Collusion-Proof and Sybil-Proof Reward Mechanisms for Query Incentive Networks

Youjia Zhang, Pingzhong Tang

IIIS, Tsinghua University
zhangyou19@mails.tsinghua.edu.cn, kenshinping@gmail.com

Abstract

This paper explores reward mechanisms for a query incentive network in which agents seek information from social networks. In a query tree issued by the task owner, each agent is rewarded by the owner for contributing to the solution, for instance, solving the task or inviting others to solve it. The reward mechanism determines the reward for each agent and motivates all agents to propagate and report their information truthfully. In particular, the reward cannot exceed the budget set by the task owner. However, our impossibility results demonstrate that a reward mechanism cannot simultaneously achieve Sybil-proof (agents benefit from manipulating multiple fake identities), collusion-proof (multiple agents pretend as a single agent to improve the reward), and other essential properties. In order to address these issues, we propose two novel reward mechanisms. The first mechanism achieves Sybil-proof and collusion-proof, respectively; the second mechanism sacrifices Sybil-proof to achieve the approximate versions of Sybil-proof and collusion-proof. Additionally, we show experimentally that our second reward mechanism outperforms the existing ones.

Introduction

There is an old proverb that states, "Many hands make light work." In other words, the more people involved in a task, the quicker and easier it will be completed. With the growing popularity of social media, many online platforms (e.g., Quora and Stack Overflow) provide the opportunity for people to ask a question online. As opposed to traditional search engines (e.g., Google and Bing), Q&A platforms allow users to post questions and have them answered by other users rather than a central system that provides several related solutions. A major advantage of online Q&A platforms is that the task is handled by humans, who are capable of solving more rare questions. Additionally, it is unnecessary for the questioner to simplify the question in order to make it understandable by a machine.

An important open question in such query models is developing a reward mechanism that ensures the query is propagated successfully and the solution is provided if one exists. Many researchers have discussed and applied such a reward mechanism in various fields, such as peer-to-peer file-sharing systems (Golle et al. 2001), blockchains (Brünjes et al. 2020), and marketing (Drucker and Fleischer 2012) and task collaboration (Shapley 1952).

In this paper, we mainly focus on reward mechanisms for the answer querying system, in which only one agent is chosen to solve the problem. When implementing such reward mechanisms, it is critical to understand how to motivate agents to spread the information successfully and answer the question honestly. In addition, the total reward should not exceed the budget of the task owner (questioner). Apart from these basic requirements, there are two other significant challenges to be overcome.

The first challenge is Sybil-proof, avoiding agents manipulating multiple false identities to improve the overall reward. For example, an agent who knows the solution can create multiple accounts, and his other accounts can invite him to do so. As the task owner does not know the information of each participant, all these accounts are rewarded. Moreover, Sybil attacks not only increase the monetary costs of the questioner but also delay the time it takes to get a solution. This issue has been discussed in various literature, including (Babaioff et al. 2012; Drucker and Fleischer 2012; Chen et al. 2013; Lv and Moscibroda 2016; Zhang, Zhang, and Zhao 2021).

The second challenge is collusion-proof, preventing agents from colluding with each other to gain more rewards. For instance, an agent who knows the solution could report it to his parents rather than solve it directly, as such action may increase the overall reward. Additionally, his parents can also take the same action. In consequence, the questioner spends more time solving the problem. A question of this nature is not well studied in the query incentive network (Nath et al. 2012; Zhang, Zhang, and Zhao 2021), but has been extensively discussed in other fields, for example, auction design (Laffont and Martimort 1997; Che and Kim 2006; Marshall and Marx 2007).

However, it is impossible for a reward mechanism to simultaneously achieve both Sybil-proof (SP) and collusion-proof (CP) along with a set of other desirable properties (as illustrated in Section). Hence, we not only formally define these properties but also provide approximate versions of SP and CP. In this paper, we propose two novel families of reward mechanisms that are implementable in dominant strategies. The first mechanism is inspired by a geometric mechanism (Emek et al. 2011). In addition to the basic properties, it also achieves Sybil-proof and collusion-proof, respectively. The second mechanism sacrifices Sybil-proof in order to achieve approximate versions of Sybil-proof and collusion-proof simultaneously. Furthermore, both mechanisms provide a fair allocation/outcome, which is not well-defined in the query incentive network (Rahwan et al. 2014). Finally, our numerical experiments indicate that our second reward mechanism performs better than the existing ones.

