

A Detailed Background

In this section, we explain the terminologies related to Ethereum blockchain, transactions and NFTs, etc.

A.1 Blockchain and Ethereum

Blockchain, a distributed ledger technology, has drawn continuous attention recent years. Blockchains are made up of securely linked blocks with cryptography techniques [53], where each block contains information of the previous block (e.g., cryptographic hash). Then, consensus algorithms or protocols are applied to validate transactions and keep them being consistent. By this way, the blockchain transactions are immutable, traceable and publicly available. Ethereum is a decentralized, programmable blockchain, which means users can construct various decentralized applications on the blockchain. Ether (ETH) is the native cryptocurrency of Ethereum, and every transaction incurred in the Ethereum needs a specific fee paid in ETH. According to the market capitalization, ETH is the second-largest cryptocurrency behind Bitcoin.

A.2 Account and Transaction

Ethereum account is the key to access and explore the Ethereum ecosystem. Every Ethereum account is associated with a unique address, akin to an email address for an inbox. The address can be used to receive or send funds to the corresponding account. Accounts in the Ethereum can be classified into two types: (1) Externally-owned account (EOA) and (2) Contract account. All the accounts are denoted as a 64 character hex string.

Transactions are messages sent from one account to another account. One of the simplest transaction is transferring ETH from one account to another, which will change the state of the Ethereum Virtual Machine (EVM) and need to be broadcast to the whole Ethereum network. Each transaction requires amount of fees to pay for the computation. The key information in a transaction includes the receive address, the sender address, value, data, gas limit, and the max fee per gas, etc. There are three categories of transactions within the Ethereum: 1) regular transactions, which indicates the transactions between two externally-owned accounts; 2) contract deployment transactions, which are special transactions without receive addresses; 3) execution of a contract, whose receive address is the smart contract address. When the transaction is submitted, its life cycle can be simplified into the following three steps: 1) an externally-owned account sends a transaction and generates a transaction hash; 2) the transaction is broadcast across the network; 3) a validator verifies the transaction and includes it into a block. Once the transaction is successfully executed, it can never be altered.

A.3 Smart Contract and Non-Fungible Token

Smart contract is an important feature in the Ethereum blockchain [88], which is a computer program that runs on the Ethereum to automatically execute or control relevant events and actions according to its logic. Smart contracts have largely reduced the requirements for trusted intermediaries, fraud losses and arbitration costs, etc. As we have mentioned before, smart contracts also belong to a type of Ethereum account, and it can interact with user accounts. Moreover, anyone can program a smart contract and deploy it to Ethereum, as long as the code is complied successfully and can be executed by the EVM. Smart contracts are the fundamental building blocks for various applications like decentralized finance (DeFi) and game finance (GameFi).

Non-Fungible Token (NFT) is one of the most successful applications on Ethereum. NFTs are tokens that can be used to represent the ownership of any unique asset such as image, video and audio. Different from fungible items where one dollar is exchangeable for another one dollar, NFTs are not interchangeable with each other, since they all have unique properties and are not divisible. Smart contracts manage the ownership and the transferability of NFTs. Specifically, when NFTs are minted or transferred, it triggers the code stored in the smart contract, and then the relevant actions are executed. Each NFT token will have an owner after mint, and this information can be easily verified in Ethereum. The NFTs can be bought and sold on any NFT market like OpenSea, LooksRare or X2Y2. More recently, the NFT and DeFi have been combined to form a number of interesting applications including NFT-backed loans, fractional ownership and certificates of authenticity.

B Related Work

In this section, we present the related work on graph analysis.

Analyses of Social Networks, Citation Networks and the Internet. There exist lots of prior works focusing on analysing social networks and citation networks like Flickr, Yahoo! 360 and LiveJournal [5, 38, 39, 41]. These studies investigate the graph properties including density, degree distribution and clustering coefficient, etc. Among them, [17] observed that the node degrees followed a power-law distribution in most real-world networks. [9] studied the graphs from the perspective of connectivity, where large strongly connected components (i.e., SCC) widely existed in the graphs. [23] classified social network's links into strong ties and weak ties, with strong ties indicating tighter clustering. [77, 37] explained the social networks' small-world phenomenon. [41] showed that the citation graphs demonstrated denser densities and decreasing diameters as time goes by. Then, a forest-fire graph generation model was proposed to simulate these phenomena. [64] proposed a jellyfish model to describe the topology of the Internet, which abstracted the structure in a human understandable way.

