Sybil-Proof Diffusion Auction in Social Networks

Hongyin Chen
Center on Frontiers of Computing Studies, Peking University
Beijing, China
chenhongyin@pku.edu.cn

Xiaotie Deng
Center on Frontiers of Computing Studies, Peking University
Center for Multi-Agent Research, Peking University
Beijing, China
xiaotie@pku.edu.cn

Ying Wang
Center on Frontiers of Computing Studies, Peking University
Beijing, China
wying2000@pku.edu.cn

Yue Wu
Center on Frontiers of Computing Studies, Peking University
Beijing, China
wu.y@pku.edu.cn

Dengji Zhao
ShanghaiTech University
Shanghai, China
zhaodj@shanghaitech.edu.cn

ABSTRACT

A diffusion auction is a market to sell commodities over a social network, where the challenge is to incentivize existing buyers to invite their neighbors in the network to join the market. Existing mechanisms have been designed to solve the challenge in various settings, aiming at desirable properties such as non-deficiency, incentive compatibility and social welfare maximization. Since the mechanisms are employed in dynamic networks with ever-changing structures, buyers could easily generate fake nodes in the network to manipulate the mechanisms for their own benefits, which is commonly known as the Sybil attack. We observe that strategic agents may gain an unfair advantage in existing mechanisms through such attacks. To resist this potential attack, we propose two diffusion auction mechanisms, the Sybil tax mechanism (STM) and the Sybil cluster mechanism (SCM), to achieve both Sybil-proofness and incentive compatibility in the single-item setting. Our proposal provides the first mechanisms to protect the interests of buyers against Sybil attacks with a mild sacrifice of social welfare and revenue.

KEYWORDS

Sybil attack; Mechanism design; Social network; Diffusion auction

ACM Reference Format:
Hongyin Chen, Xiaotie Deng, Ying Wang, Yue Wu, and Dengji Zhao. 2023. Sybil-Proof Diffusion Auction in Social Networks. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Auction is an important method of selling commodities where the seller collects bids from buyers and allocates commodities according to these bids. Previous works [1] have shown that more buyers would significantly lead to higher social welfare and revenue in auctions. However, buyers have no incentive to invite others to their auctions because it would cause tougher competition and hurt their own interests. Recently, there has been an emergence of studies on the diffusion auction over social networks [6], which studies mechanisms that incentivize buyers to invite new agents to an auction via a social network. In these works, while the seller only knows her neighbors, any buyer who is informed of the auction may bid to buy, as well as diffuse the information about the auction to her neighbors to improve her utility.

The first work of diffusion auctions [11] proposed the information diffusion mechanism (IDM) for selling one item in a social network, focusing on the incentive compatibility of its information propagation action. Under IDM, it's a dominant strategy for each bidder to truthfully bid her private valuation and to diffuse the auction information to all her neighbors. Zhao *et al.* [22] and Kawasaki *et al.* [8] further designed diffusion mechanisms in selling multiple homogeneous items in a social network. Those works, however, did not consider a common threat known as the Sybil attack.

The first study on Sybil attacks [5] considered a situation in peer-to-peer systems where malicious agents may gain an unfair advantage by creating fake identities. One such example is presented in Figure 1 for a social network where the agent x creates six false-name identities. In this example, a directed edge from a vertex x to y means that agent x knows the existence of y. When agent x creates Sybil identities, they cannot connect to other agents that she does not know. The Sybil attack is a significant threat to

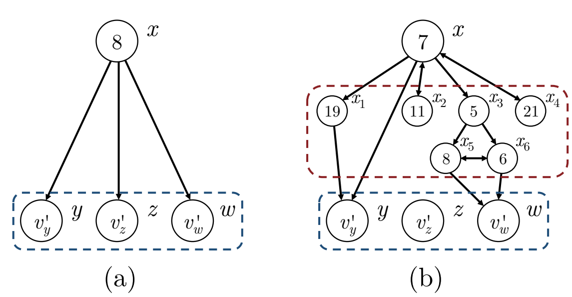

Figure 1: (a) the true type of x; (b) a Sybil attack of x involving six Sybil identities (in the red rectangle).

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

auctions (commonly called *false-name bids*) and has been extensively investigated in traditional auction settings [18]. One example proven to be vulnerable to Sybil attacks is the well-known Vickrey-Clarke-Groves (VCG) auction for combinatorial auctions of at least two items.

Our work studies this fundamental issue in diffusion auctions on social networks, where fake nodes can be easily created. In existing mechanisms such as IDM and FPDM [20], intermediate buyers are rewarded for inviting more buyers. This makes Sybil attack highly profitable and harmful. Intermediate buyers under IDM can use Sybil identities as shill bidders to exploit their descendants; under FPDM, buyers with more neighbors are prioritized when allocating the item, and Sybil attack can increase the number of neighbors.

We consider a strong adversarial model that allows every buyer to create fake identities, which can link with other such identities or link to the buyer's neighbors. But they don't have incoming edges from other agents, since they are only visible to the creator. Our goal is to incentivize diffusion without encouraging Sybil attacks in diffusion auctions.

Contribution: We propose two Sybil-proof diffusion mechanisms, the Sybil tax mechanism (STM) and the Sybil cluster mechanism (SCM). STM achieves Sybil-proofness by identifying trustworthy agents. In STM, diffusing to "suspicious" agents is not beneficial to buyers. However, agents in STM do not have a strong incentive to diffuse information. In proposing SCM, we provide a stronger incentive for diffusion, where the reachability of each non-Sybil vertex is credited to some selected agents. By a mild sacrifice on the seller's revenue, SCM creates a strong incentive to invite new buyers.

Our work overcomes several difficulties. We are the first to identify and model the Sybil attack in diffusion auctions. Our adversarial model is the most general form of Sybil attack possible without collusion. Existing diffusion mechanisms cannot resist such Sybil attacks, even after removing "suspicious" agents.

Additionally, we discuss the social welfare and revenue of Sybil-proof diffusion auctions. We prove that there is no optimal SP diffusion auction mechanism for social welfare, and all SP mechanisms perform poorly in the worst case. We conduct experiments under different settings to evaluate the performance of STM and SCM. The results suggest that STM and SCM do not significantly sacrifice welfare and revenue when compared to non-SP mechanisms.

Related Literature. The Sybil attack has become a fundamental issue in traditional social networks [15], where nodes are usually divided into two types: honest ones and Sybil ones. In such settings, various protocols have been proposed to identify Sybil nodes and maintain honest nodes through the graph structures [12, 19].

The Sybil attack is also destructive in auctions. The pioneering work to study Sybil attacks on combinatorial auctions [18] proved that Sybil-proofness and Pareto optimality can't be achieved simultaneously. Many other works have followed. For example, Iwasaki *et al.* [7] have shown that a Sybil-proof combinatorial auction mechanism may result in extremely low social welfare in some cases. In dynamic spectrum access auctions, Sybil-proofness has been only achieved when severe restrictions are imposed on Sybil agents. For example, PRAM [4] requires that if an agent performs the Sybil attack, the sum of bids given by herself and her Sybil identities is equal to her private valuation.

2 PRELIMINARIES

In a social network, a seller s is selling one item to a buyer among the set of potential buyers $N = \{1, 2, \dots, n\}$. The set N is unknown to the seller; instead, she only knows some buyers $r(s) \subseteq N$. Likewise, each buyer $i \in N$ has her private social connections, represented as a set of neighbors $r(i) \subseteq N$. Each buyer $i \in N$ also has a private valuation of the item, which is denoted as v_i. Collectively, each buyer i owns a private type $\theta_i = (v_i, r(i)) \in \mathbb{R}_{\geq 0} \times 2^N$.

