

# Understanding (Mis)Behavior on the EOSIO Blockchain

YUHENG HUANG, Beijing University of Posts and Telecommunications, China

HAOYU WANG\*, Beijing University of Posts and Telecommunications, China

LEI WU\*, Zhejiang University, China

GARETH TYSON, Queen Mary University of London, United Kingdom

XIAPU LUO, The Hong Kong Polytechnic University, China

RUN ZHANG, Beijing University of Posts and Telecommunications, China

XUANZHE LIU, Peking University, China

GANG HUANG, Peking University, China

XUXIAN JIANG, PeckShield, Inc., China

EOSIO has become one of the most popular blockchain platforms since its mainnet launch in June 2018. In contrast to the traditional PoW-based systems (e.g., Bitcoin and Ethereum), which are limited by low throughput, EOSIO is the first high throughput Delegated Proof of Stake system that has been widely adopted by many decentralized applications. Although EOSIO has millions of accounts and billions of transactions, little is known about its ecosystem, especially related to security and fraud. In this paper, we perform a large-scale measurement study of the EOSIO blockchain and its associated DApps. We gather a large-scale dataset of EOSIO and characterize activities including money transfers, account creation and contract invocation. Using our insights, we then develop techniques to automatically detect bots and fraudulent activity. We discover thousands of bot accounts (over 30% of the accounts in the platform) and a number of real-world attacks (301 attack accounts). By the time of our study, 80 attack accounts we identified have been confirmed by DApp teams, causing 828,824 EOS tokens losses (roughly \$2.6 million) in total.

CCS Concepts: • **Information systems** → **Web mining**; • **Security and privacy** → **Intrusion/anomaly detection and malware mitigation**; **Web application security**;

Keywords: EOSIO; blockchain; bot account; DApp; attack detection

## ACM Reference Format:

Yuheng Huang, Haoyu Wang, Lei Wu, Gareth Tyson, Xiapu Luo, Run Zhang, Xuanzhe Liu, Gang Huang, and Xuxian Jiang. 2020. Understanding (Mis)Behavior on the EOSIO Blockchain. In *Proc. ACM Meas. Anal. Comput. Syst.*, Vol. 4, 2, Article 37 (June 2020). ACM, New York, NY. 29 pages. <https://doi.org/10.1145/3392155>

\*Corresponding Authors: Haoyu Wang ([haoyuwang@bupt.edu.cn](mailto:haoyuwang@bupt.edu.cn)) and Lei Wu.

Authors' addresses: Yuheng Huang, Beijing University of Posts and Telecommunications, Beijing, China; Haoyu Wang, Beijing University of Posts and Telecommunications, Beijing, China, [haoyuwang@bupt.edu.cn](mailto:haoyuwang@bupt.edu.cn); Lei Wu, Zhejiang University, Hangzhou, China; Gareth Tyson, Queen Mary University of London, London, United Kingdom; Xiapu Luo, The Hong Kong Polytechnic University, HongKong, China; Run Zhang, Beijing University of Posts and Telecommunications, Beijing, China; Xuanzhe Liu, Peking University, Beijing, China; Gang Huang, Peking University, Beijing, China; Xuxian Jiang, PeckShield, Inc. Beijing, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

2476-1249/2020/6-ART37 \$15.00

<https://doi.org/10.1145/3392155>

## 1 INTRODUCTION

Blockchain technologies, such as Bitcoin [25], have experienced much hype in recent years. They have been proposed for applications in many areas, including financial and logistical systems. For example, the Internet giant Facebook recently announced their plans for a cryptocurrency [24]. Blockchain technologies have found particular popularity with decentralized application (DApp) developers, most notably for creating smart contracts. These consist of a decentralised protocol which is capable of digitally negotiating an agreement in a cryptographically secure manner. As the most widely used blockchain system after Bitcoin, Ethereum [26] offers support for executing smart contracts, yet it suffers from poor performance (as with Bitcoin) due to the reliance on Proof-of-Work (PoW) consensus protocols.

This has motivated researchers to propose new blockchain approaches that employ more efficient consensus mechanisms. One particularly prominent example is *EOSIO*, the largest Initial Coin Offering (ICO) project to date (over \$4 billion). *EOSIO* adopts a Delegated Proof-of-Stake (DPoS) consensus protocol. This allows *EOSIO* to achieve far higher performance throughput, i.e., up to 8,000 Transactions Per Second (TPS) within a single thread, and unlimited for multiple-threaded cases [2]. As a result, *EOSIO* has grown rapidly and successfully surpassed Ethereum in DApp transactions just three months after its launch (in June 2018). For example, a recent report demonstrated that the average amount of EOS (the *EOSIO* currency) traded in 24 hours has achieved 57 million (with a peak exceeding 80 million) [18]. As a comparison, Bitcoin has an average of 825 thousand transactions and Ethereum has an average of 717 thousand transactions.

Consequently, *EOSIO* has attracted significant attention from both industry and research communities alike. Criticisms, however, have started to emerge, accusing *EOSIO* of suffering from *superficial prosperity*, i.e., it has a large number of transactions, yet the majority of users are inactive [3, 27]. Furthermore, recent years have witnessed attacks against *EOSIO*, exploiting vulnerabilities in DApps. This has resulted in millions of dollars lost [14, 17, 23]. Despite these anecdotes [56, 63], *we still lack a comprehensive understanding of EOSIO's operation in the wild, especially the severity of problems that EOSIO faces.*

To rectify this, we present a detailed study of the *EOSIO* ecosystem at scale, longitudinally and across various dimensions. To this end, we first gather a large-scale dataset containing both on-chain data of *EOSIO* and off-chain data related to DApps and attacks (**Section 3**). Our dataset consists of over 3 billion transactions, over 1 million *EOSIO* accounts, thousands of bots and a number of attack reports. Based on the collected dataset, we then perform an explorative study to characterize the activities on *EOSIO* (**Section 4**), including money transfer, account creation, and contract invocation. Following this, we identify bot-like accounts and fraudulent activities by mining the relations and behavioral similarities among millions of accounts (**Section 5**). We further investigate their incentives and purposes. Finally, we characterize security issues in *EOSIO*, including permission misuse issues and attacks (**Section 6**).

To the best of our knowledge, this is the first comprehensive study of the *EOSIO* blockchain at *scale, longitudinally, and across various dimensions*. We have revealed a range of serious misbehaviors. Among many interesting results and observations, the following are prominent:

- **The overall ecosystem follows the Pareto principle.** Although the overall ecosystem shows a growing volume of transactions (over 1 billion transfers), *EOSIO* is dominated by a small percentage of accounts. The top 0.47% of accounts constitute 89% of the total transaction volume. Cryptocurrency exchanges and gambling DApps dominate the transactions. **Over 32% of the accounts are “silent”**: they have never actively initiated any transaction with other accounts.

- **Bot-like accounts are prevalent.** We flag over 30.82% of the accounts (381,837) as bot-like, with over 188 million transactions, and 776 million EOS transferred. These bots are mainly used for malicious and fraudulent purposes including Bonus Hunting, Clicking Fraud, etc.
- **Permission misuse issues are overlooked by users.** We identify permission misuse issues of 5,541 accounts, i.e., granting their “eosio.code” permissions to other accounts, which could cause serious security issues (e.g., the accounts with that granted permission can stealthily transfer users’ EOS tokens even without their attentions).
- **EOSIO suffers from a number of serious attacks.** We identify over 301 suspicious attack accounts, causing over 1.5 million EOS losses. We have reported them to the DApp teams (developers): 80 of the attacks (with 828,824 EOS losses) have been confirmed by the time of our study. We further assist the DApp teams in tracing the losses.

We develop core methodologies to trace attacks and fraud, as well as deriving key insights into EOSIO. Our efforts contribute developer awareness, inform the activities of the research community and regulators, and promote better operational practices. We have released our dataset and experiment results to the research community at Github:

[https://github.com/EOSIO-analysis/EOSIO\\_open\\_source\\_data](https://github.com/EOSIO-analysis/EOSIO_open_source_data)

## 2 BACKGROUND

We start by briefly presenting key concepts. Note that this is intended as an overview; we refer readers to [15] for full details.

### 2.1 Overview of EOSIO

EOSIO is a decentralized enterprise system that executes industrial-scale DApps [20] — software which relies on the EOSIO blockchain to cryptographically record *transactions*. The most common transaction is transferring the EOSIO currency token, named *EOS*. In contrast to Bitcoin or Ethereum, EOSIO is a DPoS-based system, which can scale to millions of transactions per second, making EOSIO an attractive option for new DApp developers.

There are four key concepts to understand within EOSIO. If an entity wishes to interact with the EOSIO blockchain, it must first create an *account*. This unique identity can then invoke a *smart contract* through a *transaction*, which consists of one or more *actions* to perform. For example, an action might be transferring an EOS token from one account holder to another. To enable this process, accounts wishing to invoke a contract must first delegate appropriate *permissions*, granting it the privileges to act on its behalf. The rest of this section describes in detail these four concepts.

### 2.2 EOSIO’s Smart Contract and Transactions

EOSIO adopts C++ as the official language for DApp developers to develop smart contracts. In particular, the source code of smart contract is first compiled down to WebAssembly (aka Wasm) bytecode. Upon invocation, the bytecode will then be executed in EOSIO’s Wasm VM, resulting in transactions recorded on the blockchain, e.g., transferring EOS. As the basic element of communication between smart contracts, an *action* is a base32 encoded 64-bit integer which can be used to represent a single operation. There are two types of actions: *external action*, when the user calls an action directly from the outside; and an *inline action*, where an inline action refers to a call to another action in a smart contract (same or external).

Each transaction can consist of one or more actions. In EOSIO, there are two ways to send actions for communication: an *inline action* that performs an operation within the same transaction as the original action, and *deferred action* that might be scheduled to perform operation in a future (deferred) transaction. Inline actions are guaranteed to execute synchronously, while deferred

actions will be executed asynchronously if being scheduled. The transaction will be rolled back if an inline action fails or raises an exception.

### 2.3 EOSIO's Account Management

In EOSIO, accounts are the entities that can execute transactions. An EOSIO account is a human-readable name (up to 12 characters) recorded on the blockchain. In practice, accounts are authorization structures that can define senders and receivers of contracts. Accounts can also grant permissions to contracts and be configured to provide individual or group access to transfer/push any valid transactions to the blockchain.

Note that an EOSIO account is different from (and more complicated than) that of Ethereum. Most notably, accounts are hierarchical and can only be created by an existing account, thereby creating a tree structure. This means the resources required to create new accounts must be allocated by existing accounts (except the first privileged account named *eosio* – the root of the tree – created by the blockchain system when the mainnet was launched). This inevitably consumes system resource (RAM) and therefore the account creation of EOSIO is not free.<sup>1</sup> This characteristic does affect the behaviors of the blockchain, and determines the necessity and methods to perform the relevant analyses in Section 4 and Section 5.

Permissions associated with an EOSIO account are used to authorize actions and transactions to other accounts [36]. Specifically, the account can assign public/private keys to specific actions, and a particular key pair will only be able to execute the corresponding action. By default, an EOSIO account is attached to two public keys: the *owner* key (which specifies the ownership of the account) and the *active* key (which grants access to activities with the account). These two keys authorize two native named permissions: the *owner* and *active* permission, to manage accounts. Apart from the native permissions, EOSIO also allows customized named permissions for advanced account management.

