the last 40% of its activity within the next 12 hours, becoming inactive after this point. This puts the active lifespan of traditional scams at approximately 22 hours. These attacks are built for rapid use and disposal, meaning they are often profitable despite their short lifespan.

While traditional phishing attacks primarily revolve around rapid deployment and a short lifespan, we find that blockchain-based phishing scams, or *cryptoscams* in short, follow a significantly different pattern. Figure 2 illustrates the general flow of cryptoscam based on the data we collected and discovered in our study. The timeline of cryptoscam exhibits similarities to traditional phishing at the onset, with a rapid surge in activity observed within an hour of the first victim transaction as the peak activity is attained much faster than in a traditional phishing, typically within the first hour of the first victim transaction.

The initial drop-off in peak activity is also considerably quicker, with approximately 80% of total victim traffic being reached within the first hour. Subsequently, the scam's activity takes a reverse trajectory, declining at a slower pace over the following few days. However, after the first week, there are consistent, slight upticks in activity. The scam accounts may continue this trend for several years before ultimately ceasing. Due to the immutable nature of the blockchain, scam accounts and contracts cannot be easily removed or universally blacklisted by external authorities, similar to how a phishing website can be taken down. This characteristic makes the chain aspects of the scam more persistent and challenging to mitigate.

The "EasyInvest" contract. The "EasyInvest" contract's behavior is an exemplary case of the cryptoscam life-cycle that is present in our collected data. On September 25th, 2018, an advertisement of the "EasyInvest" contract (`0x7b307c1f0039f5d38770e15f8043b3dd26da5e8f`) was posted on the Telegram group "Cryptominers community". The message promised a daily 4% Ether returns for those who invested funds in the address. The provided link in the message led victims directly to this contract's page on the Etherscan platform directing the user's to trade to the address through instructions included in comments in the contract's source code, as shown in Figure 4. After the Telegram message was posted, the balance for this account spiked sharply during the initial week reaching approximately $1,060,000 USD worth of Ether (632 ETH). This rapid surge in victim activity is similar to traditional phishing, but unlike the conventional phishing, this surge persists for several days. After a period of 20-23 days, the entire amount was withdrawn. The sender of this message, and the website linked to this contract by Etherscan both no longer exist. Yet, the contract remains on the chain. Hence, some activity continued to occur for years after the initial spike of activity (the latest transfer of 0.02 Ether has occurred on April 26, 2020) before ceasing completely.

META coin imitation. Another example highlighting a crypto scam of a different nature but a similar life-cycle is the imitation of the 2021 META coin. This scam first appeared on Twitter on October 29, 2021, through

Table 1. Collected data

Source	Covered time period	Total collected messages	Total unique messages	Total messages with Ethereum accounts	Total messages with malicious accounts	Unique users posting malicious accounts	Ave. msg per user
Telegram	July 2016 - March 2022	27,145,027	27,145,027 (100%)	52,530	247	165	1.5
Twitter	Nov. 2019 - Nov. 2022	5,109,351,116	5,108,163,599 (99%)	1,187,566	322	303	1
Reddit	Jan. 01, 2018 - Nov. 24, 2022	5,524,734	3,988,036 (72%)	309,734	155	155	1
Total		5,142,020,877	5,139,296,662 (99%)	1,549,830	724	623	

a now-deleted account. The account advertised the popular META coin, a legitimate digital asset. However, the posts directed users to an Ethereum address (`0x17a459bfF9277E945354fc32b2DaEf5211fE801B`) via the Uniswap service. This contract represented a fake ERC20 META coin. After the initial Twitter post, token transactions to the address spiked over the next three days before tapering off. Nonetheless, a steady stream of transactions persisted for 32 days from the launch of the account. The total loss from this scam is estimated to be $354,000 USD [35]. While the social media account distributing the scam advertisement has been taken down, the fraudulent contract remains on the blockchain, allowing for small amounts of activity to continue as recently as January 2024, over two years after the scam's original deployment.

6 Dataset analysis

To understand the full life cycle of cryptoscam activity, we collected 5,142,020,877 messages from Telegram, Reddit, and Twitter platforms. They are summarized in Table 1. Among them, over 27 million messages were obtained from Telegram, over 5 million from Reddit, and over 5.1 billion from the Twitter platform. Interestingly, approximately 99% of them were unique. The largest amount of duplication was 28% in our Reddit set (by far our smallest set).

Across these messages, we retrieved 1,549,830 messages with Ethereum accounts. Out of these messages, 724 contained accounts found in our malicious set. The rest of the messages contained accounts that were not labeled as malicious by Etherscan nor appeared in our blacklist. The sheer volume of messages prevented us from asserting that all messages that were not detected as malicious were indeed benign. This aspect of the dataset lies outside of the scope of this work and we leave its analysis to future research.

The majority of the 724 messages contained more than one Ethereum account. These 724 messages were posted by 623 users across three social media sites and overall advertised 495 malicious accounts. Interestingly, of these 495 accounts, 242 were smart contract accounts, i.e., the phishing scams found in our set executed a deployed contract, and in all other cases, victims transferred funds directly to another account. We manually verified each message and the mentioned accounts to ensure no false positives were present.

6.1 Distribution channels

Without the use of social media sites, blockchain-based scams cannot be effectively distributed to victims. We explore the distributions of detected malicious Ethereum accounts across three social sites.

Table 2. Malicious Ethereum accounts found in advertised messages

	Unique malicious accounts	Unique shared across sites
Telegram	262 (53%)	123 (47%)
Twitter	97 (20%)	23 (24%)
Reddit	300 (61%)	127 (42%)
Total	495	452 (91%)

Table 3. Crypto scam advertisement approaches

Donation	2.46%
Impersonation	5.51%
Investment	4.35%
Promotion	6.38%
Ad	2.61%
Info & help request	21.59%
Warnings about scams	21.45%
Not in English	8.84%
No additional info	26.81%

Table 4. The redirected domains that contain malicious Ethereum accounts

Distinct domains	Distribution of domains hosting malicious accounts
etherscan.io	57.7%
youtube.com	34.7%
bscscan.com	2.3%
press.swarm.city	1.1%
bitcointalk.org	0.8%
others	3.4%

Among three, Twitter has the largest amount of messages (322) advertising 97 Ethereum phishing accounts present in our set. On average, this translates to 1 message per attacker (Table 1). Analyzing Reddit posts, we observe on average one to one mapping between users and their corresponding messages with malicious Ethereum accounts. Surprisingly, Reddit has the largest amount of advertised phishing accounts among three sites, 61% of malicious accounts are found in Reddit conversations (Table 2). It appears that due to moderation, attackers are forced to use new Reddit identities for posting phishing scams. Telegram user do no appear to face this challenge. Due to lower levels of moderation in the social platform, a small group of users can repeatedly advertise malicious accounts (1.5 messages per user).