In a diffusion auction, s can only advertise the sale to her neighbors $r(s)$ initially. Then, each buyer i with the information of the sale may diffuse it to some of her neighbors in $r(i)$. Recursively, many buyers can be informed. Each buyer is asked to give a bid on the item besides diffusing the sale. The mechanism consequently sells the item to an informed buyer and rewards some buyers for their contribution of inviting others.

We model the bid and diffusion of buyer i as the report type $\theta'_i = (v'_i, r'(i)) \in \mathbb{R}_{\geq 0} \times 2^{r(i)}$, where $v'_i \in \mathbb{R}_{\geq 0}$ is her bid and $r'(i)$ is the set of buyers she diffuses to. A buyer can only diffuse to her neighbors (i.e. $r'(i) \subseteq r(i)$). The input of the mechanism is therefore a report profile $\boldsymbol{\theta}' = (\theta'_1, \dots, \theta'_n)$, and we suppose the set of seller's neighbors $r(s)$ is provided in advance and fixed.

The set of all possible types and reports of buyer i is denoted as $\Theta_i = \mathbb{R}_{\geq 0} \times 2^N$, and we denote the set of all possible profiles as Θ.

Definition 2.1 (Diffusion auction mechanism). A diffusion auction mechanism $\mathcal{M}$ is defined as a pair of allocation and payment schemes $(\boldsymbol{\pi}(\cdot), \boldsymbol{t}(\cdot))$ for arbitrary agent set N:

- allocation scheme $\boldsymbol{\pi} : \Theta \to \{0, 1\}^n$;
- payment scheme $\boldsymbol{t} : \Theta \to \mathbb{R}^n$.

Given the reported type profile $\boldsymbol{\theta}' = (\theta'_1, \dots, \theta'_n)$, whose length is not known in advance, $\pi_i(\boldsymbol{\theta}') = 1$ means that agent i wins the item, and 0 otherwise. She then pays $t_i(\boldsymbol{\theta}') \in \mathbb{R}$ to the seller.

We assume the following feasibility conditions for diffusion auctions throughout the paper:

1. Allocation feasibility: $\sum_{i \in N} \pi_i(\boldsymbol{\theta}') \leq 1$,
2. Anonymity: except for ties, the mechanism output is invariant to any permutation on N, and
3. Ignorance of unreachable vertices: if $i \in N$ is unreachable from s on the social network represented by $\boldsymbol{\theta}'$, then $\pi_i(\boldsymbol{\theta}') = t_i(\boldsymbol{\theta}') = 0$, and the mechanism output must be invariant with respect to θ'_i.

We can use graph theory to formalize the third condition above. The social network represented by the true type profile $\boldsymbol{\theta}$ can be denoted as a graph G with vertex set $V(G) = \{s\} \cup N$ and directed edge set $E(G) = \{(x, y) \mid x \in V(G), y \in r(x)\}$. Likewise, a graph $G(\boldsymbol{\theta}')$ can be defined for the report profile $\boldsymbol{\theta}'$. The subgraph of $G(\boldsymbol{\theta}')$ with vertices that are reachable from s is denoted as $G_s(\boldsymbol{\theta}')$. All vertices unreachable from s are excluded from it. When $\boldsymbol{\theta}'$ can be inferred from context, we omit it and write $G_s(\boldsymbol{\theta}')$ as G_s. The ignorance condition means that the mechanism can only use the structural information about $G_s(\boldsymbol{\theta}')$ and the bids of $V(G_s)$ as inputs. This is a key difference between a diffusion mechanism and the traditional auction mechanism.

The agents have a quasi-linear utility model. Given a buyer's true type $\theta_i = (v_i, r(i))$ and the report profile of all agents $\boldsymbol{\theta}'$, her utility under mechanism $\mathcal{M} = (\boldsymbol{\pi}, \boldsymbol{t})$ is $u_i(\theta_i, \boldsymbol{\theta}', \mathcal{M}) = v_i \cdot \pi_i(\boldsymbol{\theta}') - t_i(\boldsymbol{\theta}')$.

2.1 Non-deficiency, Individually Rationality and Incentive Compatibility

In this section, we define the objectives of diffusion mechanisms.

An individually rational mechanism is one in which every buyer can attain a non-negative utility by reporting truthfully, regardless of what other agents do. This means that any agent is at least willing to participate.

Definition 2.2 (IR). A diffusion mechanism $\mathcal{M} = (\boldsymbol{\pi}, \boldsymbol{t})$ is *ex-post individually rational* (IR) if for all $\boldsymbol{\theta} \in \Theta$, for all $i \in N$ with $\theta_i = (v_i, r(i))$, it is guaranteed that $u_i(\theta_i, \boldsymbol{\theta}, \mathcal{M}) \geq 0$. [1]

A desired mechanism encourages agents to behave truthfully, i.e., to bid their private values and to diffuse the information to all their neighbors. In a diffusion auction, an agent may act strategically by overbidding, underbidding, or not sharing information with some neighbors, if it provides a benefit. Dominant-strategy incentive compatibility requires that reporting the true type is a dominant strategy for every buyer, ruling out these strategic reports.

Definition 2.3 (DSIC). A diffusion mechanism $\mathcal{M} = (\boldsymbol{\pi}, \boldsymbol{t})$ is *dominant-strategy incentive compatible* (DSIC, or IC for short) if, for any buyer $i \in N$ with type θ_i, any report profile of other agents $\boldsymbol{\theta'}_{-i}$ and any $\theta'_i \in \Theta_i$ satisfying $r'(i) \subseteq r(i)$, we have $u_i(\theta_i, \boldsymbol{\theta}, \mathcal{M}) \geq u_i(\theta_i, \boldsymbol{\theta'}, \mathcal{M})$, where $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta'}_{-i})$ and $\boldsymbol{\theta'} = (\theta'_i, \boldsymbol{\theta'}_{-i})$.

Some IC diffusion auction mechanisms, like VCG, may give a negative revenue to the seller [11]. We define the following non-deficiency condition to rule out these mechanisms.

Definition 2.4. A diffusion mechanism $\mathcal{M}$ is *non-deficit*, or weakly budget balanced, if its revenue for the seller is always non-negative, or formally, $\mathcal{R}^{\mathcal{M}}(\boldsymbol{\theta}) \geq 0$ for all $\boldsymbol{\theta} \in \Theta$, where $\mathcal{R}^{\mathcal{M}}(\boldsymbol{\theta}) = \sum_{i \in N} t_i(\boldsymbol{\theta})$ is the revenue to the seller.

2.2 The Sybil Attack and Sybil-Proofness

In this paper, we also want to disincentivize Sybil attacks. A Sybil attack happens when a buyer i creates multiple fake Sybil identities (or false-name identities) $i_1, i_2, \ldots, i_k$, each with its report $\theta'_{i_1}, \ldots, \theta'_{i_k}$. We call the set of all identities of i as $\phi = \{i, i_1, \ldots, i_k\}$. For each identity $i_j \in \phi$, its report $\theta'_{i_j} = (v'_{i_j}, r'(i_j))$ must satisfy $r'(i_j) \subseteq r(i) \cup \phi$ since i does not know any agent beyond herself, her neighbors, and her Sybil identities. See Figure 1 for an illustration of Sybil attacks.