To summarize, the major findings of existing works are as follows: 1) power law degree distribution, where some nodes exhibit significantly large degrees; 2) preferential attachment growth model, where the likelihood of a new node establishing a connection with an existing node is directly tied to the degree of the existing node; 3) density of the graphs follows a rapid decline, and then becomes steady. For surveys of graph analysis, interested readers can refer to [27, 76]. Although the graph analyses are extensively studied in those networks, it is not clear whether these findings are still valid in this emerging NFT transaction network.

Graph Analyses of Cryptocurrency Transaction Networks. Due to the decentralized nature of blockchain, this makes it possible for everyone to access all the transaction information. Several recent works have studied the properties of Bitcoin and other cryptocurrency transaction networks. For instance, [26, 59, 48] studied the user behaviors in Bitcoin transaction network. [66] classified and visualized the information extracted from the Bitcoin network. [80] and [1] forecast the price of BTC via modeling the local topological structure of the Bitcoin graph. Apart from Bitcoin, other cryptocurrencies like Zcash [34], EOS [29] and Monero [51] had also been conducted similar analyses. [8] analyzed the transaction linkability in Zcash, which revealed the underlying privacy concerns. [16, 2] discussed the key factors that impact the scalability of various blockchain systems. [24] identified arbitrage behaviors among multiple cryptocurrency exchange markets (e.g., Kraken, Coinbase and Gemini) through weighted cycle detection. However, the majority of these works are performed on analyzing static graphs, whereas graphs are usually evolving over time in the real-world scenarios. Moreover, the aforementioned studies that analyze Bitcoin and other cryptocurrencies only involve transactions related to value or token transfers. Different from Bitcoin, recent popular blockchains like Ethereum and Solana support deploying smart contracts to provide diverse services, where human controlled accounts and program controlled agents coexist in the network, making the transaction network even more complicated. It is of great interest to us to investigate this type of transaction network.

Analyses of the Ethereum Blockchain. Given the possibility to access comprehensive information within the blockchain transaction network, some recent efforts follow the pioneer studies on social networks, citation networks and the Internet [39, 41, 64] to analyze the static Ethereum transaction network. Particularly, [40] measured four interaction networks to give new insights on the Ethereum graph properties. [18, 11] characterized major activities including money transfer and contract creation on Ethereum via graph analysis. [46] learned from Ethereum graph to perform price anomaly prediction. [86, 6] investigated the evolutionary dynamics of Ethereum activities through the lens of temporal graphs.

Instead of studying all the transactions in the Ethereum, [65] only considered transactions relevant to ERC20 tokens which were fungible tokens circulated in the Ethereum, and the results showed that they presented obvious social signals in the trading network. [12] performed a systematic analyses on the whole ERC20 token activities. [72] studied massive individual token networks from a graph analysis perspective, and found that they were largely dominated by a single hub and spoke pattern. Since tokens could be heavily influenced by various events like mint, burnt, transfer or staking, these aforementioned works only provide a general intuition about the structure of token distributions, the flow and spread of assets on the blockchain.

Similar graph analysis approaches have also been applied to Non-fungible tokens (NFTs), where NFT tokens are unique and not comparable with each other. For example, [78] used the Louvain algorithm [71] to extract the community based structures from the networks, which characterized the relationships between buyers and sellers. [52] showed that the NFT’s sale history and visual appearance were two good indicators to predict its price. [10] analyzed several popular NFT projects, and concluded that the structure of NFT networks was qualitatively similar to social networks. [73] quantified suspicious wash trading behaviors in NFT market via closed cycle detection in the network. Although NFTs play a crucial role in the Ethereum ecosystem, none of the aforementioned studies explore the temporal properties of the NFT transaction network from the temporal graph point of view.

C Basic Structure Properties

The following four properties are discussed in this section:

Assortativity characterizes the tendencies of nodes getting attached to similar nodes through a specific metric. Following [4], we calculate the degree assortativity as follows:

$$\alpha = \frac{\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} k_i k_j - [\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} \frac{1}{2}(k_i + k_j)]^2}{\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} \frac{1}{2}[k_i^2 + k_j^2] - [\frac{1}{|\mathcal{E}|} \sum_{(i,j) \in \mathcal{E}} \frac{1}{2}(k_i + k_j)]^2} \quad (1)$$

where k_i and k_j represent the degrees at the ends of edge $(i, j) \in \mathcal{E}$. $|\mathcal{E}|$ denotes the number of edges. The assortativity α lies in the range of $[-1, 1]$, where a positive α indicates that high degree nodes have high probabilities of linking to other nodes with high degrees on average. In contrast, when α is negative, it is a disassortative network and the high-degree nodes are more likely to link to low-degree nodes. More specifically, when $\alpha = 0$, we say the network is neutral, and neither the tendencies of linking to high-degree nor low-degree nodes are observed. Figure 4a shows that the assortativity of NFT transaction network is negative in the recent six years, and it increases year by year, gradually approaching to zero. This indicates that the emerging transaction network is evolving from disassortative to assortative, which means there are more and more hub nodes available for themselves to connect to each other to increase the assortativity. A detailed analysis is given in Section 4.2.