Overall, the vast majority of malicious accounts (91%) are present across three social media sites indicating that attackers leverage multiple avenues for distribution (Table 2). Among three, Telegram has the most amount of shared accounts (47%) which we suspect again is mostly due to low level of moderation.

The majority of the collected messages (70%) advertised the malicious accounts directly in the message text, while the remaining 30% had accounts posted with a linked webpage. We crawled all links found within the social media messages. Since some of the links redirected visitors to other sites, our focus was on the landing sites, as they contained the content that could potentially harm a victim. The analysis of landing site showed that only a handful of domains is used for distribution of Ethereum scams. Vast majority of messages with links containing malicious accounts led to Youtube sites (34.7%) or directly to various Ethereum platforms that allow to trade Ether (e.g., Etherscan.io, Ethplorer.io, Zapper.fi, Bscscan.com). Table 4 shows a list of these landing sites' domains.

As opposed to traditional spam and phishing attacks that primarily leverage disposable domains, it appears that cryptoscams rely on the stable domains. We suspect this is due to the expected longevity of the attacks. Since accounts are immutable, even if malicious activity is identified, the accounts persist on

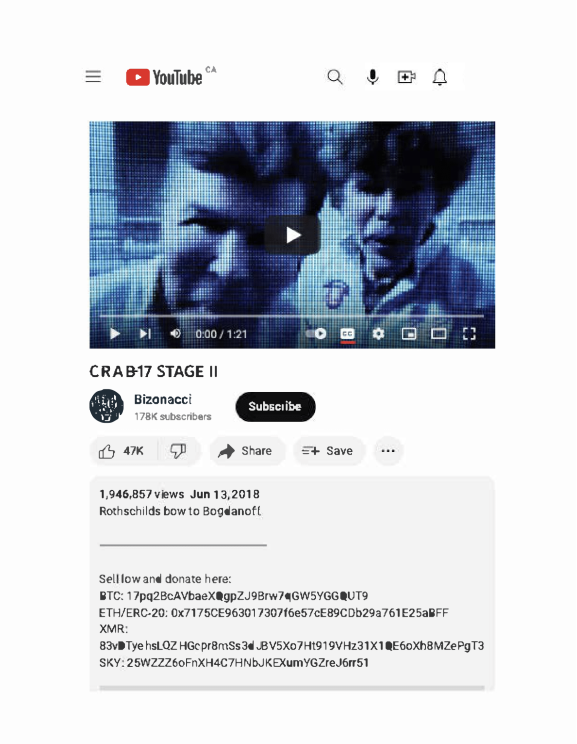

Fig. 5. An example of donation phishing scam on Youtube.

the blockchain and can continue to receive funds transferred by the victims. In contrast to conventional phishing techniques that rely on extensive website impersonation to attract victims [30], we observed that Ethereum phishing scams employ brief and succinct messages that can be easily distributed. We were surprised to see the brevity of messages which is likely due to length constraints posed by distribution channels, e.g., `PRE SALE. Send ETH to-0xe55e749770F1aa743D6ce30a14F84105090398e3.`

We manually inspected landing sites to understand how these advertisement messages are distributed. Interestingly, the vast majority of the landing sites contained messages within the public comments sections. For example, on YouTube sites, the attackers advertise malicious accounts in a description section that appears underneath the video post and/or among comments open to visitors. Similarly, on Etherscan sites, these messages are posted in comments section of the page. Etherscan platform permits registered visitors to leave comments on any Ethereum account page. When it comes to malicious accounts, the site moderators and visitors typically use these comments to alert others about the accounts' malicious activities. Nevertheless, we have also observed instances where this section is utilized to advertise other malicious Ethereum accounts.

6.2 Crypto phishing scam advertisement

Among the phishing scams in our set, we observed several approaches used by attackers to lure the victims:

Donation. Cryptocurrency donations have been particularly appealing to a variety of non-profit organizations and charities in recent years. These donations are anonymous, the process is accessible and simple. Anyone can make a donation to support causes they care about, regardless of their location, from anywhere in the world, without the need for intermediaries e.g., payment processors. Hence, cryptocurrency donation scams have flourished. Cryptocurrency donation posts often request users to sell Ether or tokens at low price or donate them to an address. These donations posts may also promise returns on investment and/or protection of invested funds. For example, `"Sell low and donate here:ETH/ERC-20: 0x7175CE963017307f6e57cE89CDb29a761E25aBFF"`. (Figure 5). This post also advertised a Secure Asset Fund for Users (SAFU) to be setup to protect users' investment, e.g., 'Funds are SAFU'. SAFU is a concept proposed by Binance that refers to protection of funds from irregular trading. It indicates that 10% of all received money will be allocated to offer protection to users' funds.

Investment. Investment scams invite users to invest cryptocurrency promising returns on investment. These investment scams sometimes also referred to as Ponzi schemes. We have observed two variations of investment scams. (1) *Ether investments*, these are direct advertisements inviting potential victims to invest Ether in the contract. For example, a message posted on the Telegram group "Cryptominers community" advertised an EasyInvest account (`0x7b307c1f0039f5d38770e15f8043b3dd26da5e8f`) and promised 4% returns for every 24 hours (or 5900 blocks). (2) *Token investments*. This scheme is designed to lure victims by advertising new tokens. Buying tokens at the initial lower price, promises victims higher returns if they decide to sell them later. Some of these advertisements include referral programs promising bonus tokens in addition to an initial investment. Usually, both schemes convince users to invest indicating that contract's code is open, tested and verified. The contact code is indeed available and verified on Etherscan. However, this verification refers to the ability of developers (in this case, attackers) to prove that the bytecode deployed on the chain comes from the publicly available contract source code. In other words, this verification does not provide any security guarantees.

Promotion scheme presents a variation of an investment scam. It promotes giveaways that are time-bound and promise a specific amount of Ether, e.g., 5 to 100 Ether, or a free token in return for a transfer of 0.5 to 10.0 Ether to the specified address.