We define Sybil-proofness as a criterion for ruling out such attacks. A mechanism is Sybil-proof if, for every buyer, any form of Sybil attack cannot bring a higher utility.

Definition 2.5 (SP). A diffusion mechanism $\mathcal{M} = (\boldsymbol{\pi}, \boldsymbol{t})$ is *Sybil-proof* (SP) if, for any type profile $\boldsymbol{\theta}$, any buyer $i \in N$, and for all $\theta'_i, \theta'_{i_1}, \ldots, \theta'_{i_k} \in \Theta_i$ satisfying $r'(i) \subseteq \phi \cup r(i)$ and $\forall i_j \in \phi : r'(i_j) \subseteq \phi \cup r(i)$, we have

$$u_i(\theta_i, \boldsymbol{\theta}, \mathcal{M}) \geq u_i(\theta'_i, \boldsymbol{\theta'}, \mathcal{M}) + \sum_{i_j \in \phi} u_{i_j}(\theta'_{i_j}, \boldsymbol{\theta'}, \mathcal{M})$$

where the Sybil-attack report profile is $\boldsymbol{\theta'} = (\theta'_i, \theta'_{i_1}, \ldots, \theta'_{i_k}, \boldsymbol{\theta}_{-i})$.

[1]The definition of IR in previous literature in diffusion auctions does not require the buyer to truthfully diffuse, which differs from the traditional definition in AGT. In the setting of this paper, the two definitions are equivalent, and the traditional definition is presented.

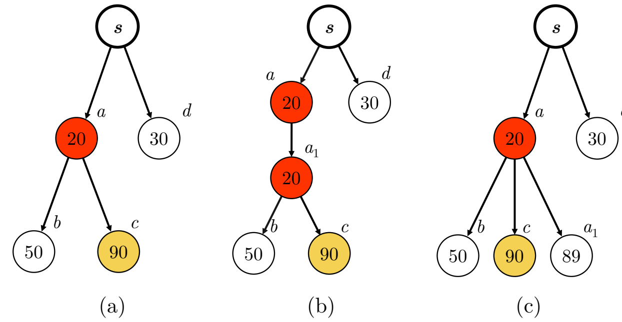

Figure 2: Sybil attack counterexamples of VCG and IDM.

A Sybil attacker can bring an arbitrary number k of Sybil identities, and each of the $k + 1$ identities (including the agent herself) can report arbitrarily. This formulation is the most general form of Sybil attacks without collusion. As a degenerate case, a Sybil attack with $k = 0$ is equivalent to a single-agent strategic play in the previous diffusion action setting. Therefore, Sybil-proofness implies incentive compatibility.

2.3 Vulnerability of Existing Mechanisms

To the best of our knowledge, none of the existing diffusion auctions in the literature is Sybil-proof; the only exception is the trivial Neighbor Second-Price Auction (NSP), where only the seller's neighbors are considered with a second price auction[2] (see Appendix B in the full version of this paper [2] for a detailed definition). As assumed, the seller's neighbors are known to the seller, so there is no chance for them to create fake identities to join NSP.

Other existing mechanisms for diffusion auctions are all vulnerable to the Sybil attack. Here we use the two typical mechanisms proposed in [11], VCG and IDM, to demonstrate the possibility of Sybil attacks. Definitions of these mechanisms are given in Appendix B.

Observation 1. *VCG and IDM are not Sybil-proof.*

The classic VCG mechanism can be easily extended as a diffusion auction. Under VCG, the item is sold to the highest bidder, and other agents are paid the social welfare increase due to their participation. In the example shown in Figure 2(a), if the intermediate node a does not participate, b and c will be unable to join, and the social welfare will be 30. With a's participation, the social welfare is 90, so VCG will pay 60 to a. Now, if a creates a fake identity a_1, then both a and a_1 will be paid 60 (a successful Sybil attack).

Since VCG paid a lot to the agents connecting the highest bidder to the seller, it cannot be non-deficit. Thus, IDM was proposed to guarantee that the seller's revenue is non-negative. IDM does not directly sell the item to the highest bidder; it uses a resale process to find the winner. It first allocates the item to the first cut point to reach the highest bidder, and the buyer pays the highest bid without her participation. In the example shown in Figure 2(a), the item is

[2]Confusingly, there are two mechanisms named VCG in the literature of diffusion auctions: 1) the single-item VCG auction among the seller's neighbors, and 2) the generic VCG mechanism applied to diffusion auction [11]. To disambiguate, the former is called Neighbor Second-Price (NSP) in this paper.

first allocated to a and a pays 30. Then a can choose to resell it to c and c has to pay the highest without c to a, which is 50. Now, if a creates a fake neighbor a_1 with bid 89, then c will need to pay 89 to a (another successful Sybil attack).

We also proved that some other existing diffusion mechanisms [20, 21] are not Sybil-proof in Appendix of our full version [2].

3 ANALYSIS OF THE SYBIL ATTACK

In this section, we study the features of Sybil attacks. In our model, a Sybil identity y created by a real agent x can only be connected by the other Sybil identities of x or by x itself. Thus, any path from s to y must include x. In graph theory [9], this is known as x *dominating* y, or $x \operatorname{dom} y$, and x is called a *dominator* of y. Every vertex y has at least two dominators: y itself and s. If it has no other dominators, one can be sure that y is not a Sybil identity.

In previous diffusion auction mechanisms like VCG and IDM, an agent is rewarded for inviting others *only if* it is a dominator of the item winner, which makes Sybil attacks (creating dominatorship) profitable. This explains why Sybil-proofness is hard to achieve in diffusion auctions.

An *immediate dominator* of x, denoted as $v = idom(x)$, is defined as the unique vertex v who dominates x and is dominated by every other dominator $w \neq v$ of x.

Theorem 3.1. *Every vertex on the graph $G_s(\theta')$ except s has an immediate dominator, and the edges $\{(idom(x), x) \mid x \in N\}$ form a directed tree with s being its root, called the* dominator tree *of $G_s(\theta')$ rooted at s.*

This is exactly Theorem 1 of [9]. The definition of dominators is identical to *diffusion critical nodes* in [11], and the path from s to x on the dominator tree is the *diffusion critical sequence* of x.

3.1 Graphical Non-Sybil Agents

In this subsection, we use graph theory to characterize the set of vertices that cannot be Sybil identities. Firstly, the seller and her neighbors are not Sybil identities. To account for trustworthy entities like public figures and centralized institutions, we allow for an optional set of vertices to be provided externally, which must be guaranteed to be free of Sybil identities. This set Γ_0 defaults to be $\emptyset$ if not provided. Allowing such external information makes our mechanisms more flexible.

For the convenience of expression, we first give the definition of meeting points.

Definition 3.2 (Meeting points). For a pair of vertices x, y, a vertex z is defined to be a *meeting point* of x and y if there are two vertex-disjoint paths to z, from x and y respectively.

If a vertex is a meeting point of two other non-Sybil vertices x, y, it must not be a Sybil identity. This is because all paths from non-Sybil vertices to a Sybil identity i_j must contain its owner i which contradicts the definition of meeting points. Therefore, we have the following definition of graphical non-Sybil agents which iteratively collects meeting points of existing members.