Literature Review

Our research contributes to three streams of literature: query incentive networks, Sybil-proof, and collusion-proof. Below is a brief overview of the research areas closest to our own.

The seminal paper (Kleinberg and Raghavan 2005) was the first to introduce the query incentive model, which treated the query network as a simple branching process. They found that the reward function exhibited threshold behavior based on the branching parameter b. Apart from the fixed-payment contracts, (Cebrian et al. 2012) studied the winning team of the Red Balloon Challenge (Pickard et al. 2011) and proposed a split contract-based mechanism, which is robust to agents' selfishness.

It should be noted that the aforementioned works have not been examined in the context of Sybil attacks. The idea of Sybil attacks was first introduced by (Douceur 2002). (Douceur and Moscibroda 2007) formalized a set of desirable properties and proposed the Lottery Tree mechanism, which is Sybil-proof and motivates agents to join in a P2P system. (Emek et al. 2011; Drucker and Fleischer 2012) studied Sybil-proof mechanisms for multi-level marketing in which agents buy a product and are rewarded for successful referrals. (Babaioff et al. 2012) studied a similar problem and proposed a reward scheme for the Bitcoin system. (Lv and Moscibroda 2015, 2016) introduced a reward mechanism in crowdsourcing or human tasking systems. (Chen et al. 2013) proposed a Sybil-proof Mechanism for query incentive networks with the Nash equilibrium implementation, while (Zhang, Zhang, and Zhao 2021) considered the dominant strategy.

As we mentioned before, collusion-proof is not well-studied in query incentive networks. (Chen et al. 2013) left it as future work, and (Nath et al. 2012; Zhang, Zhang, and Zhao 2021) discussed the impossibility of a reward mechanism to achieve Sybil-proof and collusion-proof simultaneously. Specifically, (Nath et al. 2012) proposed a collusion-proof reward mechanism. Many other works such as (Laffont and Martimort 1997; Che and Kim 2006; Marshall and Marx 2007) belong to this category. According to related studies on auction design, this research area is promising and deserves further investigation.

Preliminaries and Model

Consider the problem for a task owner (she) r employing a set of agents to solve the task (e.g., Red-balloon type of challenges and InnoCentive). Each agent (he) can solve the job by himself or ask his friends for assistance. Intuitively, the query process works as follows:

1. The task owner r announces the task and the reward scheme and propagates the information to her friends.
2. Upon receiving the query, her friends decide whether to do the task if they can and whether to continue propagating the query to their friends for assistance.
3. Once the task is solved, the agents are rewarded according to the rules specified in Step (1).

Following the query process, a query tree rooted at r is built; we denote such a tree as $T_r = (V, E)$, where V is the set of all agents, including the task owner r, and edge $e(i, j) \in E$ means that agent i informs the task to agent j. We also use the standard tree notions of *parent* and *child* in their natural sense; the parents and children of agent i are denoted as p_i and c_i, respectively. Note that an agent i is reachable from the root r if all his ancestors are in the query tree T_r and decide to propagate the query. Furthermore, we assume that both owner and agents have zero cost to inform others in our main model for analytical brevity.

In such a process, agents are asked to report two pieces of information, the response to the task $resp_i \in \{0, 1\}$ (1 for answering, otherwise, 0) and the set of children $c_i \subseteq V \setminus r$. Accordingly, an agent i's action in the mechanism is defined as $\theta'_i = (resp'_i, c'_i)$. Notice that agents cannot propagate the task to the non-existing child; therefore, $c'_i \subseteq c_i$.

Definition 1. *Given a report file θ' of all agents, let the tree generated from θ' be $T_r(\theta') = (V', E') \subseteq T_r$, where $V' \in V$ and $E' \subseteq E$.*

Given the above setting, the task owner aims to design a reward mechanism that determines how the task is solved and how the rewards are distributed among the rooted tree $T_r(\theta')$. (Note that agents cannot invite those who are already in the query. Hence, it is a rooted tree rather than a graph.) Furthermore, the mechanism only rewards agents who have made a positive contribution. Mathematically, the formal definition of such a mechanism is defined as follows.