### 2.4 Example

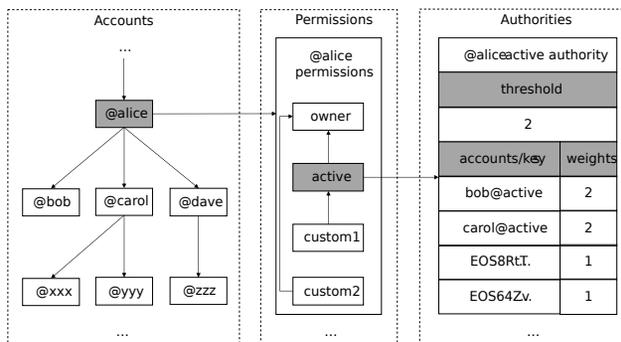


Fig. 1. An example of EOSIO accounts.

Fig. 1 gives an example to demonstrate the relationship between accounts, associated permissions and authorities. These account permissions can be delegated to actions, such that transactions can be performed on their behalf. The left most part of the figure shows the hierarchical structure (i.e., tree) of the accounts. The account named *alice* (marked as grey) is the root of the tree. Alice has created three other accounts: *bob*, *carol* and *dave*. The central part of the figure shows that *alice* has

<sup>1</sup>More precisely, one has to buy RAM to store the account data [37].

4 permissions, including two native permissions (*owner* and *active*) and two customized permission (*custom1* and *custom2*) respectively. The right most part of the figure gives the authority table of the grey marked permission *active*. This cryptographically states which permissions have been allocated to the children accounts of *alice*. Note that there exists a threshold that must be reached to authorize the execution of the action. In addition, in order to be executed by or on behalf of *alice*, a weight threshold of 2 must be reached as well [36]. The permissions and authorities will be used in Section 6.1.

### 3 STUDY DESIGN & DATA COLLECTION

We seek to focus on the following three research questions (RQs):

- RQ1 What are the characteristics of accounts and their transaction behaviors in the EOSIO ecosystem?** No previous work has characterized the EOSIO ecosystem, including the health of the platform and the various crypto-assets within. As previous work has investigated the Ethereum blockchain [45], we wish to compare them and understand the difference between these two blockchain platforms.
- RQ2 How severe is the presence of bot activities in the ecosystem?** As many applications experience a range of bot activities [9], we wish to investigate whether blockchains are also inundated with bot activities. Although millions of transactions are emerging on the EOSIO platform, it is unknown *how many of them are manipulated by bots and how many accounts/activities are fake*.
- RQ3 Can we identify real-world security issues by analyzing the accounts and transactions we collect?** As one of the most popular platforms for DApps, EOSIO has always been the target for hackers. A number of reports have already revealed attacks, leading to millions of dollars lost. Thus, it is interesting to explore whether we can identify real-world attacks and build an early warning system.

#### 3.1 Data Collection

To provide a comprehensive analysis of EOSIO, we first seek to harvest both *on-chain data* and *off-chain information* (cf. Table 1):

##### 1) On-chain data of EOSIO blockchain:

- **Transaction Records.** To enforce a fine-grained analysis, we use *actions* to measure the transaction activities because actions are the basic unit that constitute a transaction in EOSIO, as mentioned in Section 2. However, due to the volume of on-chain data, it is not feasible to fetch them either by querying public API endpoints or by crawling from the blockchain explorer. To solve the problem, we have built a customized EOSIO client, which can be used to synchronize with the mainnet in an efficient way. Specifically, the core service daemon (named *nodeos*) of the official client provides different approaches as plugins to synchronize the data. We further customize the MongoDB plugin (i.e., *mongo\_db\_plugin*) to accelerate the performance. As a result, we can archive blockchain data through MongoDB directly rather than fetching data from the client<sup>2</sup>. Note that we also collected all the notifications on the EOSIO blockchain, which could be used to facilitate the detection of specific kinds of attacks (e.g., Fake EOS Transfer attack, see Section 6).
- **Account Information.** An EOSIO account has many attributes, including creation information (e.g., creator and creation time), owned system resources, assigned keys and permissions. The creation information can be traced by crawling blockchain explorers (we take advantage

<sup>2</sup> *mongo\_db\_plugin* was deprecated on June, 2019. The *state\_history\_plugin* and the *history-tools* are recommended officially [38].

Table 1. Overview of Our Dataset (On-chain And Off-chain)

Category	Amount
Action trace	3, 208, 168, 339 actions
Account information	1, 239, 030 accounts
DApp accounts	1, 513 accounts
Bot accounts	63, 956 accounts
Attack accounts	37 accounts

of EOSPark). Other information can be queried through public API endpoints [8]. Many engineering efforts were required to overcome the heterogeneous structures of different explorers and APIs. To this end, we have implemented a generic collector that is capable of crawling and querying necessary account information.

## 2) Off-chain data related to DApps, Bots and Attacks:

- **DApp Information.** As accounts on the EOSIO all have human-readable names, finding whether an account belongs to a DApp is more straightforward than Ethereum. By crawling data from websites including DAppTotal [13] and DAppReview [12], we have annotated accounts with their DApp. To the best of our knowledge, we have collected the most complete list of Dapp accounts available.
- **Bot Accounts.** Blockchain bots have been revealed by DApp teams and blockchain security companies [27, 33]. We have collected 63,956 bot-like accounts with the help of PeckShield [29], which will be used as our ground-truth in bot accounts identification (see Section 5).
- **Attack Information.** A number of attacks on EOSIO have already been observed in the wild. To measure the severity, we collect existing attack information, including date, participants (victims and attackers) and damage, by monitoring and collecting security news and blogs from well-known blockchain security companies. Based on this ground-truth dataset, we implement a monitoring system to perform attack detection and forensics (see Section 6).

In summary, we have collected transactions from 2018.06.09 to 2019.05.31, with over 3 billion actions in total. We also collected information for all 1,239,030 accounts (including 2, 482, 192 keys). Moreover, we have collected 1, 513 DApp accounts, 63,956 bot-like accounts and 40 attack events (including 37 attack accounts).

## 3.2 Study Approach

To answer RQ1, we are focused on three important behaviors, including money transfer, account creation and contract invocation, mainly based on graph analysis. To answer RQ2, we perform an analysis of the account relationships to pinpoint bot candidates, and then perform behavior-level similarity comparison to identify real bot accounts. To answer RQ3, we first provide a taxonomy of the attacks found in the EOSIO platform. Then we further seek to explore the characteristics of different kinds of attacks with regard to their behaviors at the transaction-level. Based on the summarized behaviors, we have implemented a warning system to flag suspicious attacks and then perform manually verification.

## 4 GENERAL OVERVIEW OF EOSIO

We first conduct a comprehensive investigation of EOSIO, exploring *money transfers*, *account creation* and *contract invocation*.

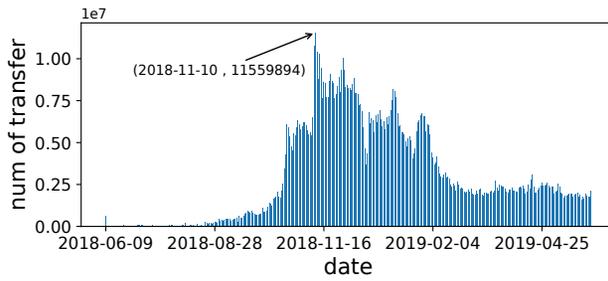


Fig. 2. The distribution of money transfers over time.

### 4.1 Money Transfers

**4.1.1 Overview of Money Transfers.** The total number of money transfer transactions is over 1 billion (1,055,690,229), and the total number of EOS tokens being transferred is 15,190,552,483.9523 EOS tokens. This represents over \$47 billion market value<sup>3</sup>. Fig. 3 shows the number of transfers and accounts involved over time. We see that the number of transfers achieves its peak between 2018.11.09 to 2018.11.14. This is mainly due to the rising popularity of four gambling games (EOS Max, Dice, FAST and EOSJacks), as they cover 59.66 % of the total transfers during that time.

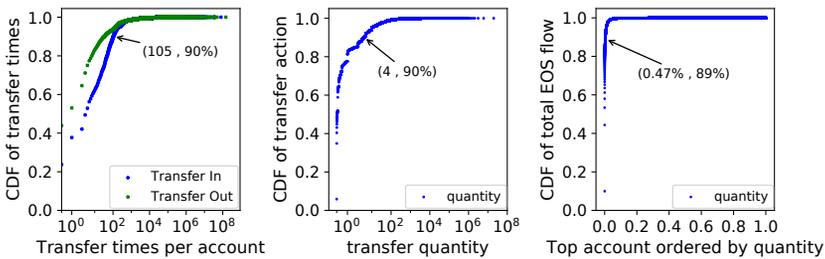


Fig. 3. Distributions of transfer frequency and quantity.

The distributions of transfer frequency and quantities are further shown in Fig. 3. The number of accounts involved in transfers is 944,907; in other words, 294,123 (23.74% of all) accounts do not perform *any* transactions. Fig. 3(a) shows that the transfer frequency of the majority of nodes is below 100, which implies that most accounts are not particularly active. One interesting detail is that 249,644 accounts (20.15% of total) exclusively *receive* EOS tokens but never transfer them out. Fig. 3(b) also shows the quantity exchanged within each transaction. For the 1 billion money transfers, most are small: over 90% are under 4 EOS. We further analyze the total amount of money transfers for each account (see Fig. 3(c)), and find that 0.47% of the accounts make up approximately 90% EOS tokens being transferred, which is a typical *Pareto effect*.

**4.1.2 Graph Modeling of Money Transfers.** We next inspect which accounts perform transfers. To achieve this, we construct a graph from the set of  $\langle from, to, value \rangle$  transaction tuples. However, this would ignore time information, which is important and necessary to capture abnormal behaviors. Therefore, we construct an Enhanced Money Flow Graph (EMFG) with timestamps, as follows:

<sup>3</sup>We use the price up to October 2019 to calculate the market value.

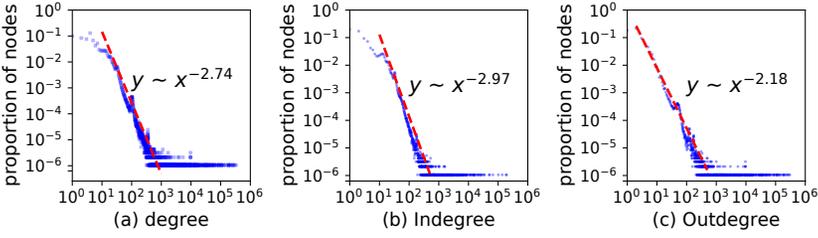


Fig. 4. The degree distributions of EMFG.

Table 2. A Comparison of The Graph Metrics. Note that the numbers in **bold** are the metrics of EOSIO graphs, while the numbers in (brackets) are for Ethereum graphs [45].

Metrics	EMFG	EACG	ECIG
Clustering	<b>0.5049</b> (0.17)	<b>0</b> (0)	<b>0.2375</b> (0.004)
Assortativity	<b>-0.3392</b> (-0.12)	<b>-0.1539</b> (-0.35)	<b>-0.1834</b> (-0.2)
Pearson	<b>0.2448</b> (0.44)	<b>/</b> (/)	<b>-0.00068</b> (0.11)
# SCC	<b>251,099</b> (466, 095)	<b>1,239,028</b> (622, 158)	<b>610,281</b> (682, 984)
Largest SCC	<b>693,632</b> (1, 822, 192)	<b>1</b> (1)	<b>688</b> (84)
# WCC	<b>1</b> (81)	<b>3</b> (22, 260)	<b>3</b> (4, 088)
Largest WCC	<b>944,907</b> (2, 291, 707)	<b>1,239,028</b> (126, 246)	<b>611,083</b> (668, 891)

$$EMFG = (V, E, D, w), E = (v_i, v_j, D_k), v_i, v_j \in V, D_k \subseteq D$$

The order of the nodes in an edge indicates the direction of transferred money. Each edge has a set of time attributes  $D_k$ . For  $d \in D_k$ ,  $2018.06.09 \leq d \leq 2019.05.31$ , indicating when the transfer occurs (UTC time). If account  $v_i$  transfers multiple EOS to  $v_j$  across multiple days, there will be more than one timestamp for edge  $(v_i, v_j)$ .  $w : (v_i, v_j, d) \rightarrow \mathbb{R}^+$  associates each edge  $(v_i, v_j)$  on a particular day  $d$  with the transfer quantity.