Density calculates the ratio of existing edges over the number of possible edges [49]. The density d is 1 for a complete network, and it can be larger than 1 when self-loops or multi-edges are taken into consideration. We compute the density of a directed network as follows:

$$d = \frac{|\mathcal{E}|}{|\mathcal{V}|(|\mathcal{V}| - 1)} \quad (2)$$

where $|\mathcal{E}|$ indicates the number of edges and $|\mathcal{V}|$ represents node numbers. As we can see in Figure 4b, the density drops rapidly, which indicates the network becomes sparser over time. The decrease of density is mainly caused by the increment of edges are less than the node increment. It also means that the network utilization is quite low and the interactions between different nodes are very limited (i.e., one node only interacts with a few other nodes). This is understandable, since creating an account (i.e., node) is free, but making transactions (i.e., creating edges) cost gas fees. This property is quite different from citation networks [41], where density becomes denser over the time.

Reciprocity in a directed network is determined by the proportion of bidirectional edges to the number of total edges [67, 3]. Formally, the reciprocity r is calculated as:

$$r = \frac{\sum \mathbb{I}((i, j) \in \mathcal{E} \wedge (j, i) \in \mathcal{E})}{|\mathcal{E}|} \quad (3)$$

where $|\mathcal{E}|$ indicates the number of edges. $\mathbb{I}(\cdot)$ is an indicator function and it returns 1 when node i and j have bidirectional edge, otherwise it returns 0. The trend of reciprocity is demonstrated in Figure 4c. The relation between reciprocity and time is not monotonous. In general, it increases at the first, and then decreases in the following several years. This may be due to the emerging of NFT swapping, which allows users to swap their NFTs with each other. However, as time goes on, the NFT market becomes more and more mature, and users are prone to trading instead of swapping their NFTs to gain more profits.

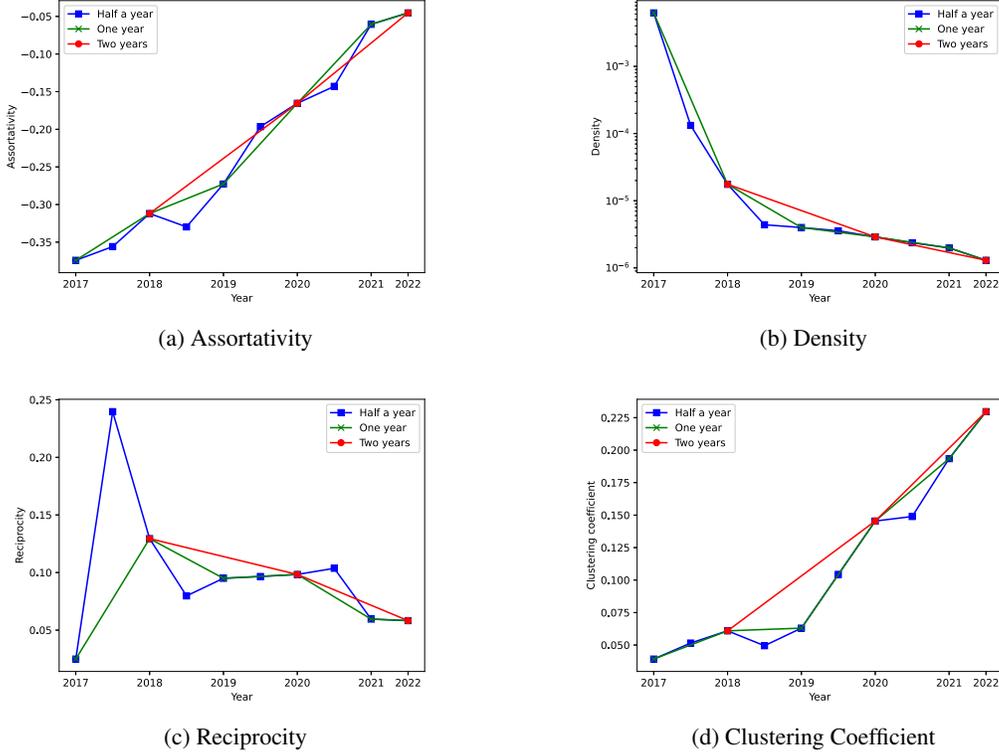


Figure 4: Evolution of network global properties with different time granularity.