Definition 3.3 (Graphical non-Sybil agents). The set $\Gamma(\theta') \subseteq V(G_s)$ is defined as follows:

(1) Initialize the set as $\Gamma(\theta') := \{s\} \cup r(s) \cup \Gamma_0$.
(2) For each pair of vertices $x, y \in \Gamma(\theta')$, if z is a meeting point of them in graph G_s, then add z to the set, i.e. $\Gamma(\theta') := \Gamma(\theta') \cup \{z\}$.
(3) Repeat step 2 until there are no more vertices to add.

It can be shown that $\Gamma(\theta')$ is precisely the maximal set of vertices that cannot be Sybil identities. This will be proven in Lemma 5.3 after the introduction of Sybil clusters.

3.2 Overly Sensitive Mechanism

Given the graphical non-Sybil agents, a straightforward idea to achieve Sybil-proofness is to apply the existing diffusion mechanisms on non-Sybil agents. This idea of detection and removal is a common solution to Sybil attacks in social networks [12, 16, 17]. However, we find that such an approach doesn't work because an agent can misreport her neighbor set and turn non-Sybil agents into suspicious ones.

We propose the *overly sensitive* mechanism (OSM) to show why such an idea does not work. In OSM, we ignore all potential Sybil identities (i.e. all $i \notin \Gamma(\boldsymbol{\theta}'))$ and focus on the reachable part of the induced subgraph inducing from $\Gamma(\boldsymbol{\theta}')$, denoted as $G_s[\Gamma(\boldsymbol{\theta}')]$. The subgraph $G_s[\Gamma(\boldsymbol{\theta}')]$ contains only vertices in $\Gamma(\boldsymbol{\theta}')$, and for each vertex x in it, there is a path from s to x that only passing non-Sybil agents. We adopt IDM on $G_s[\Gamma(\boldsymbol{\theta}')]$.

OSM seems Sybil-proof because Sybil identities are all ruled out. However, we find that OSM is not even incentive compatible. In OSM, the detection-and-removal process can be exploited by malicious agents. In Figure 3a, every vertex will be in $\Gamma(\boldsymbol{\theta})$. Under IDM, a will buy the item with the second-highest price $v'_c = 9$. However, if a chooses not to diffuse the information to c as in Figure 3b, c would be excluded from $G_s[\Gamma(\boldsymbol{\theta}')]$, and a would get the item with a lower payment of 7.

Since SP implies IC, OSM is not Sybil-proof either. Therefore, we need a new approach to resist Sybil attacks in diffusion auctions.

4 SYBIL TAX MECHANISM

In this section, we present the first main contribution of this paper, our first Sybil-proof diffusion mechanism, called Sybil Tax Mechanism (STM).

Before describing STM, we introduce some notations. We use $\text{Max}[S]$ to denote the highest bid in a set S, that is, $\text{Max}[S] = \max_{x \in S} v'_x$. We also denote the vertices she dominates as $\alpha(x) =$

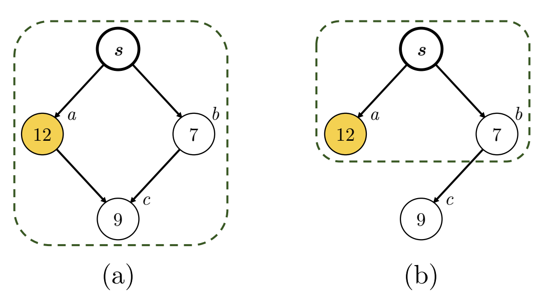

Figure 3: A counterexample for the overly sensitive mechanism. The set $\Gamma(\boldsymbol{\theta}')$ is denoted by the dashed border rectangle.

$\{y \mid x \operatorname{dom} y\}$ for every vertex x. The vertex x is critical for these $y \in \alpha(x)$, because without her diffusion, these vertices are not reachable from s. It is also known as *diffusion critical children* in the terminology of previous literature on diffusion auctions.

Sybil Tax Mechanism (STM)

(1) Given the reported type profile $\boldsymbol{\theta}'$, we first calculate the reachable reported graph $G_s(\boldsymbol{\theta}')$ and the graphical non-Sybil agent set $\Gamma(\boldsymbol{\theta}')$. Let's write them as G_s and Γ for short.
(2) Find the reachable buyer with the highest bid, denoted by x^*, where $v'_{x^*} = \mathrm{Max}[V(G_s)]$.
(3) Compute the dominator sequence $C_{x^*} = \{c_0 = s, c_1, \ldots, c_\ell = x^*\}$. Specifically, we have $c_j = idom(c_{j+1})$, for all $0 \le j < \ell$.
(4) We define p_j, the buying price of c_j, as the highest bid without the participation of c_j. The selling price of c_j, denoted as q_j, is defined as the highest bid among all vertices that are guaranteed to not be a Sybil identity of c_j. Formally,

$$p_j = \mathrm{Max}[V(G_s) \setminus \alpha(c_j)] \qquad \text{for } 1 < j \le \ell,$$
$$q_j = \mathrm{Max}[(V(G_s) \setminus \alpha(c_j)) \cup \beta_j] \qquad \text{for } 1 \le j < \ell,$$

where

$$\beta_j = \{x \mid \exists y \ne c_j, y \in (\alpha(c_j) \setminus \alpha(c_{j+1})) \cap \Gamma, x \in \alpha(y)\}.$$

(5) Pick a c_d with the lowest index d that satisfies $v'_{c_d} \ge q_d$. If such d does not exist among index $1 \le d < \ell$, we set $d = \ell$. This agent c_d wins the item with $\pi_{c_d}(\boldsymbol{\theta}') = 1$. The payment function is calculated as

$$t_{c_j}(\boldsymbol{\theta}') = \begin{cases} p_j - q_j & \text{for } 1 \le j < d, \\ p_j & \text{for } j = d. \end{cases}$$

(6) The payment and allocation of all other buyers are zero.

In this mechanism, the item is sold along the dominator sequence from s to x^* as a series of successive transactions between neighboring agents. Agent j's buying price p_j is set as other agents' optimal social welfare (i.e. the highest bid of them) when she does not participate in the auction. This ensures that her report cannot lower her buying price. The agent j can sell the item further down the critical sequence to reach more potential buyers with a selling price of q_j. To achieve Sybil-proofness, we need Sybil-attacking to be not profitable, i.e. not able to increase q_j. When the item is passed from c_j to c_{j+1}, since the latter may be a Sybil identity of the former, the selling price q_j of c_j must be irrelevant to the report of c_{j+1}. Indeed, q_j is defined as the highest bid among those who are guaranteed not the Sybil identity of her. The set β_j is defined in a way that it is monotonically increasing with the report of c_j to incentivize diffusion, and that it contains no Sybil identity.

Conceptually, a buyer who gets the item can choose to keep it or to resell. She will pass the item only when her selling price is higher than her private value. STM simulates this choice based on buyers' bid through the choice of the winner d.

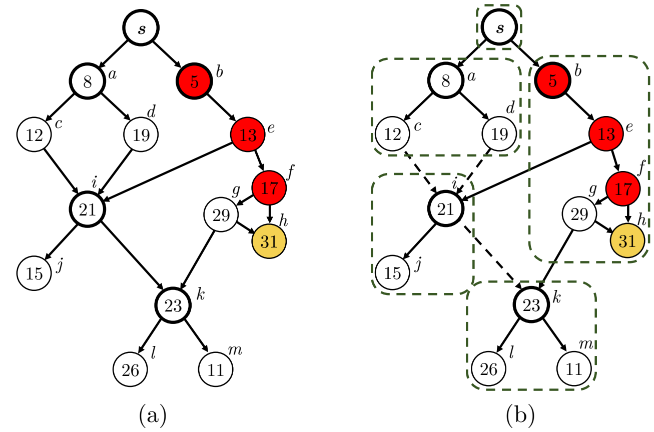

Figure 4: Examples of the Sybil tax mechanism and the Sybil cluster mechanism.