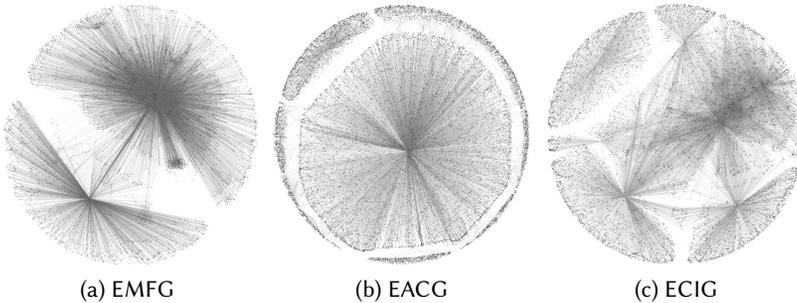


Fig. 5. Visualization of the constructed graphs.

To measure the properties of the financial transactions, we apply some well-defined network metrics. The clustering coefficient [11] measures the tendency that two nodes in a network cluster together; this is calculated using the approach of [47]. The assortativity coefficient [61] measures the

preference for nodes to attach to others, i.e., a network is assortative when high degree nodes are, on average, connected to other nodes with high degree, and vice versa. Assortativity is calculated based on [60]. We use Pearson coefficient [28] to measure the linear correlation between indegree and outdegree. We also extract the Weakly Connected Components (WCC), and Strongly Connected Components (SCC). Note that, *we will use the same metrics to measure the graphs we construct for money transfer, account creation and contract invocation.*

**4.1.3 Results of EMFG Graph Modeling.** Overall, there are 10, 438, 158 edges and 944, 907 nodes. Fig. 5 is a partial visualization of EMFG with 5, 000 nodes (19, 099 edges) randomly selected. The overall-degree, in-degree and out-degree distributions of EMFG are shown in Fig. 4, and all of them satisfy the power law distribution.

The metrics for the constructed EMFG are shown in Table 2. For context, we compare against equivalent metrics for Ethereum, taken from [45] (note that although this study was performed in 2018, it still offers a useful comparison). First, we see the EMFG is a single WCC because of the specific account named *eosio* (the first privileged account created by the system), which transferred to 163, 937 other accounts during the mainnet launch. This is different from Ethereum, which contains 81 WCCs. The largest SCC contains 73.41% of all nodes, which is similar to that of Ethereum (86.59%). This result suggests that EOS tokens can flow from one SCC to another one but will not be transferred back (i.e., unidirectional transfer).

The clustering coefficient of EMFG is 0.5049, i.e., the likelihood of triadic closure between accounts, is almost three times as large as that of Ethereum (0.17 [45]). The negative assortativity coefficient also implies that high-degree nodes are more likely to connect to nodes with lower degree. Such a phenomenon may come from the existence of accounts having one-to-many mapping relationship with others, such as exchanges, DApp games and some specific accounts like *eosio* mentioned earlier.

We also identify the most *central* accounts (measured by PageRank) in the EMFG. For the top-10 accounts, 3 accounts are cryptocurrency exchanges and 4 accounts are gambling DApps. Comparing with the previous Ethereum study [45], i.e., 8 exchanges and 0 gambling DApp, the composition of top DApps suggests that EOSIO is very much about speculation of value because nearly half of the top DApps are gambling games. We further analyze the categories of DApp accounts, and find that the total volume of Gambling DApps has occupied 78.88% of the overall DApp volume, which further supports our observation.

**Findings #1:** *Most money transfers via EOSIO are small transfers (< 4 EOS). Although the overall ecosystem shows a promising volume, EOSIO is dominated by a small percentage of accounts (i.e., the top 0.47% of accounts cover 90% of the total volume). Cryptocurrency exchanges and gambling DApps dominate the transactions.*

## 4.2 Account Creation

**4.2.1 Overview of Account Creation.** We next inspect the account creation properties that we observe. The times series of daily account creation is shown in Fig. 6. There are several interesting peaks. The first stems from the token registration during the mainnet launch. Another peak, arising between 2019.04.23 and 2019.04.29, may seem a little strange though. 178, 018 accounts were created during this period, belonging to two wallet accounts: *trxcashstart* and *eastokenhome*. However, only a few accounts were directly created by them: 13, 292 for the former and 13, 014 for the latter. The majority of the remaining accounts (from the 178, 018) were then indirectly created by the children

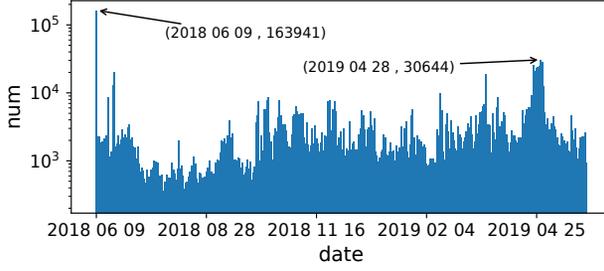


Fig. 6. Distribution of account creation.

of *eastokenhome*. This “spawning” process is a bi-product of how accounts are created, hence we give further analysis in Section 5.

This sporadic spawning of new accounts may suggest that not all are live and active. Hence, we next inspect if they, indeed, perform any transactions. Although creating an account expends system resources (like RAM, see Section 2.3), *the liveness of accounts on EOSIO is far below expectation*.

Overall, 31.75% accounts (i.e., 393, 430) are “**silent**” – they never actively initiate any transactions with other accounts and never invoke any smart contracts. The proportion of silent accounts can therefore be used to reflect the liveness of EOSIO accounts. Even those accounts that *do* perform transactions, are rarely used. Over 66% of accounts performed 10 or fewer transactions. Such a circumstance inevitably reveals the superficial prosperity of the EOSIO blockchain, as most of the accounts are silent.

**4.2.2 Graph Modeling of Account Creation.** Intuitively, the relationship formed during account creation activities can be modelled as a graph. This is because new accounts can only be created by existing ones (see Section 2). Hence, we define and construct an Enhanced Account Creation Graph (EACG), as follows:

$$EACG = (V, E, D), E = (v_i, v_j, d), v_i, v_j \in V, d \in D.$$

An edge  $(v_i, v_j)$  indicates that account  $v_i$  creates an account  $v_j$  on day  $d$ , where  $2018.6.9 \leq d \leq 2019.5.31$  (UTC time zone), indicating when the account was created.

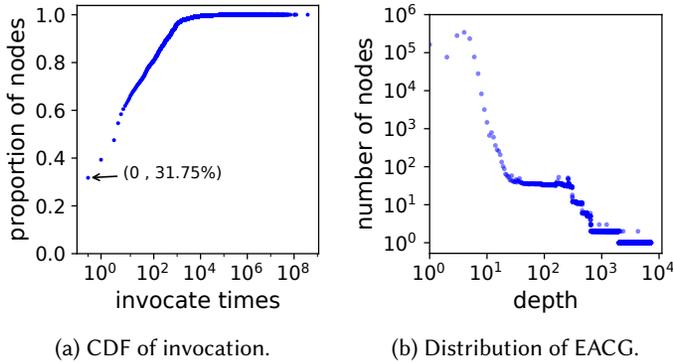


Fig. 7. Account invocation and creation.

Table 3. Top-5 Nodes of EACG Using Degree Centrality

Account	Outdegree	Identity
eosio	163,946	Official account
senseaccount	50,531	Sense Chat
eosaccountwm	45,618	MEET.ONE
eostokenhome	35,908	ET Wallet
lynxlynxlynx	33,018	EOS LYNX

4.2.3 *Result of EACG Graph Modeling.* We compute our earlier graph metrics for the EACG, and present them in Table 2. The EACG consists of 1,239,027 edges and 1,239,030 accounts in total. In order to provide a more intuitive understanding of the graph, Fig. 5(b) gives a visualization for part of EACG. The account named *imtokenstart* is selected as the root node, and the graph consists of 16,382 nodes and their corresponding 16,381 edges. This account is an example of a public service for account creation.

We see that the EACG is a single tree when excluding a few isolated official accounts (*eosio.prods*, *eosio.null*), and the root node of the whole EACG, *eosio*. Fig. 7b shows the distribution of account trees by depth. We see that EACG is a wide tree with a few deep paths. We see that from depth 651 to depth 7,198, there only exist *two* paths (see Fig. 8). The first path has a root node named *dogaigaohvwj*; it has a height of 7,195, and was created to transfer illegal profits of a publicly known hacker account *hnihpyadbunv* [6]. This attack was initially reported to include just 2,190 sub-accounts [7], yet we find this misses 5,005 accounts. The second path has a root node named *chengcheng21*; it has a height of 2,000. Interestingly, a DApp named *VSbet* received 95.30% of the total EOS it transferred out. The activities of these accounts are quite suspicious. Almost all of them had transactions with *VSbet*, and the total amount of EOS they transfer to *VSbet* was almost identical to what they received; we later revisit this trend (see Section 5).

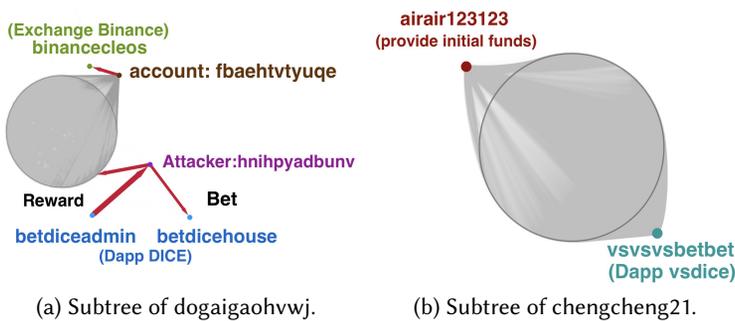


Fig. 8. Visualization of the two anomalous paths.

Briefly, Table 3 lists the top-5 central nodes. All listed accounts are related to public services for account creation, allowing users to obtain new accounts. The existence of such services breaks the assumption that any new accounts are controlled by their creators. As we later revisit, this becomes an obstacle to identify attackers and bots on EOSIO.

**Findings #2:** Over 30% of accounts are “silent” and have never initiated any transactions. The constructed EACG graph is a wide tree with several deep paths, while the outliers mainly belong to attack and fraudulent accounts.

### 4.3 Contract Invocation

Finally, we inspect the contract invocation activities observed within EOSIO. There are 4,453 smart contracts in EOSIO, which have been invoked 2,140,945,703 times. The time series of contract invocation is shown in Fig. 9, which shows that creation has been stable after rapid growth in 2018.

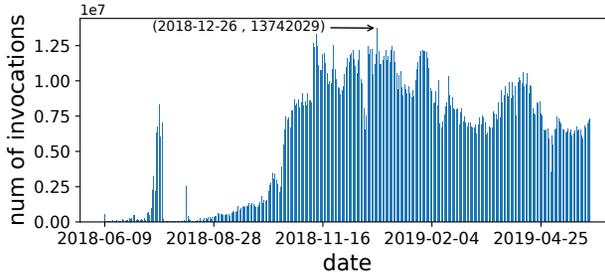


Fig. 9. The number of contract invocations over time.