Impersonation probably the most closely resembles traditional phishing attacks that attempt to mimic legitimate brands. Cryptophishing similarly aims to impersonate known people or reputable sources, e.g., tokens. However, while traditional phishing relies heavily on visual and structural resemblance of the impersonated website to deceive users, crypto phishing requires comparatively less effort. In cryptophishing, the marketing approach typically revolves around utilizing the reputation of a well-known brand that is being impersonated, coupled with a sense of urgency related to a limited pre-sale period, with the aim of persuading victims to promptly purchase tokens. In this case, it can be an existing token or a promise of a

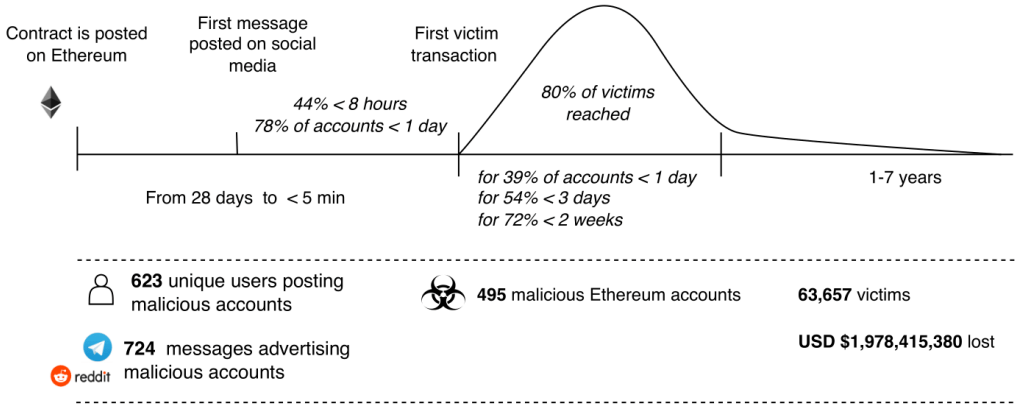

Fig. 6. The crypto phishing scam pipeline on the example of 495 malicious accounts

new token to be released by a known brand such Nike, Meta, and JPMorgan seen in out set. Many of the impersonation scams advertisements include airdrop scams, which impersonate popular Ethereum tokens.

Ad is the advertisement of contracts that does not exhibit any characteristics of other techniques, e.g., does not mention limited sale or percent on investment. For example, `Seriously it brings so much profits. Contract Address: ETH 0x17a459bff9277e945354fc32b2daef5211fe801b`.

Info & help request. We noticed that several messages (21.59%) were not promoting crypto scams, instead, they appeared to be seeking assistance or providing general information that looked innocuous. The assistance requests usually came in a form of asking for help with a malicious crypto wallet, transfer/investment of funds to a malicious account, help to test a deployed malicious contract, etc. We suspect that some of these requests are attempts to take a control over private wallets or trick victims into transferring funds. Since these message typically do not contain user's accounts, a further investigation would require a manual interaction with the users behind these requests.

Warning about scams. 21% of messages shared information about malicious accounts (acknowledging their malicious nature) or contained warnings. Reddit by far was the primary source of these messages.

A fraction of all gathered messages (~35%) were either not written in English or lacked any information besides the malicious Ethereum account. The large portion of all Telegram messages (65%) fell into this group, i.e., contained no additional information beyond the malicious address. Table 3 shows the distribution of messages in our set across categories.

6.3 Life span analysis

The analysis of crypto scam activity across 7915 malicious accounts provides only a partial understanding, as it exposes just one aspect of the malicious behavior that is visible on the blockchain. We selected 495 accounts together with the messages advertising them, in order to gain a complete picture of the entire

lifecycle of crypto scams within those 495 accounts, taking into account both social media activity and behavior on the blockchain.

The general timeline of an Ethereum phishing scam lifespan shown in Figure 6.

We analyze the timeline of messages advertising malicious accounts on three social media platforms. For each malicious account, we analyze timestamps of all messages advertising this account. To align the activity on social platforms and the blockchain, we show the distribution of the messages relative to the first and last victims' transactions for malicious Ethereum accounts. For the accounts that have an associated smart contract, we also indicate the contract's timestamp. Figure 7 shows the timeline of messages advertising malicious accounts on social platforms. There are several observations that lead us to the timeline of a crypto scam lifespan.

Contracts are deployed on average 10 hours before the first victim's transaction, although we observed several cases where contracts are deployed immediately before the first transfer of funds. 242 out of 495 accounts have an associated contract. 3.7% of our malicious contracts include self-destruct functionality that allows developers to remove their contract from the blockchain by calling `selfdestruct()` function. Upon execution, `selfdestruct(address)` transfers all remaining money stored in the contract to an indicated address and removes the bytecode of the contract which triggered it from the state trie of the chain. Since selfdestruct does not remove contract's history already stored on the blockchain, it only serves as a measure to prevent further funds going into this account. This is also potentially a way for attackers to introduce a new variation of the same contract.

Once the contract is activated on the blockchain, the messaging campaign starts, followed by the first transaction from the victim. *The initial transaction is seen within 24 hours for the majority of malicious accounts (78%), and within 8 hours for almost half of our accounts (44%).* Although this is not surprising considering the urgency in some of the posted messages, it is still slower than the rate seen in traditional phishing attacks where 50% of victim activity occurs within 6 hours of their initial visit [26]. This however gives a longer time frame for detection and mitigation efforts.

After the initial transaction for the 495 fully analyzed accounts, *80% of victims are reached within 1 day for 39% of our accounts*. For some accounts it takes longer to accumulate this portion of victims. On average, it takes approximately three days for 54% of the accounts and around two weeks for 72% of malicious accounts to cover 80% of victims. After this, the incoming transactions continue for several years.

The fund transfers out of the 495 malicious accounts are less uniform. On average, we saw *the transfer of funds out of malicious accounts initiated in 4.2 hours after the first incoming transaction.* Several accounts had outgoing transactions occur before the first transfer in. While this behaviour seemed counterintuitive, manual inspection of these accounts showed that these outgoing transactions contained no exchanged Ether. We explore and discuss these transactions next. Overall, the entire lifespan of crypto scams we investigated took up to 5-7 years. This is drastically different from the traditional flow of generic phishing attacks that spans 22 hours, or less than one day from their first emergence to the last victim interactions before going offline [26].

The close analysis of the distributions of messages on each social media platform shown in Figure 7 reveals several additional observations:

First, there is a regular presence of social media messages being posted for weeks after the beginning of the malicious addresses lifespan. The frequency of these postings differs between platforms which we suspect is due to difference in moderation. The postings appear to be sparse and random on Reddit platform, but on Twitter and Telegram, there is a noticeable pattern. Twitter activity is well-coordinated, messages are repeated every several hours over the first 5 days. On the other hand, Telegram activity is more rapid with a rapid succession of messages appearing in the first 1 to 2 days.

Second, in some instances, we observe messages posted after the last activity on the address. The last transaction is a likely indicator that malicious nature of an account is revealed. Although accounts remain on the chain, the community can alert other users, e.g., through social media, comments section on Etherscan platform, hence preventing further transfer of funds. The fact that some messages continue to appear after the last transaction is an indication that attackers are perhaps not aware of an account being detected. Interestingly, in most of these cases the messages are posted on Reddit platform.