This series of transactions are summed up by STM. In a single transaction, buyer c_j will receive q_j units of money and c_{j+1} pays p_{j+1} for it. The price difference $p_{j+1} - q_j$ can be considered as a "tax" paid by the intermediate buyers (which we call *brokers*) to prove their innocence.

Since $V(G_s) \setminus \alpha(c_j) \subseteq (V(G_s) \setminus \alpha(c_j)) \cup \beta_j \subseteq V(G_s) \setminus \alpha(c_{j+1})$, we have $p_j \le q_j \le p_{j+1}$, so the monetary gain of brokers and the tax are all non-negative. This leads to individual rationality and non-deficiency.

Figure 4(a) illustrates STM with an example. We assume that the externally provided set Γ_0 is empty. When all buyers report their type truthfully, the mechanism runs as follows.

The set of graphical non-Sybil agents $\Gamma(\boldsymbol{\theta})$ is calculated as $\{s, a, b, i, k\}$. The mechanism identifies the highest-bidder h and calculates the dominator sequence $C_h = \{c_0 = s, c_1 = b, c_2 = e, c_3 = f, c_4 = h\}$. The item is sold to $c_4 = h$ because c_4 is the only buyer on the dominator sequence that satisfies $v'_{c_j} \ge q_j$. Then we calculate the payments. For brokers c_1, c_2, c_3, $p_1 = p_2 = p_3 = 26 = q_1 = q_2 = q_3$, so they get paid 0. The winner $c_4 = h$ pays $p_4 = 29$. The seller gets a revenue of 29. In short, the buyer h will pay 29 to buy the item. Other buyers get zero utility.

Theorem 4.1 (Main). *STM is IR, non-deficit and Sybil-proof.*

The proof of Theorem 4.1 will be elaborated in Appendix E of [2]. We provide a proof sketch here.

Intuitively, STM is individually rational because $q_j \ge p_j$ and $v'_{c_d} \ge p_d$. A non-Sybil-attacking buyer c_j would want to maximize β_j to maximize her utility, which can be achieved by maximal diffusion. By the graph-theoretic properties of Sybil attacks, if a Sybil attack happens on the dominator sequence, the identities of the same buyer must be contiguous on the sequence, and the tax paid by such brokers would disincentivize this attack.

The above theorem shows that STM is incentive compatible because Sybil-proofness implies IC. In a previous work [10], Bin Li *et al.* identified one class of diffusion mechanisms called critical diffusion mechanism (CDM) on social graphs, which covers a large class of incentive compatible mechanisms. The successive reselling in STM resembles CDM, but STM is not a member of that class.

By introducing non-Sybil agents externally (i.e. $\Gamma_0 \neq \emptyset$), STM can contribute the occurrence of some "isolated" non-Sybil agents to buyers in the dominator sequence.

Recall the example in Figure 4, and we can see that every buyer other than the item's winner has zero utility. The following lemma shows that this is not a fluke. In essence, all possible profits of the brokers are taxed by the seller. The proof can also be found in Appendix E in the full version [2].

Lemma 4.2. *In* STM *when* $\Gamma_0 = \emptyset$*, every buyer, except the item winner, has a payment of zero, and thus zero utility.*

5 SYBIL CLUSTER MECHANISM

In STM, we reward the brokers for their contribution to introducing agents in $\Gamma(\boldsymbol{\theta}')$. However, when Γ_0 is empty, a broker cannot bring a graphical non-Sybil agent on her own. This leads to zero profit for all brokers, as shown in Lemma 4.2, and their incentive to invite other agents is weak.

To create a positive incentive without sacrificing Sybil-proofness, we propose a clustering process that removes edges from $G_s(\boldsymbol{\theta}')$ while keeping $\Gamma(\boldsymbol{\theta}')$ unchanged. This is used to reward brokers who introduce non-Sybil agents.

5.1 Clustering

Definition 5.1 (Sybil clusters). For every $x \in \Gamma(\boldsymbol{\theta}')$, we define its *Sybil cluster* K_x as below:

The cluster K_x contains vertex t if and only if there is a path from x to t on $G_s(\boldsymbol{\theta}')$ that does not contain any vertex in $\Gamma(\boldsymbol{\theta}')$ other than x itself.

The vertex x is called the *root* of K_x. The Sybil cluster rooted at x includes all vertices that might be the Sybil identities of x. All non-Sybil vertices in $\Gamma(\boldsymbol{\theta}')$ other than x are not in K_x, and all vertices in $K_x \setminus \{x\}$ are excluded from $\Gamma(\boldsymbol{\theta}')$.

The clusters $\{K_x \mid x \in \Gamma(\boldsymbol{\theta}')\}$ get the name because they form a partition of $V(G_s(\boldsymbol{\theta}'))$, which is shown in the following lemma. Its proof is deferred to Appendix D in the full version [2].

Lemma 5.2. *Sybil clusters are disjoint, and every vertex* t *in* $V(G_s(\boldsymbol{\theta}'))$ *belongs to some Sybil cluster* K_x*.*

Using Sybil clusters, one can prove that $\Gamma(\boldsymbol{\theta}')$ is the maximal set of guaranteed non-Sybil vertices.

Lemma 5.3. *Any vertex* t *in* $G_s(\boldsymbol{\theta}') \setminus \Gamma(\boldsymbol{\theta}')$ *may be a Sybil identity of some other vertex in* $\Gamma(\boldsymbol{\theta}')$*.*

Proof. Given a report profile $\boldsymbol{\theta}'$, we can compute $\Gamma(\boldsymbol{\theta}')$ and the Sybil clusters by definition. For any $t \notin \Gamma(\boldsymbol{\theta}')$, there exists $x \in \Gamma(\boldsymbol{\theta}')$ such that $t \in K_x$ from Lemma 5.2. Let $\tilde{r}(x) = \bigcup_{u \in K_x} r(u) \setminus K_x$, and $\tilde{\theta}_x = (\tilde{r}(x), v'_x)$. We can see that, under the true type profile $\tilde{\boldsymbol{\theta}} = (\tilde{\theta}_x, \boldsymbol{\theta}'_{-K_x})$, the agent x may create Sybil identities $\phi = K_x$ and make the report profile identical to $\boldsymbol{\theta}'$. This shows that t may be a Sybil identity of x. □

5.2 SCM Mechanism

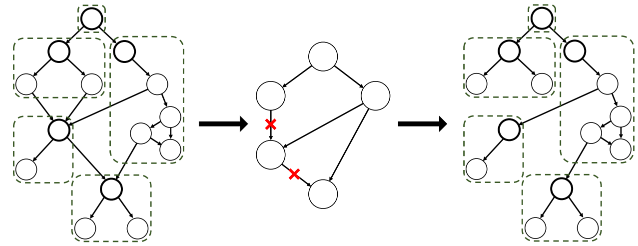

(a) (b) (c)

Figure 5: A visualization of SCM.