**4.3.1 Graph Modeling of Contract Invocation.** Again, we can model contract invocations as a graph. Specifically, we define the Enhanced Contract Invocation Graph (ECIG), as follows:

$$ECIG = (V, E, D, A, f), E = (v_i, v_j, D_k, A_k), \\ v_i, v_j \in V, D_k \subseteq D, A_k \subseteq A.$$

An edge  $(v_i$  and  $v_j)$  represents that an account  $v_i$  invokes a smart contract  $v_j$ . To facilitate the identification of abnormal behavior, we annotate each edge with extra attributes. Each edge has a set of timestamp attributes  $D_k$ . For  $d \in D_k$ ,  $2018.6.9 \leq d \leq 2019.5.31$ , indicating when the invocation occurs (in UTC time zone). The name of the function for each invocation is also recorded as  $a$ ,  $a \in A$ . Each edge has a set of invocation actions, namely  $A_k$ .  $f : (v_i, v_j, d, a) \rightarrow \mathbb{Z}^+$  assigns each edge  $(v_i, v_j)$  with a particular day  $d$  and given action  $a$  the number of invocation.

**4.3.2 Results of ECIG Graph Modeling.** Overall, there are 4,313,739 edges and 611,085 nodes. Fig. 5(c) is a partial visualization of the ECIG with 10,626 nodes (and 19,079 edges) being randomly selected. We compute our earlier metrics, and present the results in Table 2. Because of the existence of the system contract, the largest SCC of EOSIO is larger than that of Ethereum, although there are only 4,453 contracts being invoked on EOSIO.

Table 4 lists the top-5 most important nodes in the ECIG. An intriguing observation is that no exchanges are present in Table 4. This is quite different from Ethereum (7 reported by a recent study [45]). Notice that the number of invocations by these nodes are quite high compared to their degrees. Combined with the observation from Fig. 7a, we can confirm that most accounts rarely invoke smart contracts. Obviously, a small set of accounts are much more active than the others: the top 0.01% of accounts perform 80% of contract invocations. Our manual investigation suggests that there are many bots on EOSIO, as we will explore in Section 5.

Table 4. Top-5 Nodes of ECIG Using Degree Centrality

Account	Degree	# of Invocation	Identity
eosio	383,590	8,925,999	Official account
betdicelucky	109,151	48,694,997	DApp Dice
pornhashbaby	107,417	64,072,844	DApp Hash Baby
hashbabycoin	96,826	31,380,78	DApp Hash Baby
endlessdicex	93,469	23,955,881	DApp Endless Game

**Findings #3:** *Although there are only 4,453 contracts in the platform, EOSIO has a large number of contract invocations. A small number of accounts are much more active than the others. The top 0.01% of accounts perform 80% of contract invocations, while the majority of accounts rarely invoke contracts.*

## 5 CHARACTERIZING BLOCKCHAIN BOTS

It is well-known that online activities are impacted by bots [9, 49]. We posit that these might heavily also impact the EOSIO ecosystem, and therefore we next investigate *how many activities are driven by bots, and what are their incentives (purposes)*.

### 5.1 A Primer on Bots

We define **bot-like account** as those operated by machines (e.g., programs). These are already known to exist in EOSIO, e.g., the famous DApp team “pornhashbaby” has reported 8 groups of bots, and each group has hundreds to thousands of accounts [10]. These past findings have shown that bot accounts usually operate in groups, and have repetitive behavioral patterns. In this paper, we refer to a group containing **bot-like accounts** operated by a controller as a *bot-like community*. Note, the bot accounts may be created by the same controller account (in the same sub-tree of EACG) or by a number of different accounts (across different sub-trees of EACG). We will have two complementary approaches to detect them.

### 5.2 Preliminary Observations

To help distinguish bot accounts from other normal accounts (especially public account creation services), we perform a preliminary study of bot activities. We do this by harvesting a number of ground-truth bot accounts from [10], as well as PeckShield, a well-known blockchain security company. From this, we gather 63,863 EOSIO bot accounts flagged by existing efforts. The accounts were created within 21 bot-like communities. To further differentiate these accounts from normal ones, we label 229,907 normal accounts that belong to 25 public contract services (e.g., official account, exchanges, and pocket) as a white-list for comparison. We then manually inspect these two groups and make several observations from two perspectives: *community-level* and *account-level*.

**5.2.1 Community-level Observation.** First, controller accounts usually create a large number of children accounts (bots). For the flagged 21 controller accounts, the out-degree of them in the EACG varies from 108 to 15,025. In contrast, the average out-degree of the EACG graph is 30, while the median is only 1. Thus, we argue that a “shortlist” of potential bot controllers could be identified using the EACG graph. Note that this is *not* definitive though — many non-bots also have high out-degrees. Hence, we turn to our second observation, where we find that accounts belonging to the same bot-like community tend to share similar behaviors. For instance, they might perform

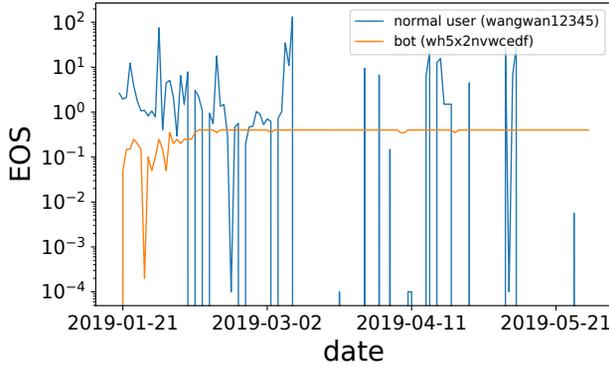


Fig. 10. An example of difference between bots and normal users in EOS transfer.

transactions on the same day using the same contract. We posit that this similarity could be used to automatically group bots belonging to the same controller.

**5.2.2 Account-level Observation.** We further analyze the bots at a per-account level. As bot accounts are controlled by machines, their behaviors usually have regular patterns. Figure 10 shows an example of a labelled bot account and a normal account. The bot account has relatively small but very frequent incoming transfer activities from February 2019 to May 2019. By analyzing its activities, we find that it is performing Click Fraud. Each day it transfers a few EOS to eosvegasjack (an account that belongs to a gambling game) several times. After each transfer, the bot instantaneously gets its money back. The time interval between each transfer is also regular (usually three hours). Thus, we believe that the bot accounts could be classified based on their behavior patterns, including the active time, frequency and volume of transfers, etc.

### 5.3 Detecting bots from community-level

Based on the aforementioned patterns, we next devise a simple algorithm capable of identifying bot-like communities.

**5.3.1 Algorithm Design.** Our approach combines both *account relations* and *behavioral similarities* to flag suspicious bot accounts. The algorithm has two stages. In the *first* step, we extract all accounts that have created in excess of 30 new accounts (which is the average out-degree for nodes in the EACG). This provides a shortlist of accounts that *may* be spawning many accounts for use as bots. However, it is not necessarily accurate, as we have already identified legitimate services that also spawn many accounts. Hence, the *second* step computes the similarity between accounts on the shortlist to identify groups that operate in similar ways. To measure this similarity, we compare accounts across two dimensions.

The first dimension is the **invocation time and frequency**, i.e., the time and frequency of invoking smart contracts and transferring money. Specifically, for each account, we summarize its money transfer actions and smart contract invocations in a feature vector. For each account on day  $i$ , we calculate the frequency of its money transfer actions, and its smart contract invocations, respectively. As we have collected all the action records since the launch of the EOSIO mainnet (which lasts 357 days), the behavior of each account can be represented in a 714 dimension feature vector  $\vec{t}_i$  (357 for the money transfer actions, and 357 for the contract invocations).

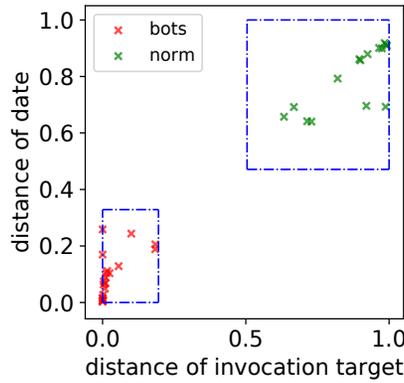


Fig. 11. Bot-like communities VS. normal communities.

Table 5. Bot-like Communities VS. Normal Ones

	Metrics	Bots	Normal
Time Features	Average	0.09	0.78
	Std.Deviation	0.08	0.10
Target Features	Average	0.03	0.87
	Std.Deviation	0.05	0.12

The second dimension is the **contract target** of the transactions. This is because accounts belonging to the same bot operator might invoke the same smart contracts. There are currently 4,453 unique smart contracts in the EOS ecosystem. Thus, for each account  $i$ , we represent its targets in a 4,453 dimension vector  $\vec{s}_i$ . Each dimension represents the number of times that the account has invoked the corresponding smart contract.

The above provides two dimensions for which we can compute similarity between potential bot accounts. We use *cosine distance* to measure the similarity between the feature vectors. Considering there are a number of accounts in a potential bot-like community, we further calculate a *group similarity* for each of them as follows:

- (1) For a community with  $N$  non-silent nodes<sup>4</sup>, we calculate the feature vectors  $\vec{s}_i \in S$  and  $\vec{t}_i \in T$  for each node.
- (2) We further calculate the median vector  $\vec{m}$ ,  $\vec{m} = \frac{1}{N} \sum_{j=1}^N \vec{v}_j$  for vectors in  $S$  and  $T$ .
- (3) We then measure the average distance between each vector and median vector.  $\overline{dist} = \frac{1}{N} \sum_{j=1}^N \text{cosineDist}(\vec{v}_j, \vec{m})$ . As a result we get the overall  $\overline{dist}_S$  and  $\overline{dist}_T$  for  $S$  and  $T$ . Note that  $\overline{dist} \in [0, 1]$ , where larger values means less similarity.

**5.3.2 Defining the Similarity Threshold.** The above allows us to compute the similarity between accounts. The next step is to define an appropriate threshold of similarity that warrants a community of accounts being classified as bots. Here we take an empirical approach, relying on our ground-truth dataset of bots and normal accounts (from Section 5.2). Note that this covers 21 bot communities (63,863 accounts) and the 25 normal ones (314,277 accounts). Figure 11 presents the similarity across these 46 groups. Groups that fall on the lower left have tightly correlated behavior. Specifically, the

<sup>4</sup>We have removed “silent” nodes, as such nodes have no actual behaviors to measure.

Table 6. Features Used to Classify Bot Accounts

feature	bots (avg)	normal (avg)
ACG depth	4.3	1.8
transferIn std	40.8	954
transferOut std	57.3	1267.3
volume per transferIn	2.2	104
volume per transferOut	6.2	553.1
transfer target num	5.9	12.6
invoke contract num	16.1	3.2
invocation num	518.5	2312.7
invocation std	6.2	40.3
activate time	74.6%	13.4%
siblings in same day	1724.1	117024.4

x-axis represents the average distance of the action target feature vectors, and y-axis represents the average distance of the action time and frequency vectors.

Note that by Chebyshev's Theorem, even for non-normally distributed variables, at least 88.8% of cases should fall within properly calculated  $3\sigma$  intervals. [54] Hence, we use the similarity score ranges between  $-3\sigma$  and  $+3\sigma$  of the average as our threshold to filter bots (see the rectangular frame in Fig. 11). Any groups of accounts that fall into this range are deemed to belong to the same bot groups.