Average Clustering Coefficient evaluates to what extent the nodes in a network tend to tightly cluster together [61], which is computed by averaging local clustering coefficient across all nodes. Node v_i 's local clustering coefficient c_i is the percentage of edges among its neighborhood divided by all the possible edges between its neighborhood. We formulate the local clustering coefficient for directed network as:

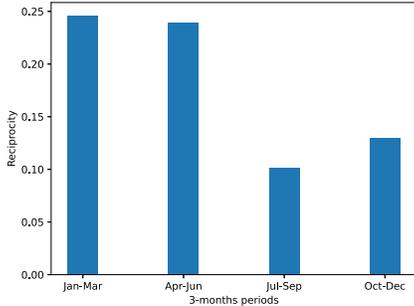
$$c_i = \frac{|\{e_{jk} : v_j, v_k \in \mathcal{N}_i, e_{jk} \in \mathcal{E}\}|}{|\mathcal{N}_i|(|\mathcal{N}_i| - 1)} \quad (4)$$

where $\mathcal{N}_i = \{v_j : e_{ij} \in \mathcal{E} \vee e_{ji} \in \mathcal{E}\}$ is the neighborhood of node i . Edge e_{jk} indicates the link between node j and node k . Thus, the equation for the average clustering coefficient is as follows:

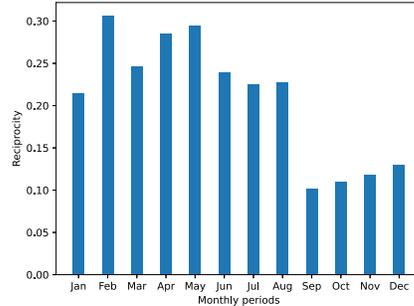
$$c = \frac{1}{|\mathcal{V}|} \sum_{i=1}^{|\mathcal{V}|} c_i \quad (5)$$

where $|\mathcal{V}|$ denotes the number of nodes. In Figure 4d, we present the clustering coefficient for the NFT transaction network across six years. We can observe that the average clustering coefficient is growing stably, which indicates the network is forming an increasing number of tightly connected clusters or communities. One possible explanation is that those accounts in a tightly connected clusters are actually controlled by the same person, and they utilize multiple accounts to conduct money laundering and wash trading [73], etc.

Furthermore, Figure 4 also illustrates the trends of different properties under different time granularities. In particular, the networks constructed with different time granularities could highlight the anomalies during its evolving procedure. Here, the anomaly means a property value that is much more larger or smaller than the value in its neighborhood time periods. For instance, the reciprocity in Figure 4c indicates that its value in the middle of 2018 is twice larger than that of the other time period. Since this abnormal value is observed in a half-year granularity, we can dive into a finer time scale, which is 3-month and monthly time scopes. The results are illustrated in Figure 5. As we can see in Figure 5a, the reciprocity values are significantly different between the first and the second

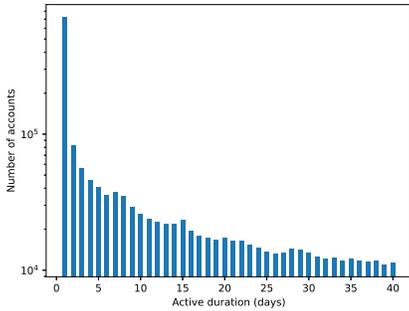


(a) 3-months granularity of year 2018

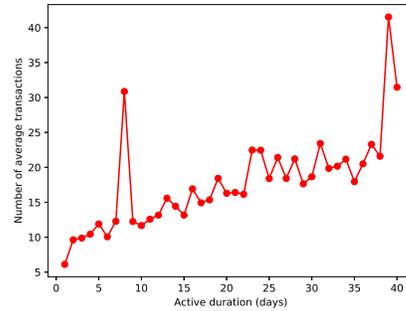


(b) Monthly granularity of year 2018

Figure 5: Finer time granularity analysis on reciprocity.



(a) Number of active accounts



(b) Number of average transactions

Figure 6: Accounts' active periods.

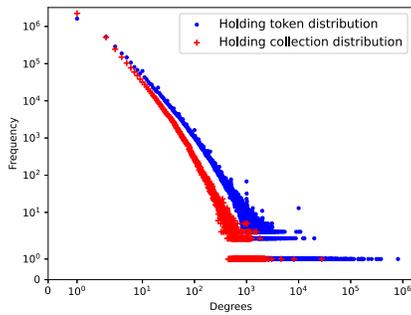
half of year 2018. Meanwhile, we further investigate the data with a monthly basis, which makes it possible for us to locate the specific months. Figure 5b demonstrates that the anomaly is lasting from January to August, and reaches its peak at February. These finer granularity analyses demonstrate that monthly data may be more helpful in the scenario of anomaly detection.