Third, the majority of accounts being advertised in social media posts overlap. Even though the content of messages varies across platforms and users, the messages across all three platforms advertise mostly the same (91%) malicious accounts from our 495 account subset. This is a strong indication that attackers use multiple avenues for reaching victims to improve the success rate of the scams.

Validation. To provide further insight into the lifecycle patterns observed among malicious accounts in our dataset, we collected five recent Ethereum accounts labeled as phishing scams by Etherscan. We verified that none of these accounts were included in the original analysis. Additionally, we confirmed that these accounts were advertised on Reddit and Twitter. The most recently active contract in our validation set had been inactive for at least 83 days at the time of collection, indicating that the bulk of victim activity for each scam had already ceased.

The timelines of these 5 accounts are presented in Figure 5. The timelines of all five accounts align with the temporal patterns observed in the originally analyzed scams. Since the original set contained scams collected as early as 2017, these latest results confirm that Ethereum-based phishing scams continue to follow the same temporal patterns.

7 Victims

To better understand the financial impact of crypto scams, we explore users that have fallen victims of phishing scams contained in both sets of malicious accounts, i.e., 495 accounts advertised through collected social media posts and all 7915 collected Ethereum accounts. For each of these accounts, we examine the history of valid incoming transactions to identify sending accounts that correspond to the victims of the phishing scams. Since in our analysis we rely on labelled accounts, it is possible that some of the victim accounts are controlled by adversaries as well. We leave investigation of this potential collaboration for future research.

Table 5. Timelines of the validation accounts

	Contract creation date	Contract creation to first victim transaction (hours)	Duration between first to last victim transaction (days)	Time to reach 80% of victims (days)
0x0000d169f98e078b60bfb09a69d145e72dbe0000	2024-03-17	1.98	225.66	94.66
0x09e0cA7E92E2B893f07cdAE75C011C3606ecD61a	2024-02-13	8.19	37.00	0.16
0x00002D98735603b0447B46c7807C385c29230000	2023-10-28	1.186	385.00	91.50
0x0000A4998724E52F0886edFf693aCA33f9900000	2023-08-16	11.18	494.53	67.53
0x7b732e6eb24dc885db7ff417478dea2c1d4d64b2	2022-06-22	5.04	694.98	172.97
Average	–	5.52	367.43	85.37

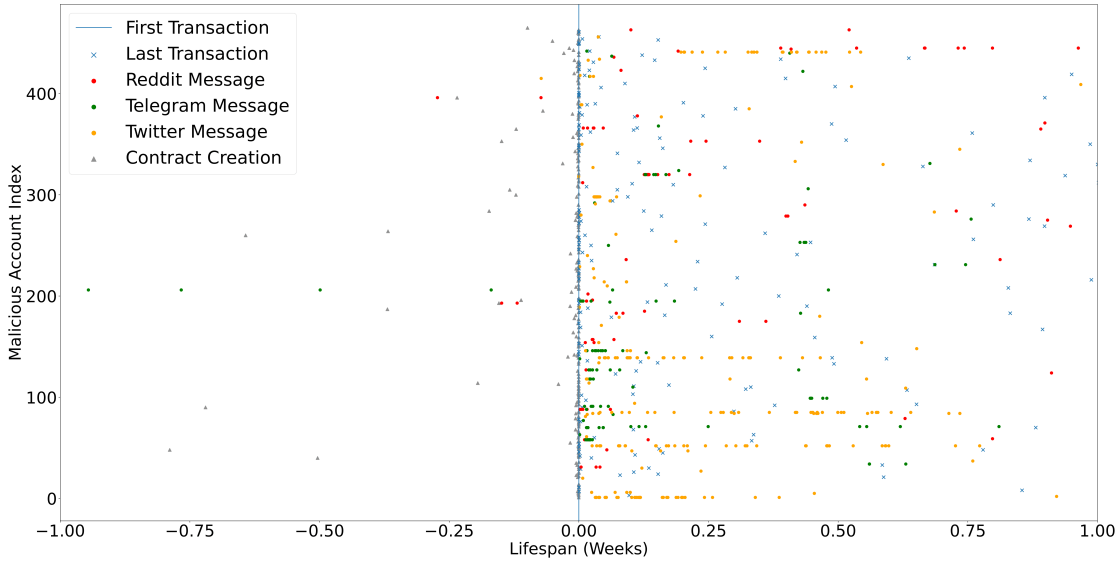

Fig. 7. Relative timeline of 724 messages posted on social media relative to the first transaction on the corresponding malicious account.

We observe that 63,657 victims fell for 495 malicious accounts found in the social media posts. Table 6 shows a breakdown of how many victims fell for how many scams, with 80.98% of victims falling for only a single scam among the 495, this meets our expectations that most users who are tricked are unlikely to be tricked again. This rate changes to 77% when considering the complete malicious set.

What is more interesting is the fact that roughly 15% and 20% of victims respectively transferred money to malicious accounts 2-5 times during their lifetime. The repeat rate is relatively high for victims falling for Ethereum scams more than 10 times, approximately 2%. These numbers are unusually high, however they may be inflated by the 14% possible collaborators mentioned above. Due to chain anonymity, we distinguish blockchain users by their accounts numbers. Since users are encouraged to generate a new account address per transaction, our calculations are likely to underestimate the number of repeat victims, which may

Table 6. Victims of phishing scams

Number of scams victims participated in	Number of victims
1	51,552 (80.98%)
2-5	9,837 (15.45%)
5-10	1,233 (1.94%)
10-20	657 (1.03%)
20-50	282 (0.44%)
50-100	46 (0.07%)
above 100	50 (0.08%)
Total victims	63,657

possibly be much higher. We also suspect these numbers are related to the nature of accounts in our set. Specifically, we observed that some fraudulent accounts are presented as investment opportunities. As a result, individuals who invest early in these accounts may receive some profits before all funds are stolen by the attackers. One possible explanation of this behaviour maybe that a few users may have invested in multiple accounts and have successfully received some initial returns, would be enticed to invest more extensively. It is though also possible that users are forced to make transfers (e.g., in case of ransomware) or are not even aware of them (e.g., in cases when the sites install malware that takes advantage of system flaws and exploits the victims' crypto wallets).

Financial losses. To understand the total amount of lost funds throughout the history of these malicious accounts, we calculate the total number of Ether transferred to malicious accounts by valid incoming transactions. In other words, we estimate how much money were transferred by victims of these scams. Since blockchain users are charged a fee for each transaction, we include the paid transaction fees into the total amount of lost funds. To provide a common reference point, we have converted the corresponding transferred Ether amounts and the paid transaction fees (in wei) to USD[11].