Sybil Cluster Mechanism (SCM)

(1) Given the reported type profile $\boldsymbol{\theta}'$ as input, we reconstruct a social network graph H with vertices in $\Gamma(\boldsymbol{\theta}')$. Formally, $H = (\Gamma(\boldsymbol{\theta}'), E(H))$, where

$$E(H) = \left\{(x, y) \mid \exists i \in K_x, j \in K_y \text{ s.t. } (i, j) \in E(G_s(\boldsymbol{\theta}'))\right\}.$$

(2) Sample a random shortest-path tree[3] of H with equal probability and denote it as T_H.

(3) We construct a subgraph $\hat{G}$ of $G_s(\boldsymbol{\theta}')$ using T_H. Formally, $\hat{G} = (V(G_s(\boldsymbol{\theta}')), E(\hat{G}))$ where $E(\hat{G})$ is defined as

$$E(G_s(\boldsymbol{\theta}')) \setminus \{(i, j) \mid i \in K_x, j \in K_y, (x, y) \in E(H) \setminus E(T_H)\}.$$

Specifically, edge (i, j) on graph $G_s(\boldsymbol{\theta}')$ is deleted if $i \in K_x, j \in K_y, x \neq y$ and $(x, y) \notin T_H$. All the remaining edges form a new graph $\hat{G}$.

(4) Perform STM with $G_s = \hat{G}$, $\Gamma = \Gamma(\boldsymbol{\theta}')$ rather than $G_s = G_s(\boldsymbol{\theta}')$, $\Gamma = \Gamma(\boldsymbol{\theta}')$ on the agents' reports.

In SCM, we remove some edges in $G_s(\boldsymbol{\theta}')$ according to the randomly selected shortest-path tree T_H and keep $\Gamma(\boldsymbol{\theta}')$ as graphical non-Sybil agents. The appearance of some vertices in Γ can be attributed to some brokers, thus increasing their profit.

The example in Figure 4(b) gives an example of the Sybil cluster mechanism. The clustering process and a possible edge-removing process are shown in Figure 5. Assuming that all buyers report their true type, SCM runs as follows:

The mechanism divides $V(G_s(\boldsymbol{\theta}))$ into five Sybil clusters $\{K_s, K_a, K_b, K_i, K_k\}$, where $K_s = \{s\}, K_a = \{a, c, d\}, K_b = \{b, e, f, g, h\}, K_i = \{i, j\}$, and $K_k = \{k, l, m\}$. The mechanism randomly picks a shortest-path tree T_H and constructs a subgraph $\hat{G}$. We only show the case when the mechanism picks the tree T_H as Figure 5(b), where the mechanism deletes edges (a, i) and (i, k). In this case, edges (c, i), (d, i) and (i, k) are removed from $G_s(\boldsymbol{\theta})$. With $\Gamma = \{s, a, b, i, k\}$, we perform STM on $\hat{G}$.

[3]For every vertex $x \in V(H)$, we denote the shortest-path length from s to it on graph H as $dis_x(H)$. A spanning tree T_H of H is a subgraph of H with $V(T_H) = V(H)$, which is also a directed tree. A spanning tree T_H is said to be a shortest-path tree if, for every vertex $x \in V(H)$, $dis_x(T_H) = dis_x(H)$. A uniformly distributed random shortest-path tree can be generated by independently selecting a parent y for each $x \neq s$, where y is selected from $\{y \mid dis_x(H) = dis_y(H) + 1, (x, y) \in E(H)\}$ with equal probability.

STM identifies the buyer with the highest bidder to be h and calculates the dominator sequence $C_h = \{c_0 = s, c_1 = b, c_2 = e, c_3 = f, c_4 = h\}$. After comparing the bids, we select $c_4 = h$ as the winner of the item.

For the payments, broker $c_1 = b$ pays $t_b = p_1 - q_1 = 19 - 19 = 0$, $c_2 = e$ gets $-t_e = q_2 - p_2 = 21 - 19 = 2$ units of money, and $c_3 = f$ gets $-t_f = 5$. The winner $c_4 = h$ pays $t_h = p_4 = 29$. The seller gets a revenue of 22.

Theorem 5.4. *Sybil cluster mechanism is IR, non-deficit, and Sybil-proof.*

We give an intuition why SCM retains Sybil-proofness here. Rigorous proof can be found in Appendix F of [2].

Assume an agent x performs a Sybil attack. If we fix the configuration of Sybil identities (the number of them and the social structure among them), the graph structure is fixed, and so is the clustering and the edge removal process. It is optimal for the agent to bid truthfully since the last step of SCM, i.e. STM, is IC. This is true for every configuration of Sybil identities.

Now the agent would pick a favorable configuration. If there is a descendant y connected to x through some Sybil identities, y would have a longer distance from the seller, thus less likely to be chosen in the subtree of x. It is favorable for x to diffuse to y directly without any Sybil identities in the middle. In this case, the Sybil identities contribute nothing to x, and can be removed. Thus, SCM is Sybil-Proof.

6 DISCUSSION

In this paper, we propose two Sybil-proof mechanisms, STM and SCM. In this section, we evaluate their performance on social welfare and revenue. Comparing our mechanism with the non-diffusion mechanism (i.e., NSP), other potential SP mechanisms and existing diffusion mechanisms (e.g., IDM, VCG) which are not SP, we raise three key questions.

(1) Do our diffusion mechanisms have better performance than non-diffusion ones?
(2) Does STM or SCM achieve optimal social welfare and revenue among all SP mechanisms?
(3) Compared with existing diffusion mechanisms, how much do our mechanisms sacrifice to achieve Sybil-proofness?

In this section, we conduct theoretical and experimental analysis to answer these questions. To eliminate the external effect, we assume that $\Gamma_0 = \emptyset$ from now on.

6.1 Comparison

We use $\mathcal{SW}^{\mathcal{M}}(\boldsymbol{\theta})$ and $\mathcal{R}^{\mathcal{M}}(\boldsymbol{\theta})$ to denote the social welfare and revenue of the mechanism $\mathcal{M}$ under $\boldsymbol{\theta}$ respectively. We have

$$\mathcal{SW}^{\mathcal{M}}(\boldsymbol{\theta}) \triangleq \sum_{x\in N} \pi_x^{\mathcal{M}}(\boldsymbol{\theta}) \cdot v_x.$$

Recall that we have defined $\mathcal{R}(\boldsymbol{\theta}) = \sum_{i\in N} t_i(\boldsymbol{\theta})$ in Section 2.1.

The following theorem shows that both of our mechanisms outperform the non-diffusion NSP mechanism. Under our mechanisms, agents' invitations indeed benefit the seller and the society.

Theorem 6.1. *For all possible type profile $\boldsymbol{\theta}$, we have*

$$\mathcal{R}^{\mathrm{STM}}(\boldsymbol{\theta}) \geq \mathcal{R}^{\mathrm{SCM}}(\boldsymbol{\theta}) \geq \mathcal{R}^{\mathrm{NSP}}(\boldsymbol{\theta}),$$
$$\mathcal{SW}^{\mathrm{STM}}(\boldsymbol{\theta}) \geq \mathcal{SW}^{\mathrm{SCM}}(\boldsymbol{\theta}) \geq \mathcal{SW}^{\mathrm{NSP}}(\boldsymbol{\theta}).$$

We are curious whether STM achieves higher social welfare and revenue than all SP mechanisms. However, we'll show in Section 6.2 that none of SP mechanisms always has optimal social welfare and revenue.

The following theorems qualitatively examine the cost of Sybil-proofness. In Theorem 6.2, we find that STM achieve better revenue than the most cited diffusion auction, IDM [11]. However, social welfare is sacrificed to achieve Sybil-proofness. Theorem 6.3 reflects that there is no clear-cut comparison of the seller's revenue between SCM and IDM, or between SCM and VCG.