D Dynamic Behavior Analyses

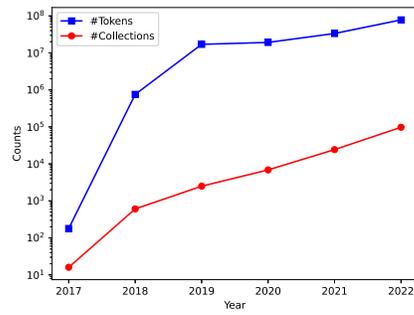
D.1 Active Period of Nodes

Based on our aforementioned analyses, the NFT transaction network is highly active with new nodes continuously adding in. We now proceed to investigate the nodes' active periods. Here, the node's active period is defined as the time interval between its first transaction and its last transaction. In this case, we discard those nodes with only one transaction in the dataset. Figure 6 presents the statistical information for node's active duration in the scale of days. Note that, we only show the active periods from 1 day to 40 days in Figure 6a. According to our data, 23.43% of the nodes have the active period of 1 day, and up to 47.25% nodes' active periods are 1 month. This is consistent with our previous observations regarding the highly increasing number of edges. It's worth noting that only one account has the active period of nearly 5 years. After checking the data, we find out that this node is associated with the Null address (i.e., $0x000...000$). It is not a surprise, since all NFT mint activities would build connections with the Null address. This also indicates that there exist new NFT tokens being continuously minted, which reveals the fast growing speed of NFT market.

Then, we explore the impact of a node's active period on the number of transactions it engages in. One natural hypothesis is that the longer the node's active period, the more transactions it will create, which is consistent with the user behaviours in social networks. Figure 6b shows the number of average transactions with different active periods. As can be seen, the active period and its average



(a) Distributions of tokens and collections



(b) Number of tokens and collections

Figure 7: Distributions of tokens and collections.

transactions are positively correlated. The number of transactions increase exponentially as the active periods become longer. We also notice two spikes with active periods of 8 days and 39 days. Among them, address $0x610...662$ ⁷ causes the spikes of average transactions in active periods of 8 days, which is a Ethereum Name Service (ENS) migration contract to migrate second-level names from the old registry and registrar to the new ones. It creates more than 685,000 transactions within 8 days. To summarize, even through there are some addresses with rather limited active periods, we can generally conclude that most of the addresses are quite active with lots of transactions.

D.2 Holding Tokens and Collections

To reveal the characteristics of the NFT economy, we analyse the distributions of holding tokens and collections for different accounts. Formally, we trace all the received tokens and sent tokens for each account. Through this way, we can know the owner of each token and the quantity of tokens held by an account. Figure 7a shows that both the accounts’ holding tokens and collections are power law distributions. Specifically, 42.10% of the accounts hold only one token, and 58.50% of the accounts hold tokens from one NFT collection. Moreover, there are about 81.70% accounts holding no more than 10 tokens, and 91.40% of accounts hold tokens from no more than 10 collections. In the middle of year 2022, the largest “holder” is the Null address (i.e., $0x000...000$), which involves 809,125 tokens from 8,126 collections. Note that the Null address is a special account, and sending tokens to Null address means destroying the tokens whereas receiving tokens from Null address indicates minting tokens. There are several reasons that people destroy their tokens, including reducing the supply to increase a collection’s value or rectifying error information in the tokens. Then, we look at the next valid holding address. The second largest holder is associated with address $0x000...7a2$ ⁸, which holds 381,570 tokens from 1,454 collections. Since it holds so many tokens, we are interested in uncovering its identity. Therefore, we try to uncover all the relevant activities associated with the address $0x000...7a2$. First of all, we search the address in the Ethereum Name Service and find that the address is bound with the name *stronghands.eth*⁹. Then, we search the keyword “stronghands” with Google, and the results show that it is a blockchain community¹⁰, which supports issuing tokens, trading NFTs and bridge integrations, etc.