We estimate that the total amount of 1,180,509 ETH (USD $1,978,415,380) was transferred into the 495 malicious accounts throughout their lifetime. These amounts include 624 ETH (USD$1,046,571) in transaction fees paid by victims. This evaluates to an average loss of 22 ETH ($37K) per victim. In addition, this shows that 6.3% of the malicious accounts caused nearly 68% of all losses. While this seems extreme, this follows the traditional distribution of phishing scam profitability observed by Oest et al.[26]. Token transactions have been excluded from our analysis. Due to the drastic volatility of token prices, it would be challenging to accurately estimate them. Therefore, it is likely that the actual losses incurred by victims are considerably higher than our estimations.

By linking each of the 495 selected scam addresses to the social media channels where they appeared, we observed certain patterns related to the financial losses associated with these scams. Table 7 displays the losses from scams based on combinations of social media distribution channels. Interestingly, although

[11]The value of wei fluctuated very rapidly over the years. To estimate the financial losses, we used the price of wei given by Etherscan at the time of our analysis, i.e., 1Eth = USD$1675.9. While this results in a slight overestimation of early losses from 2015-2016 scams, it also underestimates more recent losses.

Table 7. Estimated loss per scam found in each type of social media channel.

Social Media Channels	Number of Scams	Total Loss	Mean Loss per Scam
Telegram exclusive	37	$671,123,540 (USD)	$18,138,474.05 (USD)
Twitter exclusive	85	$34,620,439 (USD)	$407,299.28 (USD)
Reddit exclusive	46	$425,019,598 (USD)	$9,239,556.48 (USD)
Twitter and Reddit	8	$549,718 (USD)	$68,714.50 (USD)
Telegram and Reddit	3	$30,744,033 (USD)	$10,248,011.00 (USD)
Twitter and Telegram	9	$190,691,852 (USD)	$21,187,983.56 (USD)
All sources	307	$625,666,200 (USD)	$2,038,000.65 (USD)
Total	495	$1,978,415,380 (USD)	$3,996,798.75 (USD)

most scams appeared across all three social media channels, these scams incurred lower financial losses on average than those that appeared solely on Telegram or Reddit. This aligns with traditional phishing patterns, where broader, more generic scams tend to be less targeted, less successful, and thus incur smaller losses [26]. Additionally, scams that appeared exclusively on Telegram, rather than across all three platforms, had a higher average loss. This suggests that scams tailored for Telegram are more effective than those on other platforms, likely due to the platform's lower moderation levels and higher popularity among cryptocurrency users.

Although the total victim losses are significant, not all attackers are successful in obtaining these profits. Only a handful of accounts receive profits above 50,000 ETH (approximately 84 million USD). The top accounts that received the highest amount of funds are 0xcb36b1ee0af68dce5578a487ff2da81282512233 that is associated with Rari Capital hack and received USD$333,958,057 over its lifetime and 0x489a8756c18c0b8b24e c2a2b9ff3d4d447f79bec that is labelled as BNB Bridge Hack and consequently received USD$56,598,551 in transactions. However, the median amount of Ether transferred to all these accounts over their lifetime is 15.4 ETH (approximately USD $25,808) with 28% of accounts seeing less than 0.3 ETH (approximately USD $500).

The analysis of victims' behavior over time reveals some interesting patterns (Figure 8). In the initial years (2015-2016), victims tended to transfer relatively large amounts of money, with 72-89% of transactions being in the range of 0.6-6 ETH (approximately USD$1000-9999). However, as time went on, victims started transferring smaller amounts of money more frequently. By 2022-2023, the majority of victims only sent small amounts of funds 0.0006-0.006 ETH (approximately USD$1-9), indicating a change in behaviour and perhaps a greater awareness about crypto scams. This may also be an indication of major blockchain scams shifting their focus to token and NFT based scams. Token transfers skyrocketed in 2018 [3] which coincides with pattern shift towards smaller transactions in Ether. While this could be partially explained by the increase in Ether price driving down transaction values, this increase in small transactions is not seen during any other spike in Ether price. Small value transactions including zero Ether transfers are used to trigger token manipulating functions in contracts. This connection is another indication of an emergence of token-based scams.

Our study also shows losses greater to other blockchain based scams on a per scam basis. Vasek et al. [36] found a loss of around $11 million USD lost among 3900 victims in their survey of bitcoin based scams in 2015. This shows an average of $2,821 USD lost per victim. Our work covered 63,657 victims in our 495 scam subset alone and showed an approximate loss of $37,000 USD per victim.

Zero Value Transactions. One of the most unexpected aspects of our results is the presence of transactions where 0 Ether was traded between parties (referred to as "zero value transactions"). These transactions made up 41% of our total transactions destined to 495 malicious addresses, and 37% of the total transactions for the entire malicious set. Zero value transactions in Ethereum are almost exclusively made to trigger specific contract functions, although there are some exceptions to this rule, e.g., transactions which create contracts. We have parsed the transaction data and specifically data field to extract function call and its arguments. Since the data field is optional, not all transactions carry this information. 51% of incoming transactions had no data field. The rest of transactions called transfer, update, extract functions that appear to be related to the transfer of funds based on their descriptive names.

A closer analysis of these zero value transactions in our set revealed that these transactions often happened after the main activity on the phishing accounts has subsided. For example, the majority of transactions to EasyInvest contract (`0x7b307c1f0039f5d38770e15f8043b3dd26da5e8f`) that occurred after the first 20 days are zero Ether transactions. We manually analyzed some phishing scams and found that they advertise zero value transactions as a way of retrieving the interest on the investment. Hence, in these cases the zero value transaction activity indicates victims' unsuccessful attempts to recover their funds.

8 Discussion

Summary of identified behavior. Our analysis of phishing-related scams in the Ethereum blockchain ecosystem revealed novel patterns indicating that its unique nature has made phishing scams more damaging than those in traditional contexts:

- *Crypto scams on Ethereum have a significantly longer lifespan than traditional scams.* Specifically,
  - Scam contracts are deployed on average 10 hours before the first victim's transaction.
  - The initial victim transaction is seen within 24 hours for the majority of malicious accounts.
  - In most cases, the majority of victims make transfers to phishing accounts within two weeks of its deployment.
  - Typically, the funds are transferred out of malicious accounts in 4.2 hours after the first victim transaction.
  - Overall, the scams may last up to 5-7 years from their emergence to the last victim transaction compared to 22 hour duration of conventional phishing scams.
- Crypto scam victims who fall for more than one phishing scam are unusually susceptible to falling for multiple crypto-scams.