**5.3.3 Analyzing the Public Key.** We also observe that some users re-use the same public-private key pair across multiple accounts. Hence, as a final step, we group all the flagged bot-like accounts sharing the same public key into a community.

**5.3.4 Results.** Using the above techniques, we identify 351,740 bot-like accounts, which cover 28.39% of all accounts in the EOSIO platform, and belong to 2,362 communities. The number of accounts in the identified bot communities varies from 30 to 15,025.

## 5.4 Detecting bots from per-account level

Bot accounts are not always created by the same controller account. In this case, the bot accounts may reside within the ACG trees mixed up with normal accounts. To overcome this, we must also perform per-account detection.

**5.4.1 Algorithm Design.** To differentiate accounts within the same ACG tree, we formulate the problem as a binary classification task. Based on the graphs we constructed, we summarize 11 key account features in Table 6. We describe some representative features here. For EACG, the *ACG depth* describes the depth of the account in EACG. Bots appear to be deeper in the tree. For EMFG, as illustrated in Fig.10, bots usually receive EOS more frequently than normal accounts. So we use features like *transferIn std* to describe the standard deviation of transfer volume over time. The amount of EOS per transfer for bots is also much lower than that of normal ones. Thus, we summarize the *volume per transferIn* as a feature. For ECIG, it appears that bots are more active than normal accounts, i.e., the active day is 74.6% (bots) vs. 13.4% (normal) on average. In other words, bots are more active than normal accounts. We use *activate time* to describe the activity. The average values of aforementioned features in the labeled dataset are shown in Table 6.

Thus, we generate a feature vector for each ground truth account, and then train a classifier based on the random forest algorithm. To evaluate the effectiveness of our classifier, we perform the 10-fold cross validation. The results suggest that our approach is promising in identifying bot accounts, i.e., with an average accuracy of 99.55%.

**5.4.2 Results.** We use the classifier to identify bot accounts in the remaining 485,742 accounts. We discover 30,419 new bots. When merged with the community-level results, **we identify 381,837 bot accounts in total.** These accounts have invoked 188,501,976 times, and the amount of transferred EOS is 776,837,813.

## 5.5 Validating Bot Accounts

Before continuing, it is important to validate that the above approaches effectively identifies bot accounts. We use three methods. *First*, we compare our results against the ground-truth bot accounts (see Section 5.2), to find that we gain 100% accuracy — all 63,956 ground-truth bot accounts are identified by our techniques. In fact, by only using the random forest approach on the per-account level, we could achieve an accuracy of 99.54%.

*Second*, we have reported a number of bot communities to the DApp developers and one anonymized blockchain security company, and all of them were confirmed as bots. *Third*, we randomly select 100 bot-like account communities for manual examination. We are confident that over 86% are indeed bots, and we are able to know their purposes (see Section 5.6). For a further 14% of the bot-like accounts, although we cannot reverse engineer what they are doing, we are able to identify some bot-control clues, e.g., the accounts have similar names, a large number of accounts were created at the same time, etc.

## 5.6 Applications of Bot Accounts

We next examine the purposes of these bot accounts. We manually sample 100 bot communities, and examine their actions, identifying 4 major categories, which we subsequently formulate automated methods to detect. Table 7 provides a summary of these categories, highlighting their significant impact of the wider ecosystem. Below, we discuss each category.

**5.6.1 Bonus Hunters.** We find many bot groups performing *bonus hunting*. This involves exploiting incentives in DApps, in order to gain profit. This is because EOSIO DApps often offer incentives to attract users, including: (1) *Login bonus*: a reward given to users for logging into DApps daily. For example, DApp BingoBet offers users one free lucky draw every day. Users can get free EOS from 0.0005 to 50.0. (2) *Free tokens*: some DApps hand out tokens to attract active users. For example, PRA Candybox offers each user a limited chance to get free tokens. In general, each account can only get a limited number of tokens, so by creating multiple bots, a controller can make more profit. (3) *Invitation rewards*: some DApps give rewards for inviting friends. For example, Hash Baby offers 5% of its total tokens to those who invite friends. Thus, bots could create a large number of accounts and invite them in order to gain profit.

Although we have manually identified many bonus hunters, it is non-trivial to automatically identify them. Here, we propose a semi-automated but effective approach. One feature of bonus hunter bots is that they invoke some specific contracts with a high frequency, or they will invite other “similar” accounts to invoke the contracts. Thus, we first rank the bot-like accounts by their frequency of contract invocations, and then check the most popular DApps they invoke. For the top DApps, we then check to confirm whether they offer the aforementioned incentives. As a heuristic, we assume that bot communities targeting these apps are likely bonus hunters. To gain an upper estimate of how many bot communities are bonus hunters, we compute the proportion of money hunting actions for each account. We consider anything above 50% to be indicative of a likely

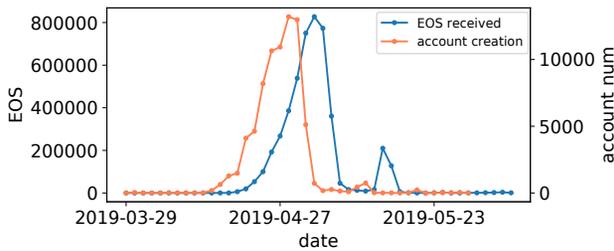


Fig. 12. An example of the click fraud (EOS Global).

bonus hunter. Through this, we identify 162,147 accounts belonging to this category. It accounts for 35.79% of all bot-like accounts. These accounts contribute the second largest number of invocations (56.42%).

**5.6.2 Click Fraud.** The second largest use of bots is for *click fraud*. These accounts create “fake” traffic in order to promote the ranking of certain DApps. One major characteristic is that *the amount of EOS they transfer to DApps is almost identical to what they receive from the same DApps*. We also observe that these accounts are quite ephemeral, i.e., they become active for a few days and then become “silent”. This results in the rank of the targeted DApp being boosted temporarily during the active days.

The most extreme case is the *EOS Global*, which ranks 3rd among all DApps on 2019.4.27 (on DappTotal). Using our EMFG, we find that account “egtradeadmin” (belonging to *EOS Global*), ranks #1 among all accounts using PageRank. However, the transaction behaviors of “egtradeadmin” are typically bot-like. To highlight this, Fig. 12 shows the money transfer and account creation behaviors of *EOS Global*. A significant peak can be witnessed between 2019.04.23 and 2019.04.30. The accounts contributing were flagged as bot-like accounts using our approach, as since their creation, they have transferred a great deal of EOS to the *EOS Global* DApp, and received almost identical EOS from *EOS Global* in return.

We define a heuristic to classify bot-like accounts in this category. We simply check their balance (transfer in vs. transfer out) to identify whether they are suspicious. Accounts with a similar ratio (over 95%) are deemed to be performing click fraud. Note that, we have filtered bot-like communities that have very few transfers (e.g., <10 EOS) to remove potential false positives. We identify 136, 670 bot-like accounts belonging to this category. They have performed 32, 606, 630 invocations, with an amount of 372, 613, 472 EOS.

**5.6.3 DApp Team Controlled Accounts.** We also observe that some bot-like accounts are controlled by DApp teams. By analyzing the behaviors of such accounts, we see that they are mainly used for: (1) *Debugging*: These bots are used to run automated product quality assurance tests within the DApp. For example, account “zhaojingdong” has created 2,001 bot-like accounts, and these accounts are used to invoke the debug interface of the smart contract “fastwinhold1” (fastwinhold1::debug), with 19,299 invocations. The “fastwinhold1” is a smart contract that belongs to *Dapp Fast*. (2) *Interacting with users*: These accounts are created to support user-based activities within DApps. There are also other cases where DApp Games cannot find enough real players (especially during cold start), so they generate bots to interact with real users. (3) *Control*: These bot accounts are mainly used to either vote for themselves or run DApps of their own. These accounts have the same public key with DApp accounts, e.g. “eosiomeetone”, “eosasia11111”. As the EOS inherits DPoS and only 21 block producers can be elected, these accounts are used to perform self-voting.

Table 7. A Summary of the Bot-like Accounts

Category	# Accounts (%)	# Invocations (%)	# Volume (EOS)
Bonus Hunter	162,147 (42.46%)	106,358,320 (56.42%)	58,299,990 (7.50%)
Click Fraud	136,670 (35.79%)	32,606,630 (17.30%)	372,613,472 (47.97%)
Dapp Team	18,828 (4.93%)	48,067,077 (25.5%)	324,244,893 (41.74%)
Account Seller	10,174 (2.66%)	328,787 (0.17%)	2,493,316 (0.32%)
Others	54,018 (14.15%)	1,141,162 (0.61%)	19,186,141 (2.47%)
Total	381,837	188,501,976	776,837,813

As we have collected the DApp info, Team Controlled Accounts can be verified easily. We have identified 18,828 accounts (4.93%) belonging to this category. Note that these accounts invoke smart contracts with high frequency (48,067,077 times in total), which represents 25.5% of the overall bot-like invocations.

**5.6.4 Account Sellers.** Account Sellers involve creating large numbers of accounts with names that are perceived to have potential future value, as names are unique in EOSIO and cannot be reused. Account Sellers therefore bid for account names from the EOS authority, and then sell them for a higher price. The public key of the accounts are identical and only after they are sold out will they change their keys. We identify these account sellers by manually examining the websites of each DApp, as they will advertise their ability to sell accounts. From this, we have identified 10,174 accounts in total.

**5.6.5 Others.** 14.15% of the bot-like accounts (54,018 accounts, 355 communities) remain uncategorized, as we cannot identify generalizable traits that might reveal their motivation. That said, we are nevertheless confident that they are bots. From this group, we manually sample 16,558 bot accounts (194 communities), and observe several clues that give weight to this confidence: (1) A large portion of these bot communities (33.47%) have confusingly similar names. For example, accounts created by *gotolab12345* are named things like *gotolabms221*, *gotolabms222*, while for normal accounts, we rarely see this. (2) The account creation time of most of the bot communities (66.53%) is extremely intense, making normal human involvement unlikely, e.g., account *staroverlord* created 980 accounts in a few hours.

## 6 MEASUREMENT OF SECURITY ISSUES

In this section, we study the security issues of EOSIO by investigating both the on-chain and off-chain data we collected.

### 6.1 eosio.code Permission Misuse

**6.1.1 Overview of Permission Mechanism.** EOSIO supports a permission system that can control access to contracts (see Section 2.3). Users can modify their account permission group and link each permission to different kinds of actions. For the default setting, only the *owner* key can sign the *updateauth* action to change the private key of an account, yet custom permissions are incredibly flexible and address numerous possible use cases when implemented. Most notably, the *eosio.code* permission is designed to enable contracts to execute inline actions. However, once a user grants this permission to a contract, it can call system contracts in the name of the user. This means that *the eosio.token permission can be assigned to transfer EOS tokens without notifying the user*. For example, this type of permission misuse has been discovered in a popular DApp named *EOS Fomo3D* (whose account name is *eosfoiowolfs*), which led to many users stop using it [34].

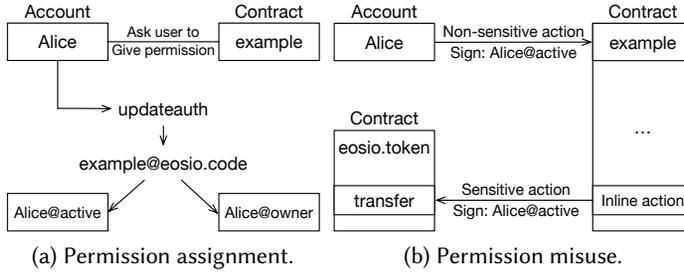


Fig. 13. Illustration of permission misuse.