Table 6 and Table 7 list the top-10 largest holders in the year of 2020 and 2021, respectively. As can be seen, most top-10 holders in 2020 continue to be top-10 holders in 2021 as well. It is interesting to note that the relative positions are almost same, except that address $0xd9a...6a5$ and $0x721...ace$ swap their positions in 2021, and the Null address becomes the second largest “holder”. Among them, address $0xff1...2c8$ drops out of the top-10 list in 2021, and address $0xe05...9d5$ enters the top-10 list in 2021, which are annotated as bold. The little difference in top-10 list for year 2020 and 2021 indicates that it is very difficult to be one of the top holders for new users, since it costs a lot of money to mint or buy NFTs. Moreover, when looking at the numbers in Table 6 and Table 7, we can find that

⁷ $0x6109DD117AA5486605FC85e040ab00163a75c662$

⁸ $0x0008d343091ef8bd3efa730f6aae5a26a285c7a2$

⁹<https://app.ens.domains/address/0x0008d343091ef8bd3efa730f6aae5a26a285c7a2>

¹⁰<https://www.stronghands.io/>

Table 6: Top-10 accounts’ holding tokens and collections in 2020.

Account Address	Tokens	Collections
0x0008d343091ef8bd3efa730f6aae5a26a285c7a2	363,962	36
0x26cdee4269273e1ea5dfac6b5791df2656897738	343,413	14
0xe4a8dfca175cdca4ae370f5b7aaff24bd1c9c8ef	308,916	13
0xf7ee6c2f811b52c72efd167a1bb3f4adaa1e0f89	216,477	39
0x09c1e4c1adad99436b5c22a395174a1320ee716b	166,321	1
0xf33bd4edc6dcd7240966f20401014ad0018d065b	161,368	19
0xd9ab699e5e196139b8a1c8f70ead01b2137fc6a5	152,788	16
0x721931508df2764fd4f70c53da646cb8aed16ace	149,115	49
0x00	143,849	683
0xff18298382948028f9d93c4e32be1382204022c8	140,025	22

Table 7: Top-10 accounts’ holding tokens and collections in 2021.

Account Address	Tokens	Collections
0x0008d343091ef8bd3efa730f6aae5a26a285c7a2	378,893	198
0x00	365,565	1,914
0x26cdee4269273e1ea5dfac6b5791df2656897738	343,413	14
0xe4a8dfca175cdca4ae370f5b7aaff24bd1c9c8ef	308,880	13
0xf7ee6c2f811b52c72efd167a1bb3f4adaa1e0f89	217,172	83
0x09c1e4c1adad99436b5c22a395174a1320ee716b	166,321	1
0xe052113bd7d7700d623414a0a4585bcae754e9d5	163,499	1,467
0xf33bd4edc6dcd7240966f20401014ad0018d065b	161,413	20
0x721931508df2764fd4f70c53da646cb8aed16ace	159,464	204
0xd9ab699e5e196139b8a1c8f70ead01b2137fc6a5	152,788	16

the top holders’ tokens and collections are rapidly growing. One exception is the address *0x09c...16b*, which holds 166K tokens from the same collection (i.e., *DozerDoll*, a game dirven NFT). On average, they hold 10K more tokens in 2021 compared with the year 2020. This is because there are more and more tokens as well as collections. Figure 7b also verifies this observation. Take the year of 2021 as an example, it increases 17,428 collections and 14 million tokens. Thus, the whole NFT ecosystem is still in its bull market.

D.3 Evolution of Diameters

We study the diameter of the NFT transaction network in this section, which reflects the communication efficiency among different nodes. Generally, the network’s diameter is defined as the largest shortest path among all pairs of nodes in the network. As pointed out by [41, 42], this metric is very sensitive to the noise in the network. For instance, a single long path would result in a large diameter. Thus, we resort to the *effective diameter* used in [41], which is defined as the 90-th percentile of the shortest path length among all pairs of nodes. Figure 8 presents the diameters as it evolves over time by years. The blue line shows the diameter calculated with the whole transaction network, which includes the special Null address. In contrast, the red line indicates the diameter computed without the Null address, which means all the NFT minting and destroying activities are removed from the network. We can observe that these two lines show totally different trends, i.e., one is increasing and the other is decreasing. The final diameter is about 3.0 in 2022 when including the Null address, and it is only half of the value when removing the Null address. It is surprised to see that the Null address has such large influence on this property. With in-depth analysis, we find that there are about 3.5 million addresses connecting with the Null address, which accounts for 77.3% of the total addresses. Thus, the Null address is a huge hub node, and provides a shortcut for nodes to reach each other. As a result, the diameter is smaller in this situation, and every non-mint action would increase its value when the network expands.