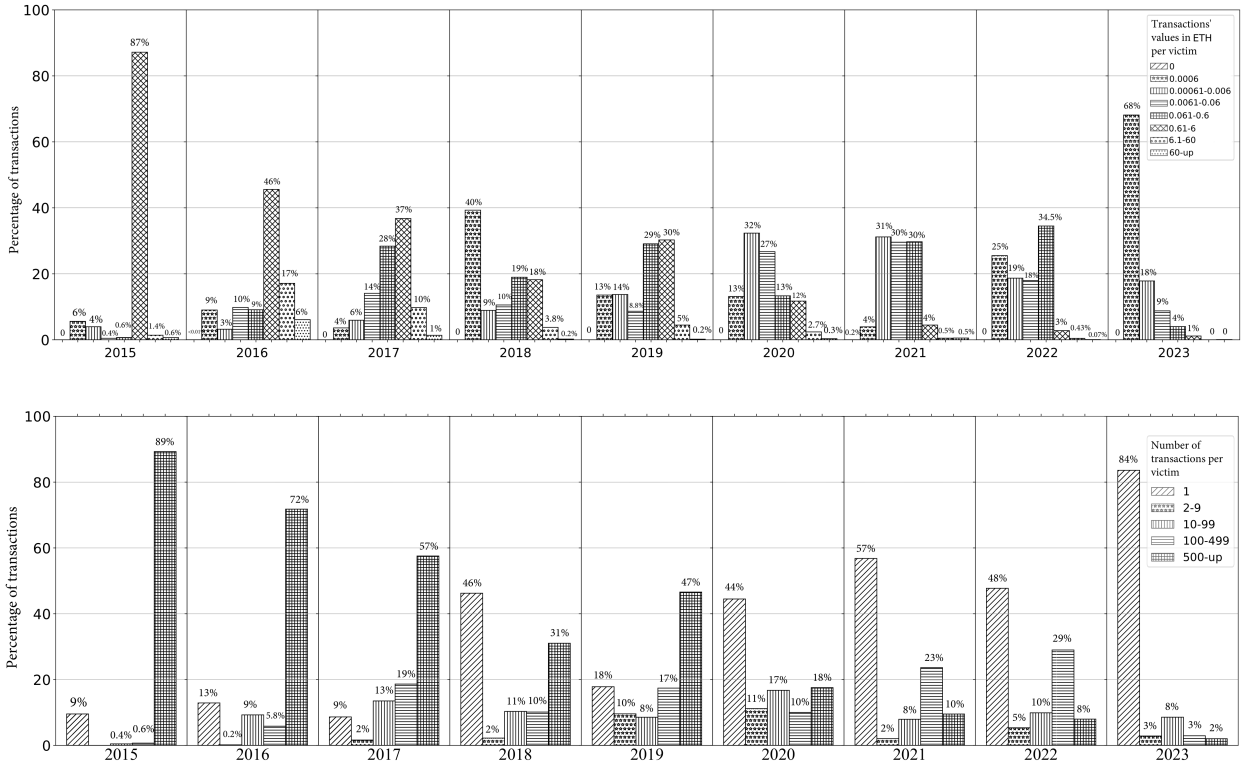

Fig. 8. Frequency of transactions per victim.

- Crypto scams have shifted to larger, more dispersed scams with smaller per victim losses over the past 8 years.
- Crypto scams on Ethereum result in higher average losses compared to those on other platforms and in traditional phishing contexts.

These patterns are critical in guiding the design of new strategies to detect fraudulent activities early and minimize their impact. These insights help refine automated detection systems, improve user awareness campaigns, and enhance security protocols to better protect users from evolving threats.

Mitigation measures. Mitigation of scams in blockchain community is still largely absent. The clearer view of the crypto scam ecosystem and timelines provided by our work provides an ideal platform for the development of effective detection and mitigation strategies. We envision the following strategies:

- *Early-stage detection.* Detecting crypto scams is primarily a manual process, often dependent on feedback from victims. Leveraging social networks as a delivery channel makes early-stage detection methods viable. Our examination reveals that in most cases, the initial transaction by a victim occurs

within a day of a message being posted. This presents an opportunity for proactive measures before the scam campaign on Ethereum blockchain claims any victims. By monitoring social media for scam advertisement messages, potentially malicious Ethereum addresses can be identified and labeled early in the scam lifecycle. This makes detection less reactive and more preemptive.

- *Automated detection of scam messages.* Identifying patterns in scam messages and associated Ethereum addresses on social media enables the development of automated or semi-automated tools to detect and block these scams, reducing the reliance on manual processes and victim feedback.
- *Enhanced context for blockchain analysis.* Knowing when and where scam messages are posted can improve blockchain analysis, allowing for more accurate mapping of transaction timelines and scam strategies, ultimately refining detection models.
- *Contract isolation.* Once deployed on the chain, contracts are immutable. However, Ethereum has recently introduced a few strategies that allow updating or destructing a contact already deployed on the chain. Although the current options require participation of contract developers (in our case adversaries), advancing these strategies to allow rerouting transactions destined for malicious contracts will offer more effective mitigation strategies even when contracts remain on the chain. This strategies promise to be particularly effective for re-occurring victims that are deceived by scam on multiple occasions.
- *Centralized blacklists.* The long and consistent activity patterns we observed in our data as well as the unusually large losses show that the decentralized community based efforts to secure the Ethereum chain against attacks are insufficient. A central collective effort to report and maintain malicious accounts' and contracts' repositories is necessary.

Limitations. Our analysis provides a glimpse into the malicious activity occurring on Ethereum. Despite utilizing various techniques to gather data on social media messages and associated malicious behavior on the blockchain, we could only reconstruct the lifespan of some of the accounts involved. There are possibly other malicious Ethereum accounts that we did not examine. Our estimates are based solely on the transaction activity of the accounts we studied. We believe that the actual losses suffered by victims and gains achieved by attackers are greater than what our analysis suggests.

Our study also shows losses greater to other blockchain based scams on a per scam basis, specifically those on Ethereum. Vasek et al. [36] found a loss of around $11 million USD lost among 3900 victims in their survey of Bitcoin based scams in 2015. This shows an average of $2,821 USD lost per victim. Our work covered 63,657 victims in our 495 scam subset alone and showed an approximate loss of $37,000 USD per victim.

Our set does not include domains that are exclusively registered and established for specific cryptoscam accounts. However, during our analysis of comments, we encountered multiple mentions of websites that offer users additional information about contract's details, company information, and success stories of

early 'investors'. Unfortunately, all these websites had been taken down before we were able to access them. We suspect that these sites were only operational during active scam period and were abandoned afterwards.

In addition, our analysis shows lifespan analysis only of attacks within the Ethereum blockchain. A comparative study to show if scam lifespans differ based on the structure of the chain would be an interesting future work, however it is outside of the scope of this paper.