Theorem 6.2. *For any possible type profile $\boldsymbol{\theta}$, we have*

$$\mathcal{R}^{\mathrm{STM}}(\boldsymbol{\theta}) \geq \mathcal{R}^{\mathrm{IDM}}(\boldsymbol{\theta}) \geq \mathcal{R}^{\mathrm{VCG}}(\boldsymbol{\theta})$$
$$\mathcal{SW}^{\mathrm{SCM}}(\boldsymbol{\theta}) \leq \mathcal{SW}^{\mathrm{STM}}(\boldsymbol{\theta}) \leq \mathcal{SW}^{\mathrm{IDM}}(\boldsymbol{\theta}) \leq \mathcal{SW}^{\mathrm{VCG}}(\boldsymbol{\theta}).$$

Theorem 6.3. *There exist two report profiles $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2$, such that*

$$\mathcal{R}^{\mathrm{SCM}}(\boldsymbol{\theta}_1) > \mathcal{R}^{\mathrm{IDM}}(\boldsymbol{\theta}_1), \mathcal{R}^{\mathrm{SCM}}(\boldsymbol{\theta}_1) > \mathcal{R}^{\mathrm{VCG}}(\boldsymbol{\theta}_1),$$
$$\mathcal{R}^{\mathrm{SCM}}(\boldsymbol{\theta}_2) < \mathcal{R}^{\mathrm{IDM}}(\boldsymbol{\theta}_2), \mathcal{R}^{\mathrm{SCM}}(\boldsymbol{\theta}_2) < \mathcal{R}^{\mathrm{VCG}}(\boldsymbol{\theta}_2).$$

The proofs of Theorem 6.1, 6.2, and 6.3 can be found in Appendix G.

6.2 Worst-Case Efficiency Analysis and (No) Optimality

In this subsection, we conduct worst-case analysis on SP mechanisms to explore the optimality of social welfare and revenue. We consider the concept of worst-case efficiency ratio, which is adopted from previous work [7] to measure the social welfare of Sybil-proof combinatorial auctions in the worst case. The worst-case efficiency ratio of $\mathcal{M}$ indicates the ratio of $\mathcal{M}$'s social welfare and the optimal social welfare in the worst-case input.

Definition 6.4. Given a type profile $\boldsymbol{\theta}$, the optimal social welfare $\mathcal{SW}^*(\boldsymbol{\theta})$ is defined to be the highest private value $\max_{x\in V(G_s)} v_x$. The worst-case efficiency ratio of a mechanism $\mathcal{M}$ is defined as follows:

$$\inf_{\boldsymbol{\theta}} \frac{\mathcal{SW}^{\mathcal{M}}(\boldsymbol{\theta})}{\mathcal{SW}^*(\boldsymbol{\theta})}.$$

Theorem 6.5. *The worst-case efficiency ratio of any non-deficit, IR, and Sybil-proof diffusion auction mechanism is zero.*

The above theorem shows that the social welfare of every Sybil-proof mechanism is far below the social optimum in some cases. Its proof is included in Appendix H of [2].

Because every SP mechanism is sufficiently bad compared to social optimum, it is natural to compare their social welfare relative to other SP mechanisms. However, this further impossibility result indicates that every SP mechanism would perform extremely worse than another SP mechanism in some cases. Therefore, we cannot find any optimal diffusion auction, even when the optimality is relative to each other. The proof can also be found in Appendix H of [2].

Theorem 6.6. *For any non-deficit, SP, and IR diffusion auction mechanism $\mathcal{M}$, and for any $\varepsilon > 0$, there exists another non-deficit, SP, and IR diffusion auction mechanism $\mathcal{M}'$ such that*

$$\inf_{\theta} \frac{\mathcal{SW}^{\mathcal{M}}(\theta)}{\mathcal{SW}^{\mathcal{M}'}(\theta)} < \varepsilon.$$

We can derive a similar result in terms of the seller's revenue.

Theorem 6.7. *For any non-deficit, SP, and IR diffusion auction mechanism $\mathcal{M}$, and for any $\varepsilon > 0$, there exists another non-deficit, SP, and IR diffusion auction mechanism $\mathcal{M}'$ such that*

$$\inf_{\theta} \frac{\mathcal{R}^{\mathcal{M}}(\theta)}{\mathcal{R}^{\mathcal{M}'}(\theta)} < \varepsilon.$$

The theorems above indicate that all SP mechanisms have extremely low social welfare and revenue compared to some other SP mechanisms. These impossibility results are surprising and show the drastic difference between diffusion mechanisms and traditional auctions.

6.3 Experiments

Despite the qualitative comparison results in Section 6.1, we still wonder how much our mechanisms are better than NSP, and how much social welfare and revenue is sacrificed for Sybil-proofness. Therefore, we conduct simulations to analyze the performance of mechanisms in the average case. Such experiments have never been performed on diffusion auctions in previous literature, so we have to be innovative in the settings.

To test the diffusion auction mechanisms, we must specify the private value vector of buyers and the social network structure. For simplicity, we assume the private values are drawn i.i.d. from a uniform distribution on $[0, 1]$. The graph structure in diffusion auctions can be highly complex. Since diffusion auctions are held on social networks, we take inspirations from network science to create distributions for graphs. Price's model [3] is a simple and classical model for directed networks, used to describe various scale-free networks in the real world [14]. It generates a graph of n vertices, each with a degree of m.

The mechanisms are tested with graphs with $n = 100$ vertices, and the density can be controlled by changing the parameter m. For each m, 1,000 inputs are generated as specified above. We assume that all agents act truthfully in the experiment. Five mechanisms are tested: NSP, STM, SCM, IDM and VCG. We calculate and analyze their social welfare and revenue. The results are visualized with box plots in Figure 6.

We have the following observations. Firstly, our mechanisms achieve significantly higher social welfare and revenue than the non-diffusion NSP mechanism. Secondly, the average-case social welfare distribution of either STM or SCM is very close to the social optimum (VCG), especially when the graph is denser. Thirdly, STM has the highest revenue, which is consistent with theoretical analysis. Finally, seller's revenue of SCM is slightly lower than IDM, and higher than VCG.

Experimental results indicate that our diffusion mechanisms have significantly better performance than NSP, and we do not sacrifice seller's revenue and social welfare much to achieve Sybil-proofness.

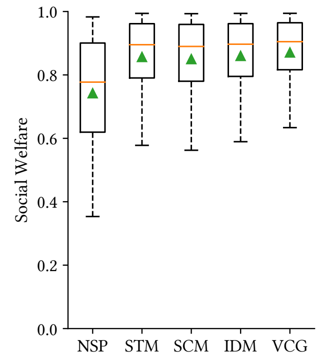

(a) Social welfare, $m = 3$.

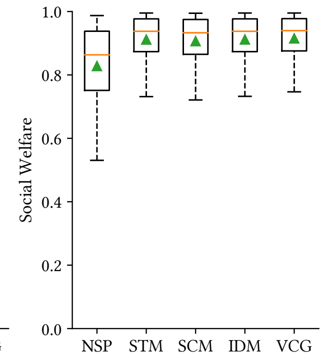

(b) Social welfare, $m = 5$.