To eliminate the effect of the Null address, we focus on the analysis of diameter without Null address, which reveals the transaction behavior patterns without mint. Similar to citation networks that the diameter is shrinking observed by Leskovec et al [41], our results show that the diameter also shrinks in NFT transaction network, i.e., decreasing from 5.94 to 4.72. Then, we explore the graph structure to see whether we can explain the diameter shrinking phenomenon. Specifically, we observe that

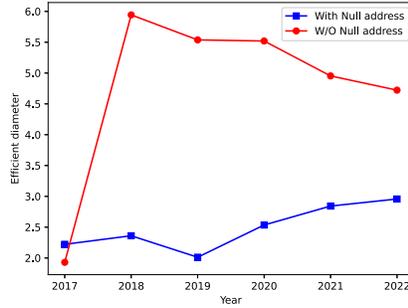


Figure 8: Effective diameter by years.

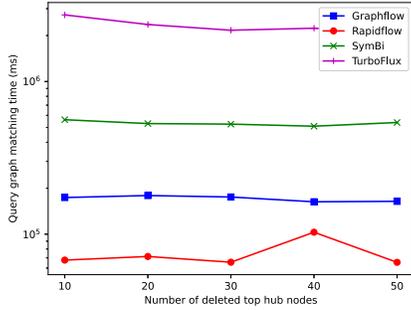
about 42.28% of the nodes have the degree of 1 in the final network. Those nodes with degree 1 are the key factor that leads to the increase of diameter. Thus, we remove these degree 1 nodes, and calculate the effective diameter of the remaining part. The value is 4.20, which shrinks again. After that, we repeat this process again, and find that it only has 2.26% nodes with degree 1 in the remaining part. Similarly, we delete those degree 1 nodes, and compute its effective diameter again, which is 4.28 and stays the same as before. This suggests that there is a giant component with extremely high connectivity, and nodes with degree 1 is only a thin layer at the outside of the well-connected giant component. This observation brings new challenges for some downstream tasks.

E More Analyses on Continuous Subgraph Matching

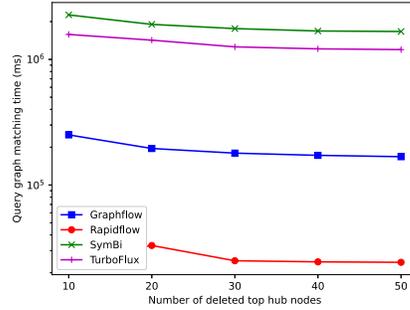
Frameworks. We evaluate the performance of six recent CSM frameworks in this section. (1) SJ-Tree [13] proposes a lazy search algorithm, and the search strategy is determined by a vertex-to-vertex basis, which depends on the likelihood of a matching in the vertex neighborhood. (2) IEDyn [30, 31] randomly selects a node from query graph, then it products matching order by conducting DFS on the query graph. (3) SymBi [50] employs a dynamic candidate space as an auxiliary data structure for filtering. (4) Graphflow [33] first generates matching order offline, and then retrieves it in online processing. (5) TurboFlux [35] employs a concise representation of the intermediate results, and a novel edge transition model is proposed to identify the update operations that may affect the current solutions. (6) RapidFlow [68] performs batch subgraph matching via designing a query reduction technique, then dual matching is utilized to leverage the duality of the graph in the matching procedure.

Settings. Similar to the link prediction task, we also remove all the transactions related to the Null address, which gets rid of the impact of the extreme large degree node. In previous studies [69, 35, 13], the initial graph is constructed using the first 90% of edges, while the rest 10% of edges are employed as insertion streams. In our scenario, we know the exact time of each transaction, thus we use NFT transactions from year 2017 to the end of 2021 as the initial graph, and then the transactions in the year of 2022 are regarded as the insertion streams. Since the original nodes and edges do not have fine-grained labels, each node is assigned a label randomly selected from a pool of 30 labels. We do not assign labels for edges following the previous works [69, 35]. For evaluation metrics, we report the query time, which denotes the time taken by the online matching procedure to execute. We discard the graph update time, since it’s same for all the frameworks. To complete the experiments under an affordable time, a one-hour time limit is imposed for query processing (i.e., 3.6×10^6 ms). Additionally, we also calculate the number of matched subgraphs for each query graph.

Results. As we have discussed in Section 4.2, there exist lots of hub nodes in the NFT transaction network. Therefore, we are interested in whether updating the index of hub nodes will have a great impact on the query time. Specifically, except Graphflow, all the remaining frameworks belong to the index-based incremental computation. In this context, the query time consists of the index time, which signifies the time spent on updating the index, and the enumeration time that enumerates the matched results. Therefore, we first sort the node degrees in the descending order, and the results show that the top-50 largest hub nodes have the largest degree with 433,114 and the smallest degree



(a) Query time on pattern 1



(b) Query time on pattern 3

Figure 9: Query time with deleted hub nodes.