In this work, we relied on other sources for the detection of malicious accounts. In the absence of ground truth and a reliable single source for labeling, it is possible that some were mislabeled. Given that all accounts we considered were reported by all sources as malicious, the possible fraction of false positives is likely to be small.

Throughout this analysis, we observed that malicious activity is broadly defined. Most sources that we used to derive our set of malicious accounts use coarse labels to identify a variety of activity, e.g., phishing, scam, exploit. Etherscan platform commonly labels the majority of malicious activity as phishing. It is possible that this labeling reflects a dual nature of the malicious activity conducted by these accounts. On the other hand, a broad labeling and treatment in our view demonstrates the need for better understanding of the malicious behavior on the blockchain.

9 Conclusion

The raising numbers of Ethereum phishing scams have become a major concern for the blockchain community. These scams trick users into transferring their funds to fraudulent accounts using a variety of tactics such as fake social media accounts, fake giveaways, and fake brands that mimic legitimate ones to lure users. The Ethereum blockchain immutable and public nature of smart contracts has led to a niche in crypto phishing scams that was not been explored yet. In this work, we present the first large-scale study of phishing-related crypto scams. We collected 5,142,020,877 messages from three well-known social media platforms and retrieved 724 messages advertising 495 malicious Ethereum accounts. We characterized the distribution process and the life cycle of crypto scams on the Ethereum blockchain. Our observations underscored the critical stages of the phishing life cycle in Ethereum. This work provides a foundation for early phishing detection and serves as a platform for future efforts to combat blockchain-based phishing attacks.

A key future research direction is leveraging the temporal insights from our analysis to enhance reactive detection mechanisms, enabling the identification of scams much earlier in their lifecycle.

Additionally, the development of AI-guided detection heuristics based on the observed life cycle patterns could significantly advance traditional anti-phishing strategies in this domain. Finally, we hope the data collected through this work will support and inspire further research in this field.

References

[1] [n. d.]. https://ethereum.org/en/whitepaper/.
[2] [n. d.]. ethereum.org. https://ethereum.org/.

[3] [n. d.]. etherscan.io. https://etherscan.io/.

[4] Neda Abdelhamid, Aladdin Ayesh, and Fadi Thabtah. 2014. Phishing detection based associative classification data mining. *Expert Systems with Applications* 41, 13 (2014), 5948–5959.

[5] Salam Al-E'mari, Mohammed Anbar, Yousef Sanjalawe, and Selvakumar Manickam. 2021. A Labeled Transactions-Based Dataset on the Ethereum Network. In *Advances in Cyber Security*, Mohammed Anbar, Nibras Abdullah, and Selvakumar Manickam (Eds.). Springer Singapore, Singapore, 61–79.

[6] Gilberto Atondo Siu, Alice Hutchings, Marie Vasek, and Tyler Moore. 2022. "Invest in crypto!": An analysis of investment scam advertisements found in Bitcointalk. In *Symposium on Electronic Crime Research*. APEG.

[7] Massimo Bartoletti, Salvatore Carta, Tiziana Cimoli, and Roberto Saia. 2020. Dissecting Ponzi schemes on Ethereum: Identification, analysis, and impact. *Future Generation Computer Systems* 102 (2020), 259–277.

[8] Massimo Bartoletti, Stefano Lande, Andrea Loddo, Livio Pompianu, and Sergio Serusi. 2021. Cryptocurrency Scams: Analysis and Perspectives. *IEEE Access* 9 (2021), 148353–148373.

[9] Lexi Brent, Anton Jurisevic, Michael Kong, Eric Liu, Francois Gauthier, Vincent Gramoli, Ralph Holz, and Bernhard Scholz. 2018. Vandal: A scalable security analysis framework for smart contracts. *arXiv preprint arXiv:1809.03981* (2018).

[10] Vitalik Buterin. 2016. Ethereum: platform review. *Opportunities and Challenges for Private and Consortium Blockchains* 45 (2016).

[11] Federico Cernera, Massimo La Morgia, Alessandro Mei, and Francesco Sassi. 2023. Token Spammers, Rug Pulls, and Sniper Bots: An Analysis of the Ecosystem of Tokens in Ethereum and in the Binance Smart Chain (BNB). In *32nd USENIX Security Symposium*. 3349–3366.

[12] Huashan Chen, Marcus Pendleton, Laurent Njilla, and Shouhuai Xu. 2020. A Survey on Ethereum Systems Security: Vulnerabilities, Attacks, and Defenses. *ACM Comput. Surv.* (2020).

[13] Liang Chen, Jiaying Peng, Yang Liu, Jintang Li, Fenfang Xie, and Zibin Zheng. 2021. Phishing scams detection in Ethereum transaction network. *ACM Transactions on Internet Technology (TOIT)* 21, 1 (2021), 1–16.

[14] Ting Chen et al. 2020. SODA: A generic online detection framework for smart contracts. In *NDSS*.

[15] Zhen Cheng, Xinrui Hou, Runhuai Li, Yajin Zhou, Xiapu Luo, Jinku Li, and Kui Ren. 2019. Towards a First Step to Understand the Cryptocurrency Stealing Attack on Ethereum. In *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*. USENIX Association, Chaoyang District, Beijing, 47–60.

[16] Steven Farrugia, Joshua Ellul, and George Azzopardi. 2020. Detection of illicit accounts over the Ethereum blockchain. *Expert Systems with Applications* 150 (2020), 113318.

[17] Chris Grier, Kurt Thomas, Vern Paxson, and Michael Zhang. 2010. @spam: The Underground on 140 Characters or Less. In *Proceedings of the 17th ACM Conference on CCS* (Chicago, Illinois, USA). ACM, New York, NY, USA, 27–37.

[18] Ákos Hajdu and Dejan Jovanović. 2020. Solc-verify: A Modular Verifier for Solidity Smart Contracts. In *VSTTE*.

[19] Bowen He, Yuan Chen, Zhuo Chen, Xiaohui Hu, Yufeng Hu, Lei Wu, Rui Chang, Haoyu Wang, and Yajin Zhou. 2023. TxPhishScope: Towards Detecting and Understanding Transaction-based Phishing on Ethereum. In *Proceedings of the 2023 ACM SIGSAC CCS*. ACM, New York, NY, USA, 120–134.

[20] Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. 2008. Spamalytics: An Empirical Analysis of Spam Marketing Conversion . In *Proceedings of the 15th ACM Conference on Computer and Communications Security*. Alexandria, Virginia, USA, 3–14.

[21] Massimo La Morgia, Alessandro Mei, Francesco Sassi, and Julinda Stefa. 2023. The Doge of Wall Street: Analysis and Detection of Pump and Dump Cryptocurrency Manipulations. *ACM Trans. Internet Technol.* 23, 1, Article 11 (feb 2023), 28 pages.