**6.1.2 Overview of eosio.code Permission Misuse.** Fig.13 illustrates how *eosio.code* permission misuse occurs. First, a smart contract declares that the *eosio.code* permission is needed in the code. A user then grants *eosio.code* during the contract invocation, through another two permissions, i.e., *owner* and *active*. Specifically, users first link the *eosio.code* of one contract to the *owner/active* permission using *updateauth* action (see Figure 13a). By doing so, for every action signed by *owner/active*, the smart contract can invoke inline actions in the name of the user (see Figure 13b).

The above is typically used for convenience, e.g., allowing a game to automatically initiate transfers. However, in practice, the permission opens up a number of attacks. The community suggests that contract developers should get rid of *eosio.code* [21, 22], and recommends users to reserve their *eosio.code* permissions unless they trust that contract.

**6.1.3 Detecting eosio.code Permission Misuse.** As permission changes are related to action *updateauth*, we can track on-chain data to see whether there exists any permission misuses. Once we find an account grants the *eosio.code* permission to smart contract accounts with different public keys<sup>5</sup>, we flag this behavior as a *potential permission misuse*. As aforementioned in Section 2.3, there exists a threshold that must be reached to authorize the execution of the action. Thus, only when the *assigned permission weight* surpasses the *permission threshold* can another account invoke the corresponding permission. Otherwise, it may require several accounts to authorize a permission invocation together. As a result, for each permission grant, we further analyze the weight of the accounts who have been granted the *eosio.code*, and then compare it with the user-defined threshold, to identify the *permission misuse* finally.

**6.1.4 Results.** In total, we have collected 327, 287 *updateauth* actions. Among them, 26, 899 are associated with the permission *eosio.code*. After filtering out actions whose granted weight is lower than the threshold, we identify 18, 321 permission grant actions in total, and 72.25% of them are linked to *active* permission. Thus, we have identified 13,237 permission misuse in total, involving 5,541 user accounts, who grant their permissions to 407 smart contracts. The top 5 contracts that have been granted *eosio.code* with *active* are shown in Table 8. Figure 14 illustrates the permission misuse graph. Each node denotes an account, and each edge represents a permission grant action. We observe that most of the smart contracts that require sensitive permissions are gambling games (i.e., 4 in Top-5). Contract *fastwincpuem* is the most representative one. To play the game, over 1,400 users have to grant their *eosio.code* permissions to it, which pose serious security risks.

<sup>5</sup>Note that we compare the public keys of the involved accounts in order to filter the *valid* permission grants for users with multiple accounts.

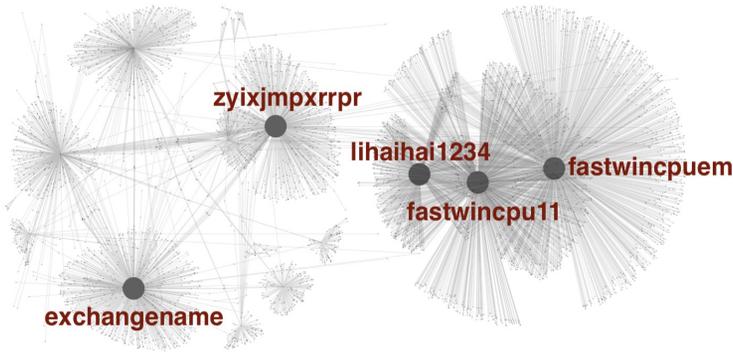


Fig. 14. Visualization of the permission grant actions.

Table 8. Top Accounts with Granted eosio.code Permission

linked permission	Account	# auth account	Identity
active	fastwincpuem	1, 453	gambling game
	fastwincpu1	1, 231	gambling game
	zyixjmpxrrpr	672	gambling game
	exchangenamename	626	Account exchange
	lihaihai1234	494	gambling game

### 6.2 Real-World Attacks

Although a number of security reports [14, 17, 23] have revealed attacks, it is still unknown how prevalent they are. By surveying 37 attack reports [29, 32], we next manually classify attacks into two main categories: *Unchecked Input Attacks* and *Predictable State Attacks*. Using these observations, we have implemented a monitoring system to identify suspicious actions and accounts.

**6.2.1 Unchecked Input Attacks.** We identify 4 samples of attacks that exploit unchecked inputs [16]. These rely on victim contracts failing to properly validate the identities of received tokens or notifications [16]. For example, an attacker can create an EOS token and named it as “EOS”, and then transfers a number of such tokens to a vulnerable (game) contract that does not verify the issuer of the tokens. As a result, that attacker might be able to make profits after successfully executing the remaining logic of the contract according to certain rules. It has two variants: (1) **Fake EOS Transfer Attack**, which transfer fake EOS tokens to deceive the victim contract into believing that it is receiving EOS tokens [4]. (2) **Fake EOS Notice Attack**, which sends fake notifications to deceive the victim contract into believing that it is receiving EOS tokens [5].

We devise a pattern-based mechanism to detect these two types of attacks, by looking for the patterns summarized in Table 9 [4, 5]. ✓ means the pattern is applicable to detect that particular attack, whereas ✗ indicates that the pattern is not applicable. In addition, pattern with *Yes* means the condition must be met; alternatively, pattern with *No* has to be otherwise satisfied. In total we have 2 patterns for the *Fake EOS Transfer Attack* and 3 patterns for the *Fake EOS Notice Attack* respectively. Note that the aforementioned patterns may lead to false positives. For instance, the transfer token symbol can always be specified as “EOS” by the sender for any non-malicious (e.g., testing) purpose. Thus, we further use heuristic strategies to alleviate this issue. Specifically, for the *fake EOS transfer attack*, if an account *A* sends fake EOS tokens to account *B*, and on the same

Table 9. Patterns of Fake EOS Transfer/Notice Attacks

Pattern	Fake EOS Transfer	Fake EOS Notice
The transfer token symbol is EOS.	✓	✓
Whether the contract executes “transfer” action is the official contract (eosio.token) or not.	No	Yes
The notice receiver is the victim, rather than any accounts involved in the transfer.	✗	✓

day account A makes profit from B, we then mark A as suspicious. Similar to the *fake EOS notice attack*, we further verify all accounts that have sent fake EOS notice to a DApp’s account. These heuristics can identify all 4 unchecked input attacks.

**6.2.2 Predictable State Attacks.** We collect 33 samples of attacks that rely on predictable state. These include variations such as: (1) The **Random Number Attack**, which exploits the vulnerabilities present in the procedure to generate random numbers. (2) The **Roll Back Attack**, which defrauds the lottery game without actually paying the bet cost by rolling back the corresponding unsatisfied reversible transaction [30, 31]. (3) The **Transaction Congestion**, which launches a DoS attack by sending a large number of deferred transactions based on the predictable state, to force the victim contract to regenerate a different result [19]. This kind of attack is usually targeted at gambling games (32 out of 33 attack gambling games).

Such attacks cannot be identified using the heuristics in Section 6.2.1. Hence, we turn to another intuition: *accounts with large profits in a short period of time with high frequency may be suspicious*. This assumption, however, could lead to many false positives. To investigate this, we manually analyze the collected attack accounts and summarize several characteristics:

- (1) *Although the profit attackers earn varies widely, their profitability ratio<sup>6</sup> is relatively high.* In contrast, even though a normal account has the potential to earn high profits (e.g., if its wins a large bet), it is unlikely that normal accounts consistently maintain a high profitability ratio. For the collected attacks, the median profit is 2, 202 EOS, and the median of the attacker’s profitability ratio is 2.4.
- (2) *Attacks usually happen over a relatively short time* (e.g., from a few minutes to an hour). This is quite different to other normal behaviors. As a reference, all the 33 Predictable State Attacks finish within an hour.
- (3) *The total amount of the “excessive” profit always makes-up a large portion of the total volume of the transactions for the account with the target DApp.* Because attack accounts are usually active for just a few days. Once attackers succeed, they usually do not use the Dapp anymore. In contrast, normal users with high profits have no reasons to be “silent” after that.

Based on these observations, we are able to propose a semi-automated approach to label suspicious attacker accounts, as follows:

**Step 1.** We first monitor the accounts (in the EMFG) that gain high profit with a high profitability ratio. For a given account, if the profit is larger than a threshold  $W1$  and the profitability ratio is larger than a threshold  $W2$  in some granularity of time period, we mark it as suspicious. In particular, we specify two kinds of granularity: *one day* for the coarse-grained detection and *an hour* for the fine-grained detection, to achieve a good accuracy with acceptable performance. Note that, we empirically set  $W1 = 400$ ,  $W2 = 1.2$ . This is a looser boundary (compared with known

<sup>6</sup>Here, we define the ratio as the amount of EOS received over the amount of EOS sent.

attacks) to identify as many suspicious accounts as possible. This step is capable of filtering out most of the unrelated normal accounts.

**Step 2.** For a given flagged victim DApp candidate, we record the profits and the corresponding dates for all suspicious accounts. For each suspicious account, we further calculate the proportion of profits (denoted by  $p$ ) over the total incoming EOS tokens received from that DApp. An account will be labelled as highly suspicious and require further verification if  $p$  is larger than a threshold  $W3$ , which gives the liveness of one account to some extent. We empirically set  $W3 = 0.9$ , i.e., the excessive profits must have occupied at least 90% of all the money transfers. As a reference, for the collected attack accounts, all of them have achieved 100% on this metric.

**Step 3.** The aforementioned *profit-driven* approaches are able to identify most of the attacks, as the sole purpose of the attacks we considered is to make profit. However, it is quite possible that we may include some benign behaviors, due to the limitation of the thresholds we defined. Thus, to filter out possible false positives, our last step is to manually analyze the remaining suspicious attack accounts and their behaviors. Note that, previous steps have filtered most of the accounts and actions, making our manual analysis possible. *Our goal is to either replay the “attacks” they performed, or find more evidences, to confirm whether a suspicious account fulfills the characteristics of the attacks..*

To replay the attacks, one can use the EOSIO official testnet or build a local testnet. The former is straightforward, whereas the latter requires more efforts to build the environment. We prefer the latter because it provides the ability to perform (and customize if necessary) deeper analysis. To the best of our knowledge, this is the only reliable way to manually label attacks. We employ three techniques here. (1) Random number attacks can be reproduced on the testnet we built. Thus, for each suspicious account, we first analyze whether it targets a gambling DApp. Then we repeat its actions to see whether we are able to launch random number attacks. We have implemented Proof-of-Concept (PoC) scripts on our testnet to verify the random number attacks are indeed performed. (2) For the roll back attack, we take advantage of the EOSIO client we customized to synchronize with the mainnet. Although roll back actions will not be put on the chain, they will be broadcast and a client could receive them. For each account, we further analyze their broadcast information. If one account has too many roll back transactions with a high profit, we believe it launches roll back attacks. (3) The transaction congestion attack is quite noticeable, as attackers usually create thousands of deferred transactions and may even paralyze the whole mainnet. Thus, for each account, we analyze its deferred transactions to see whether it performs *DDoS-like* attack. In brief, the manual investigation allows us to determine the attacks, the corresponding attack accounts, and the victim DApps.