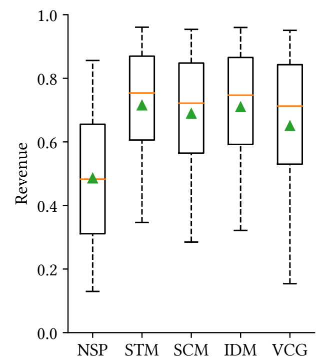

(c) Seller's revenue, $m = 3$.

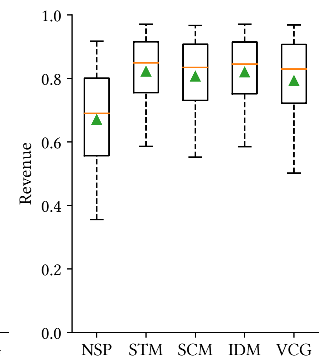

(d) Seller's revenue, $m = 5$.

Figure 6: The welfare and revenue distribution of five mechanisms on graphs of different densities. Orange line: median; green triangle: mean; box: 25% to 75%; whisker: 5% to 95%.

7 CONCLUSIONS

In this paper, we study an important issue in diffusion auctions, the Sybil attack. We find that previous diffusion mechanisms are vulnerable to Sybil attacks. We have proposed two novel solutions, STM and SCM, and proved that they are incentive compatible and Sybil-proof. We further discuss the social welfare and revenue of these two mechanisms. Theoretical analysis and experiments indicate that STM and SCM achieve Sybil-proofness with little sacrifice in the social welfare and revenue.

We also conduct worst-case analysis on all Sybil-proof diffusion mechanisms. We prove negative conclusions that the social welfare and revenue of every SP mechanism is far below some other SP mechanism in some cases.

Our work raises many open problems in the domain of Sybil-proof diffusion auctions. Firstly, how to develop Sybil-proof diffusion mechanisms for selling multiple items? Secondly, is there any other effective way to achieve Sybil-proofness? Thirdly, since we can't pick out the optimal Sybil-proof diffusion mechanism in the worst case, can we develop other methods to compare SP mechanisms? Or can we only compare a subset of all SP mechanisms to avoid such negative conclusions? Furthermore, how to reward the intermediate buyers fairly is also worth consideration.

ACKNOWLEDGMENTS

This research was partially supported by the National Natural Science Foundation of China (Grant No. 62172012) and Science and Technology Commission of Shanghai Municipality (No. 22ZR1442200).

REFERENCES

[1] Jeremy Bulow and Paul Klemperer. 1996. Auctions Versus Negotiations. *The American Economic Review* 86, 1 (1996), 180–194.

[2] Hongyin Chen, Xiaotie Deng, Ying Wang, Yue Wu, and Dengji Zhao. 2022. Sybil-Proof Diffusion Auction in Social Networks. *arXiv preprint arXiv:2211.01984* (2022).

[3] Derek J. de Solla Price. 1965. Networks of Scientific Papers. *Science* 149, 3683 (1965), 510–515. https://doi.org/10.1126/science.149.3683.510

[4] Xuewen Dong, Yuanyu Zhang, Yuanxiong Guo, Yanmin Gong, Yulong Shen, and Jianfeng Ma. 2021. PRAM: a Practical Sybil-Proof Auction Mechanism for Dynamic Spectrum Access with Untruthful Attackers. *IEEE Transactions on Mobile Computing* (2021).

[5] John R Douceur. 2002. The Sybil attack. In *International workshop on peer-to-peer systems*. Springer, 251–260.

[6] Yuhang Guo and Dong Hao. 2021. Emerging Methods of Auction Design in Social Networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 4434–4441. https://doi.org/10.24963/ijcai.2021/605

[7] Atsushi Iwasaki, Vincent Conitzer, Yoshifusa Omori, Yuko Sakurai, Taiki Todo, Mingyu Guo, and Makoto Yokoo. 2010. Worst-Case Efficiency Ratio in False-Name-Proof Combinatorial Auction Mechanisms. In *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems: Volume 1* (Toronto, Canada) *(AAMAS '10)*. Richland, SC, 633–640.

[8] Takehiro Kawasaki, Nathanaël Barrot, Seiji Takanashi, Taiki Todo, and Makoto Yokoo. 2020. Strategy-proof and non-wasteful multi-unit auction via social network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2062–2069.

[9] Thomas Lengauer and Robert Endre Tarjan. 1979. A Fast Algorithm for Finding Dominators in a Flowgraph. *ACM Trans. Program. Lang. Syst.* 1, 1 (Jan. 1979), 121–141. https://doi.org/10.1145/357062.357071

[10] Bin Li, Dong Hao, Dengji Zhao, and Makoto Yokoo. 2019. Diffusion and Auction on Graphs. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 435–441. https://doi.org/10.24963/ijcai.2019/62

[11] Bin Li, Dong Hao, Dengji Zhao, and Tao Zhou. 2017. Mechanism Design in Social Networks. *Proceedings of the AAAI Conference on Artificial Intelligence* 31, 1 (Feb. 2017). https://ojs.aaai.org/index.php/AAAI/article/view/10605

[12] Pengfei Liu, Xiaohan Wang, Xiangqian Che, Zhaoqun Chen, and Yuantao Gu. 2014. Defense against Sybil attacks in directed social networks. In *2014 19th International Conference on Digital Signal Processing*. IEEE, 239–243.

[13] Noam Nisan (Ed.). 2007. *Algorithmic game theory*. Cambridge University Press, Cambridge ; New York. OCLC: ocn122526907.

[14] Riitta Toivonen, Lauri Kovanen, Mikko Kivelä, Jukka-Pekka Onnela, Jari Saramäki, and Kimmo Kaski. 2009. A comparative study of social network models: Network evolution models and nodal attribute models. *Social Networks* 31 (Oct. 2009), 240–254. https://doi.org/10.1016/j.socnet.2009.06.004

[15] Bimal Viswanath, Ansley Post, Krishna P Gummadi, and Alan Mislove. 2010. An analysis of social network-based Sybil defenses. *ACM SIGCOMM Computer Communication Review* 40, 4 (2010), 363–374.

[16] Binghui Wang, Jinyuan Jia, Le Zhang, and Neil Zhenqiang Gong. 2020. Structure-Based Sybil Detection in Social Networks via Local Rule-Based Propagation. arXiv:1803.04321 [cs]

[17] Ling Xu, Satayapiwat Chainan, Hiroyuki Takizawa, and Hiroaki Kobayashi. 2010. Resisting Sybil Attack by Social Network and Network Clustering. In *2010 10th IEEE/IPSJ International Symposium on Applications and the Internet*. 15–21. https://doi.org/10.1109/SAINT.2010.32

[18] Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. 2004. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior* 46, 1 (2004), 174–188.

[19] Haifeng Yu, Phillip B Gibbons, Michael Kaminsky, and Feng Xiao. 2008. Sybil-Limit: A near-optimal social network defense against Sybil attacks. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 3–17.

[20] Tianyi Zhang, Dengji Zhao, Wen Zhang, and Xuming He. 2021. Fixed-price diffusion mechanism design. In *Pacific Rim International Conference on Artificial Intelligence*. Springer, 49–62.

[21] Wen Zhang, Dengji Zhao, and Hanyu Chen. 2020. Redistribution mechanism on networks. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems* (Auckland, New Zealand) *(AAMAS '20)*. Richland, SC, 1620–1628.

[22] Dengji Zhao, Bin Li, Junping Xu, Dong Hao, and Nicholas R. Jennings. 2018. Selling Multiple Items via Social Networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) *(AAMAS '18)*. Richland, SC, 68–76.