[22] Kirill Levchenko, Neha Chachra, Brandon Enright, Mark Felegyhazi, Chris Grier, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Andreas Pitsillidis, Nicholas Weaver, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage. 2011. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of 32nd annual Symposium on Security and Privacy*. IEEE.

[23] Xigao Li, Anurag Yepuri, and Nick Nikiforakis. 2023. Double and Nothing: Understanding and Detecting Cryptocurrency Giveaway Scams. In *NDSS*.

[24] Shlomi Linoy, Suprio Ray, and Natalia Stakhanova. 2021. EtherProv: Provenance-Aware Detection, Analysis, and Mitigation of Ethereum Smart Contract Security Issues. In *2021 IEEE International Conference on Blockchain*. 1–10.

[25] Yincheng Lou, Yanmei Zhang, and Shiping Chen. 2020. Ponzi Contracts Detection Based on Improved Convolutional Neural Network. In *2020 IEEE International Conference on Services Computing (SCC)*. 353–360.

[26] Adam Oest, Penghui Zhang, Brad Wardman, Eric Nunes, Jakub Burgis, Ali Zand, Kurt Thomas, Adam Doupé, and Gail-Joon Ahn. 2020. Sunrise to Sunset: Analyzing the End-to-end Life Cycle and Effectiveness of Phishing Attacks at Scale. In *Proceedings of the 29th USENIX Security Symposium*. 18 pages.

[27] Bofeng Pan, Natalia Stakhanova, and Zhongwen Zhu. 2023. EtherShield: Time Interval Analysis for Detection of Malicious Behavior on Ethereum. *ACM Trans. Internet Technol.* (nov 2023).

[28] M. Rodler, Wenting Li, G. Karame, and L. Davi. 2019. Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks. *NDSS* (2019).

[29] Sujith Somraaj. 2022. Hackers Nab $8M in Ethereum via Uniswap Phishing Attack. https://decrypt.co/104916/hackerUsnab-8m-ethereum-uniswap-phishing-attack.

[30] Karthika Subramani, William Melicher, Oleksii Starov, Phani Vadrevu, and Roberto Perdisci. 2022. PhishInPatterns: Measuring Elicited User Interactions at Scale on Phishing Websites. In *Proceedings of the 22nd ACM Internet Measurement Conference* (Nice, France) *(IMC '22)*. ACM, New York, NY, USA, 589–604.

[31] Siyuan Tang, Xianghang Mi, Ying Li, XiaoFeng Wang, and Kai Chen. 2022. Clues in Tweets: Twitter-Guided Discovery and Analysis of SMS Spam. In *Proceedings of the 2022 ACM SIGSAC CCS* (Los Angeles, CA, USA). ACM, New York, NY, USA, 2751–2764.

[32] Sergei Tikhomirov. 2018. Ethereum: state of knowledge and research perspectives. In *Foundations and Practice of Security: 10th International Symposium, FPS 2017, Nancy, France, October 23-25, 2017, Revised Selected Papers 10*. Springer, 206–221.

[33] Christof Ferreira Torres, Antonio Ken Iannillo, Arthur Gervais, and Radu State. 2021. The Eye of Horus: Spotting and Analyzing Attacks on Ethereum Smart Contracts. *arXiv preprint arXiv:2101.06204* (2021).

[34] Petar Tsankov et al. 2018. Securify: Practical Security Analysis of Smart Contracts. In *CCS*.

[35] uniswap. 2023. https://app.uniswap.org/

[36] Marie Vasek and Tyler Moore. 2015. There's no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency scams. In *Financial Cryptography and Data Security*. Springer, 44–61.

[37] Marie Vasek and Tyler Moore. 2019. Analyzing the Bitcoin Ponzi scheme ecosystem. In *Financial Cryptography and Data Security: FC 2018 International Workshops, BITCOIN, VOTING, and WTSC, Nieuwpoort, Curaçao, March 2, 2018, Revised Selected Papers 22*. Springer, 101–112.

[38] Jinhuan Wang, Pengtao Chen, Xinyao Xu, Jiajing Wu, Meng Shen, Qi Xuan, and Xiaoniu Yang. 2021. TSGN: Transaction Subgraph Networks Assisting Phishing Detection in Ethereum. In *International Conference on Blockchain and Trustworthy Systems*. 187–200.

[39] Wei Wang, Jingjing Song, Guangquan Xu, Yidong Li, Hao Wang, and Chunhua Su. 2021. ContractWard: Automated Vulnerability Detection Models for Ethereum Smart Contracts. *IEEE Transactions on Network Science and Engineering* 8, 2 (2021), 1133–1144.

[40] Xinming Wang, Jiahao He, Zhijian Xie, Gansen Zhao, and Shing-Chi Cheung. 2019. ContractGuard: Defend ethereum smart contracts with embedded intrusion detection. *IEEE TSC* (2019), 314–328.

[41] Zeli Wang, Hai Jin, Weiqi Dai, Kim-Kwang Raymond Choo, and Deqing Zou. 2021. Ethereum smart contract security research: survey and future research opportunities. *Frontiers of Computer Science* 15 (2021), 1–18.

[42] Jiajing Wu, Qi Yuan, Dan Lin, Wei You, Weili Chen, Chuan Chen, and Zibin Zheng. 2022. Who Are the Phishers? Phishing Scam Detection on Ethereum via Network Embedding. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 52, 2 (2022), 1156–1166.

[43] Pengcheng Xia, Haoyu Wang, Bingyu Gao, Weihang Su, Zhou Yu, Xiapu Luo, Chao Zhang, Xusheng Xiao, and Guoai Xu. 2021. Trade or trick? detecting and characterizing scam tokens on uniswap decentralized exchange. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5, 3 (2021), 1–26.

[44] Yijun Xia, Jieli Liu, and Jiajing Wu. 2022. Phishing Detection on Ethereum via Attributed Ego-Graph Embedding. *IEEE Transactions on Circuits and Systems II: Express Briefs* 69, 5 (2022), 2538–2542.

[45] Liping Chen Xuezhi He, Tan Yang. 2022. CTRF: Ethereum-Based Ponzi Contract Identification. *Security and Communication Networks* 2022 (2022), 10.

[46] Mengya Zhang, Xiaokuan Zhang, Yinqian Zhang, and Zhiqiang Lin. 2020. TXSpector: Uncovering Attacks in Ethereum from Transactions. In *USENIX Security*.

[47] Yanmei Zhang, Wenqiang Yu, Ziyu Li, Salman Raza, and Huaihu Cao. 2022. Detecting Ethereum Ponzi Schemes Based on Improved LightGBM Algorithm. *IEEE Transactions on Computational Social Systems* 9, 2 (2022), 624–637.