**6.2.3 Detection Result.** Using the above detection mechanisms, we identify the presence of attacks across the entire EOSIO blockchain. We discover 301 attack accounts associating with 1, 518, 401 EOS tokens (roughly \$4.8 million). Table 10 lists the top-3 of them. Specifically, we discover 24 attack accounts associated with 136, 881 EOS tokens for the *fake EOS transfer attacks*, and 28 accounts associating with 235, 867 EOS tokens for the *fake EOS notification attacks* respectively<sup>7</sup>. Furthermore, we find 251 attack accounts (1, 070, 005 EOS tokens) that rely on predictable state.

These attacks are targeting 112 victim DApps. Note that, we report all the identified attacks to the corresponding DApp teams. By the time of our study, 80 attack accounts have been confirmed by them, causing 828, 824 EOS tokens in losses (roughly \$2.6 million). For the remaining attack accounts being detected, We are still working with DApp developers to make final confirmation.<sup>8</sup> Apart from notifying the DApp developers for a timely damage control, we also assist them in

<sup>7</sup>Two accounts are related to both fake EOS transfer and fake EOS notification attacks.

<sup>8</sup>Some DApp teams are no longer active due to the financial loss of attacks.

Table 10. Top-3 Attack Accounts Identified

attackers	target	profit	confirmed
hnhlpyadbunv	Dapp Dice	195, 530	Yes
ilovedice123	EOSBet	137, 970	Yes
eykkxszdrnnc	EOS Max, BigGame	71, 731	Yes

tracing the losses and provide a digital forensic service to support advanced collaboration with other third parties like exchanges. For example, we have assisted the development teams of DApp BetDice, ToBet, EOS MAX, BigGame in **tracing the losses of 288,329 EOS** by providing them attack evidence.

## 7 LIMITATIONS AND IMPLICATIONS

### 7.1 Limitations

Our study carries some limitations. In several cases, we have relied on heuristics and manual validation. This was necessary due to the paucity of ground-truth data. For example, for the bot-like account detection, we have tried our best to validate our approach, however, it is impossible for us to make sure *all* accounts flagged are truly bots. Similarly, although we only focus on bot-like accounts that operate in groups and have similar behaviors, it is possible that there are bots that do not fall into this definition. Moreover, we did not precisely identify the boundary between legitimate bots and fraudulent bots. Instead, we examine the purposes of the bot accounts, and classify them into five main categories. In general, one would regard the *bonus hunter* bots and *click fraud* bots as fraudulent bots, as they were created explicitly with the purpose to steal currency or manipulate the ranking of DApps. Nevertheless, more precise and fine-grained techniques can be applied to distinguish between legitimate and fraudulent bots. We raise similar observations regarding our attack detection. We have applied heuristics to flag accounts that are suspicious, and relied on manual efforts to confirm them. This, of course, might not be scalable and could mean that we only offer a lower-bound. Future work will involve incorporating program analysis and dynamic testing techniques to help us automatically identify attacks.

### 7.2 Implications

Our observations are of key importance to stakeholders in the community. First, considering the large number of fraudulent behaviors and security issues we discover, the governance of EOSIO needs to be improved. Second, we argue DApp developers should take actions immediately to address the security issues introduced by vulnerabilities of smart contracts. Third, as mentioned earlier, we provide a digital forensic service to help developers recoup the losses by collaborating with other third parties like exchanges. As of this writing, we have helped four DApp teams in tracing the losses by providing them with attack evidence. Moreover, our findings could help users and investigators to understand the status quo of EOSIO ecosystem, and protect them from being deceived by some “popular” DApps. Last but not least, our findings and techniques could be generalized to other blockchain platforms as well. For example, it is reported that blockchain bots were also found in the TRON DApp ecosystem [35]. Our observations and techniques could be easily adopted to detect them.

## 8 RELATED WORK

### 8.1 EOSIO Analysis

As EOSIO was launched in June 2018, there are only a few studies looking at the EOSIO ecosystem. Lee *et al.* [56] introduced and studied four attacks stemming from the unique design of EOSIO. Quan *et al.* [63] focused on detecting two types of vulnerability for EOSIO smart contracts. He *et al.* [52] proposed a symbolic execution based static analysis framework to identify four kinds of vulnerabilities in EOSIO smart contracts. Lee *et al.* [55] found that attackers can undermine the fairness of compensation policies by manipulating the production schedule of EOSIO. They also discussed the applicability of attacks against the real EOSIO mainnet. There are also reports [1, 4, 5, 19, 30, 31] released by security companies. However, all of the existing studies were focused on security issues introduced by the EOSIO framework and smart contracts, none of them have studied security issues from the perspective of transaction-based analysis. Our paper is the first one that analyzes the EOSIO blockchain ecosystem at scale, longitudinally and across various dimensions.

### 8.2 Transaction-based Analysis of Blockchain

Previous studies have investigated other blockchain systems by performing transaction-based analyses. Several works focus on Bitcoin [40, 48, 59, 62, 64, 65, 73], including de-anonymization and money laundering detection, by using graph-based approaches. Researchers have also investigated Ethereum by using transaction-based analyses [43, 45]. For example, Chen *et al.* [45] presents a graph-based analysis of Ethereum, which is somewhat similar to our approach in Section 4. However, our work differs from previous studies. First, EOSIO is totally different to other blockchain systems, from the underlying consensus protocol, to the account system, the permission management mechanism and the smart contracts. We have compared the graph metrics of EOSIO and Ethereum, and our observations reveal the unique characteristics of EOSIO ecosystem (see Section 4). Second, we have proposed systematic approaches to successfully identify and measure bot-like and malicious accounts, while such accounts have never been identified in other platforms before. To the best of our knowledge, this is the first comprehensive work to study transaction behaviors on EOSIO.

### 8.3 Vulnerability/Attack Detection of Blockchain

A number of studies have focused on detecting the vulnerabilities in Ethereum smart contracts [53, 57, 58, 66–69]. For example, Luu *et al.* [58] proposed Oyente, a symbolic execution tool which analyses Ethereum smart contracts to detect bugs. Kalra *et al.* [66] proposed ZEUS, a framework for formal verification of smart contracts using abstract interpretation and symbolic model checking. He *et al.* [51] analyzed the copy-paste vulnerabilities introduced by code clones in smart contracts. Most of the vulnerability detection approaches were performed based on static analysis. In contrast, we adopt a statistical approach to detect attacks by analyzing the anomalies.

### 8.4 Blockchain Scam Detection

Blockchain platforms have been the target of scams. A few studies have characterized these scams. Most of them were focused on detecting Ponzi schemes [39, 46]. For example, Bartoletti *et al.* [39] summarized a set of criteria for determining whether a smart contract implements a Ponzi scheme, and they have identified four kinds of Ponzi schemes in Ethereum. Chen *et al.* [46] proposed a machine learning approach to identify Ponzi Schemes. They first extract features from user accounts and operation codes of the smart contracts and then build a classification model. A few studies have also characterized scam domains in the blockchain ecosystem. For example, Xia *et al.* [72] have studied scam domains and fake mobile apps that target cryptocurrency exchanges. They have identified over 1,500 scam domains and reveal that these scams have incurred financial losses of

\$520k at least. Although existing studies have characterized scams in the blockchain systems, no previous paper has characterized the bot accounts in the EOSIO blockchain ecosystem.

### 8.5 Bot Account and Sybil Account Detection

Bot accounts and Sybil accounts are pervasive in today's online communities. A large number of studies have proposed to identify bot accounts and sybil accounts, mainly in the domain of social networking systems (e.g., Twitter, Facebook, Instagram, and LinkedIn, etc.) [42, 44, 50, 70, 71]. For example, Wang *et al.* [71] proposed to detect sybil accounts based on server-side clickstream models (the traces of click-through events generated by online users), as sybils and real users have very different goals in their usage of online services. Thus, they build a similarity graph that captures distances between clickstream sequences and then apply clustering to identify groups of user behavior patterns. Similarly, Cao *et al.* [42] proposed to cluster user accounts according to the similarity of their actions and uncovered large groups of malicious accounts. Chavoshi *et al.* [44] proposed DeBot, an unsupervised tool to detect bots on twitter based on the key observation that highly synchronous user accounts are most likely bots. Cao *et al.* [41] proposed SybilRank, a tool that relies on social graph properties to rank accounts based on their perceived likelihood of being Sybils. Almost all the aforementioned studies were focused on social networks. Blockchain bots, however, have not been well studied in our community. Accounts in the blockchain systems have specialized behaviors compared with social accounts, e.g., money transfer and contract invocations. To the best of our knowledge, this is the first research paper that proposed to detect bots in the blockchain systems. Nevertheless, we admit that more advanced approaches in the social network domain could be exploited to identify blockchain bots in the future work.

## 9 CONCLUSION

In this paper, we have performed the first large-scale measurement study of the EOSIO blockchain. By constructing a comprehensive dataset, we first analyzed the activities including money transfer, account creation and contract invocation. We further focused on security and fraudulent issues, including bot-like accounts and attack detection. Our exploration has identified many security issues and revealed various interesting observations, including thousands of bot accounts, hundreds of real-world attacks, as well as insights for future research directions.

## ACKNOWLEDGMENT

We sincerely thank our shepherd Prof. Michael Sirivianos (Cyprus University of Technology) and all the anonymous reviewers for their valuable suggestions and comments to improve this paper. This work is supported by the National Natural Science Foundation of China (grants No.61702045 and No.61725201), Hong Kong RGC Project (No. 152193/19E), and Beijing Outstanding Young Scientist Program BJJWZYJH01201910001004.

## REFERENCES

- [1] 2018. Defeating EOS Gambling Games: The Tech Behind Random Number Loophole. <https://medium.com/@peckshield/defeating-eos-gambling-games-the-tech-behind-random-number-loophole-cf701c616dc0>.
- [2] 2018. EOSIO Dawn 3.0 Now Available. <https://medium.com/eosio/eosio-dawn-3-0-now-available-49a3b99242d7>.
- [3] 2018. EOS's Gloom: Real Users Account for 30% and 8 Million Yuan Lost to Hackers in Last Six Months. <https://news.8btc.com/eoss-gloom-real-users-account-for-30-and-8-million-yuan-lost-to-hackers-in-last-six-months>.
- [4] 2018. "Fake EOS Attack" Upgraded, 60K EOS Tokens Lost by EOSCast. <https://blog.peckshield.com/2018/11/02/eos/>.
- [5] 2018. "Fake Transfer Notice" Loophole Details Explained, 140K EOS Tokens Lost by EOSBet. <https://blog.peckshield.com/2018/10/26/eos/>.
- [6] 2018. FIBOS weekly. <https://developpaper.com/fibos-weekly/>.