Table 8: Temporal link prediction results with live-update settings in sampled subset. We repeat experiments with 3 random seeds to report the mean as well as standard deviation of AUC and MRR. We also present the results under different time snapshot granularities, e.g., days, weeks and months.

Models	Snapshot Days		Snapshot Weeks		Snapshot Months	
	AUC	MRR	AUC	MRR	AUC	MRR
Dyngraph2vec	77.39±2.04	40.29±5.75	75.33±7.15	36.84±4.97	72.03±2.73	34.83±4.20
TGCN	95.49±0.78	62.40±7.89	91.09±0.46	41.69±7.82	86.13±1.33	43.67±6.23
EvolveGCN	84.34±2.81	40.91±7.96	79.13±0.92	40.72±3.82	78.40±3.33	42.86±3.71
GCRN-GRU	91.41±1.24	33.65±9.48	90.78±2.43	42.88±0.48	82.64±0.22	38.99±0.18
GCRN-LSTM	93.97±1.66	41.63±3.22	90.16±1.83	39.20±7.19	83.38±0.71	37.22±2.63
DynGEM	95.67±0.74	49.25±6.01	93.34±0.98	45.41±5.29	90.75±1.29	41.38±0.61
Roland-MA	97.55±0.59	37.18±6.76	94.59±0.50	44.26±1.26	87.69±2.17	37.98±4.61
Roland-MLP	96.38±1.07	60.44±10.0	89.25±0.89	52.32±3.20	87.89±0.96	42.02±3.37
Roland-GRU	96.52±0.08	69.34±2.68	92.83±0.82	52.74±2.60	91.01±0.77	41.59±2.02

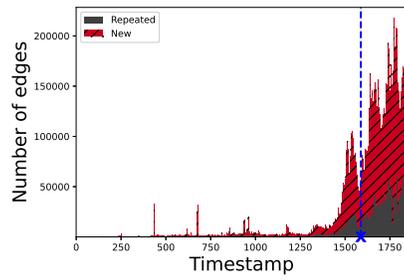
with 7,623. Then, we delete the top-50 hub nodes sequentially and report the query time in Figure 9 for pattern 1 and pattern 3. Note that, we remove the randomly assigned node labels in this situation, which counts the query time associated with the hub nodes. As can be seen, the query time almost remains the same with acceptable fluctuations across all the frameworks except for RapidFlow. This is because RapidFlow is extremely efficient (i.e., only taking several hundred milliseconds), thus a slight fluctuation will be very obvious in the figure. We can conclude that these four frameworks are applicable to graphs with dense structures. They can effectively serve as a bridge to generate ground truth for continuous subgraph matching, providing valuable support for the training of deep graph learning models.

F Temporal Link Prediction Results on Sampled Subset

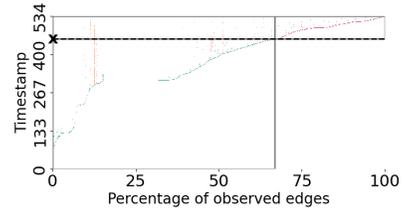
To fully evaluate the models’ performance, we also sample a subset of data spanning from January 1st 2020 to August 31st 2020. This subset contains 88,112 nodes and 203,221 edges. We provide the temporal link prediction results under live update setting in Table 8. As can be seen from the results, we can draw similar conclusions. Notably, we also observe that with the increase in time granularity, the performance of the models exhibits a corresponding decline within this sub-sampled dataset.

G TEA and TET Plots

Following the definitions in [57], we present Temporal Edge Appearance (TEA) and Temporal Edge Traffic (TET) plots in Figure 10. Specifically, a TEA plot visualizes the proportion of recurring edges compared to newly joined edges at each timestamp within a temporal graph. In Figure 10a, the gray bar represents the count of edges that are previously observed, while the red bar signifies the quantity of new edges generated at each subsequent time step. The TEA plot illustrates that our



(a) TEA plot with day intervals



(b) TET plot with sampled 2 million edges

Figure 10: TEA and TET plots.

dataset exhibits a substantial proportion of newly formed edges at each time step, which means the transaction graph is highly active. In contrast, a TET plot represents the recurrent pattern of edges across different time intervals. Edges are colored to indicate whether they appear solely in the training set (green), solely in the test set (red), or in both sets (orange). As it is time consuming to plot all the 124 million edges in the figure, we randomly sample 2 million edges and show the results in Figure 10b. We have tried to use all the data, but the plotting process was unable to complete within 12 hours. The sampled 2 million edges subset showcases a subtle recurrence pattern, which is consistent with our analyses in previous sections.