- [7] 2018. Hacker created 2190 accounts to circumvent ECAF (in Chinese). <https://www.myoschain.com/blog/134430038970859522>.
- [8] 2019. API Endpoints. <https://www.eosdocs.io/resources/apiendpoints/>.
- [9] 2019. Bots drove nearly 40% of internet traffic last year. <https://thenextweb.com/security/2019/04/17/bots-drove-nearly-40-of-internet-traffic-last-year-and-the-naughty-ones-are-getting-smarter/>.
- [10] 2019. Bots Index. <https://github.com/hashbaby-com/eos-hall-of-shame/tree/master/bots>.
- [11] 2019. Clustering coefficient. [https://en.wikipedia.org/wiki/Clustering\\_coefficient](https://en.wikipedia.org/wiki/Clustering_coefficient).
- [12] 2019. DAppReview. <https://www.dapp.review/>.
- [13] 2019. DAppTotal. <https://dapptotal.com/>.
- [14] 2019. EOS DApps Lose Almost \$1 Million to Hackers Over the Last Five Months. <https://cointelegraph.com/news/eos-dapps-lose-almost-1-million-to-hackers-over-the-last-five-months>.
- [15] 2019. EOS Developer Documentation. <https://developers.eos.io/eosio-nodes/docs>.
- [16] 2019. EOS Development Tutorials. <https://github.com/peckshield/EOS/tree/master/eos-tutorials>.
- [17] 2019. EOS news update: 2.09 million EOS disappears in a hack attack – EOS accounts blocked by Houbi.
- [18] 2019. EOS: porn blowing up transaction volumes? <https://en.cryptonomist.ch/2019/09/03/eos-porn-transaction-volumes/>.
- [19] 2019. EOS “Transaction Congestion Attack”: Attackers Could Paralyze EOS Network with Minimal Cost. [https://blog.peckshield.com/2019/01/15/eos\\_CVE-2019-6199/](https://blog.peckshield.com/2019/01/15/eos_CVE-2019-6199/).
- [20] 2019. EOSIO Official Portal. <https://eos.io/>.
- [21] 2019. EOSIO Permission Grant. <https://blog.csdn.net/zhuxiangzhidi/article/details/81635688>.
- [22] 2019. EOSIO Secure Coding. <https://github.com/peckshield/EOS/blob/master/eos-tutorials/README.md>.
- [23] 2019. EOS/USD market drops by 4% following \$7.7 million EOS hack attack. <https://www.fxstreet.com/cryptocurrencies/news/eos-usd-market-drops-by-4-following-77-million-eos-hack-attack-201902262151>.
- [24] 2019. Libra Core implements a decentralized, programmable database which provides a financial infrastructure that can empower billions of people. <https://github.com/libra/libra>.
- [25] 2019. Official Bitcoin Portal. <https://bitcoin.org/en/>.
- [26] 2019. Official Ethereum Portal. <https://www.ethereum.org/>.
- [27] 2019. Our AI Detects Your AI – Revealing the Secret Blockchain DApp World of Bots (Part 1 – EOS). <https://medium.com/@AnChain.AI/our-ai-detects-your-ai-revealing-the-secret-blockchain-dapp-world-of-bots-eed8884a07>.
- [28] 2019. Pearson correlation coefficient. [https://en.wikipedia.org/wiki/Pearson\\_correlation\\_coefficient](https://en.wikipedia.org/wiki/Pearson_correlation_coefficient).
- [29] 2019. PeckShield Official Portal. <https://www.peckshield.com/home.html?lang=en>.
- [30] 2019. Roll Back Attack about blacklist in EOS. <https://medium.com/@slowmist/roll-back-attack-about-blacklist-in-eos-adf53edd8d69>.
- [31] 2019. Roll Back Attack about replay in EOS. <https://medium.com/@slowmist/roll-back-attack-about-replay-in-eos-acddee979396>.
- [32] 2019. SlowMist Official Portal. <https://www.slowmist.com/en/index.html>.
- [33] 2019. Study: 75% of EOS Dapp Transactions Are Now Made By Bots. <https://www.coindesk.com/study-75-of-dapp-transactions-are-now-made-by-bots>.
- [34] 2019. The Security Issues of EOSIO.Code Permission for EOS Wolf. <https://bihu.com/article/992656>.
- [35] 2019. TRON Plagued By Infestation Of dApp Bots. <https://cryptobriefing.com/tron-plagued-by-infestation-of-dapp-bots-anchain-report/>.
- [36] 2020. Accounts and Permissions. [https://developers.eos.io/welcome/latest/protocol/accounts\\_and\\_permissions](https://developers.eos.io/welcome/latest/protocol/accounts_and_permissions).
- [37] 2020. Glossary of EOSIO. <https://developers.eos.io/welcome/latest/glossary/index>.
- [38] 2020. History of Histories. [https://eos.discussions.app/tag/voice/3i4rwgpi8cqal/dan\\_larimer\\_history\\_of\\_histories](https://eos.discussions.app/tag/voice/3i4rwgpi8cqal/dan_larimer_history_of_histories).
- [39] 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.
- [40] Alex Biryukov, Dmitry Khovratovich, and Ivan Pustogarov. 2014. Deanonymisation of clients in Bitcoin P2P network. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 15–29.
- [41] Qiang Cao, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. 2012. Aiding the Detection of Fake Accounts in Large Scale Social Online Services. In *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation (NSDI’12)*. USENIX Association, USA, 15.
- [42] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. 2014. Uncovering Large Groups of Active Malicious Accounts in Online Social Networks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 477–488.
- [43] Wren Chan and Aspen Olmsted. 2017. Ethereum transaction graph analysis. In *12th International Conference for Internet Technology and Secured Transactions (ICITST)*. 498–500.

- [44] N. Chavoshi, H. Hamooni, and A. Mueen. 2016. DeBot: Twitter Bot Detection via Warped Correlation. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. 817–822.
- [45] Ting Chen, Yuxiao Zhu, Zihao Li, Jiachi Chen, Xiaoqi Li, Xiapu Luo, Xiaodong Lin, and Xiaosong Zhang. 2018. Understanding Ethereum via Graph Analysis. In *IEEE International Conference on Computer Communications (INFOCOM)*. 1484–1492.
- [46] Weili Chen, Zibin Zheng, Jiahui Cui, Edith Ngai, Peilin Zheng, and Yuren Zhou. 2018. Detecting ponzi schemes on ethereum: Towards healthier blockchain technology. In *Proceedings of the 2018 World Wide Web Conference (WWW '18)*. 1409–1418.
- [47] Giorgio Fagiolo. 2007. Clustering in complex directed networks. *Physical Review E* 76, 2 (2007), 026107.
- [48] Michael Fleder, Michael S. Kester, and Sudeep Pillai. 2015. Bitcoin Transaction Graph Analysis. arXivpreprintarXiv:1502.01657
- [49] Zafar Gilani, Jon Crowcroft, Reza Farahbakhsh, and Gareth Tyson. 2017. The implications of twitterbot generated data traffic on networked systems. In *Proceedings of the SIGCOMM Posters and Demos*. 51–53.
- [50] Zafar Gilani, Reza Farahbakhsh, Gareth Tyson, and Jon Crowcroft. 2019. A large-scale behavioural analysis of bots and humans on twitter. *ACM Transactions on the Web (TWEB)* 13, 1 (2019), 1–23.
- [51] Ningyu He, Lei Wu, Haoyu Wang, Yao Guo, and Xuxian Jiang. 2020. Characterizing code clones in the Ethereum smart contract ecosystem. In *Twenty-Fourth International Conference on Financial Cryptography and Data Security (FC '20)*.
- [52] Ningyu He, Ruiyi Zhang, Lei Wu, Haoyu Wang, Xiapu Luo, Yao Guo, Ting Yu, and Xuxian Jiang. 2020. Security Analysis of EOSIO Smart Contracts. *arXiv preprint arXiv:2003.06568* (2020).
- [53] Bo Jiang, Ye Liu, and WK Chan. 2018. Contractfuzzer: Fuzzing smart contracts for vulnerability detection. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE)*. ACM, 259–269.
- [54] Alan Kvanli, Robert Pavur, and Kellie Keeling. 2005. *Concise managerial statistics*. Cengage Learning. 81–82 pages.
- [55] Dongsoo Lee and Dong Hoon Lee. 2019. Push and Pull: Manipulating a Production Schedule and Maximizing Rewards on the EOSIO Blockchain. In *Proceedings of the Third ACM Workshop on Blockchains, Cryptocurrencies and Contracts (BCC '19)*. 11–21.
- [56] Sangsup Lee, Daejun Kim, Dongkwan Kim, Sooel Son, and Yongdae Kim. 2019. Who Spent My {EOS}? On the (In) Security of Resource Management of EOS. IO. In *13th {USENIX} Workshop on Offensive Technologies ({WOOT} '19)*.
- [57] Chao Liu, Han Liu, Zhao Cao, Zhong Chen, Bangdao Chen, and Bill Roscoe. 2018. ReGuard: finding reentrancy bugs in smart contracts. In *Proceedings of the 40th International Conference on Software Engineering (ICSE-C)*. 65–68.
- [58] Loi Luu, Duc-Hiep Chu, Hrishi Olickel, Prateek Saxena, and Aquinas Hobor. 2016. Making Smart Contracts Smarter. In *ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. 254–269.
- [59] Damiano Di Francesco Maesa, Andrea Marino, and Laura Ricci. 2016. An analysis of the Bitcoin users graph: inferring unusual behaviours. In *International Workshop on Complex Networks and their Applications*. 749–760.
- [60] Mark EJ Newman. 2003. Mixing patterns in networks. *Physical Review E* 67, 2 (2003), 026126.
- [61] Rogier Noldus and Piet Van Mieghem. 2015. Assortativity in complex networks. *Journal of Complex Networks* 3, 4 (2015), 507–542.
- [62] Silivanxay Phetsouvanh, Frédérique Oggier, and Anwitaman Datta. 2018. EGRET: Extortion Graph Exploration Techniques in the Bitcoin Network. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. 244–251.
- [63] Lijin Quan, Lei Wu, and Haoyu Wang. 2019. EVulHunter: Detecting Fake Transfer Vulnerabilities for EOSIO's Smart Contracts at Webassembly-level. arXivpreprintarXiv:1906.10362
- [64] Fergal Reid and Martin Harrigan. 2011. An Analysis of Anonymity in the Bitcoin System. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*. IEEE, 1318–1326.
- [65] Dorit Ron and Adi Shamir. 2013. Quantitative Analysis of the Full Bitcoin Transaction Graph. In *International Conference on Financial Cryptography and Data Security (FC)*. 6–24.
- [66] Sukrit SKalra, Seep Goel, Mohan Dhawan, and Subodh Sharma. 2018. ZEUS: Analyzing Safety of Smart Contracts. In *Network and Distributed Systems Security Symposium (NDSS)*. 1–12.
- [67] Sergei Tikhomirov, Ekaterina Voskresenskaya, Ivan Ivanitskiy, Ramil Takhaviev, Evgeny Marchenko, and Yaroslav Alexandrov. 2018. Smartcheck: Static analysis of ethereum smart contracts. In *Proceedings of the 1st International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB '18)*. IEEE, 9–16.
- [68] Christof Ferreira Torres, Julian Schütte, and Radu State. 2018. Osiris: Hunting for Integer Bugs in Ethereum Smart Contracts. In *The 34th Annual Computer Security Applications Conference (ACSAC '18)*. 664–676.
- [69] Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Bünzli, and Martin Vechev. 2018. Securify: Practical Security Analysis of Smart Contracts. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*.

- [70] Onur Varol, Emilio Ferrara, Clayton A Davis, Filippo Menczer, and Alessandro Flammini. 2017. Online human-bot interactions: Detection, estimation, and characterization. In *Eleventh international AAAI conference on web and social media*.
- [71] Gang Wang, Tristan Konolige, Christo Wilson, Xiao Wang, Haitao Zheng, and Ben Y. Zhao. 2013. You Are How You Click: Clickstream Analysis for Sybil Detection. In *22nd USENIX Security Symposium (USENIX Security 13)*. 241–256.
- [72] Pengcheng Xia, Bowen Zhang, Ru Ji, Bingyu Gao, Lei Wu, Xiapu Luo, Haoyu Wang, and Guoai Xu. 2020. Characterizing Cryptocurrency Exchange Scams. *arXiv preprint arXiv:2003.07314* (2020).
- [73] Chen Zhao and Yong Guan. 2015. A Graph-based investigation of Bitcoin transactions. In *11th IFIP International Conference on Digital Forensics (DF)*. 79–95.

Received January 2020; revised February 2020; accepted March 2020