Definition 2. *The Reward Mechanism M on the social network is defined by a task allocation path $f : T_r \to T_r$ (T_r is the structure of the tree rooted at r), and a reward function $x = (x_i)_{i \in f(T_r(\theta'))}$, where $x_i : \Theta \to \mathbb{R}_+$ and $\theta_i \in \Theta$ is the type profile.*

In this paper, we assume the task owner r chooses only one agent to do the task. Such an assumption has no significant bearing on our results but simplifies analysis and exposition. In addition, when there exist multiple agents who can solve the task, the agent is selected via the shortest path to the root r. Since creating fake accounts increases the distance from the root to the solution, to some extent, such a task allocation rule can limit Sybil attacks (Chen et al. 2013).

Definition 3. *Given the agents profile type θ, for each $resp(i) = 1$, we define the shortest path from the root r to agent i as $P_i = \{r, a_1, a_2, ..., i\}$. The task is allocated to the agent with $P = \min_i P_i$, for all $resp(i) = 1$. If there exist multiple agents with P, then they are selected randomly.*

Properties

In this section, we define a set of important properties that a reward mechanism M on the social network should satisfy. All these properties are similar and inspired by the Lottery Tree (Douceur and Moscibroda 2007), the multi-level marketing (Emek et al. 2011), and query networks (Zhang, Zhang, and Zhao 2021); some of them are generalized to our query model.

Formally, we provide the definitions for a mechanism to be incentive compatible and individually rational. Since agents are rewarded based on their contributions, agents out of the path $P = f(T_r(\theta'))$ are not rewarded. We mainly focus our study on the agents in the path P and introduce more strict definitions of IR and IC.

Definition 4. *The Reward Mechanism M is*

- ***Profitable Opportunity (PO)**, if $x(\theta_i) > 0$,*
- ***Incentive Compatible (IC)**, if $x(\theta_i) \geq x(\theta_i')$,*

for all agents $i \in P \setminus r$ and $\theta_i, \theta_i' \in \Theta$.

PO (a.k.a. strongly IR) ensures all agents in the path P are rewarded by the owner r for participating in the mechanism. IC promises all these agents to do the task truthfully and propagate it to all children.

As mentioned before, agents receive rewards for their contributions to completing a task, which may lead to Sybil attacks. For example, an agent who knows the answer to a task could create a fake account and use it to complete the task, receiving a greater reward than if they had reported truthfully.

Definition 5. *The Reward Mechanism M is **λ-Sybil-proof (λ-SP)** ($\lambda \in \mathbb{N}$) if for all agents in the winning path Γ_w and $\lambda \geq 1$, it holds that*

$$R(i,n) \geq \sum_{k=0}^{\lambda} R(i+k, n+\lambda). \tag{1}$$

*If λ-Sybil proof holds for all positive integers λ, such a mechanism is **Sybil-proof (SP)**.*

A Sybil-proof (SP) mechanism ensures that any agent in the path P cannot benefit from pretending to be multiple agents. Moreover, a mechanism is λ-SP if it can prevent agents from improving rewards by creating λ Sybil attacks.

On the other hand, since agents are connected to each other, it may be easier for them to collude together to pretend as a single agent in order to get more rewards from the owner r. (Nath et al. 2012; Zhang, Zhang, and Zhao 2021) discuss such an observation in detail. The formal definition of collusion proof is given as follows.

Definition 6. *The Reward Mechanism M is **γ-collusion-proof (γ-CP)** ($\gamma \in \mathbb{N}$) if for all agents in the winning path Γ_w and $\gamma \geq 1$, it holds that*

$$R(i,n) \leq \sum_{k=0}^{\gamma-1} R(i+k, n+\gamma-1). \tag{2}$$

*If γ-collusion proof holds for all positive integers γ, such a mechanism is **collusion-proof (CP)**.*

A collusion-proof (CP) mechanism promises that all agents are worse off from forming a coalition group of any size. Similarly, a mechanism is γ-CP if agents cannot get more reward from creating a group with a size γ.

Definition 7. *The Reward Mechanism M is **budget balanced (BB)** if there exists a constant R_{max} such that*

$$\sum_{i \in P \setminus r} x(\theta_i) \leq R_{max},$$

and strongly BB if $\sum_{i \in P} x(\theta_i) = R_{max}$.

Budget balance ensures the total reward offered to participants never exceeds a budget set by the task owner. This is important to ensure the feasibility of the reward mechanism and to promote fairness and sustainability in crowdsourcing systems.

As discussed in the relevant literature (Cebrian et al. 2012), the reward of an agent i should depend on his child c_i and the corresponding depth i in the path P to credit the indirect referrals. For the direct referral, agent i should receive at least a certain fraction of his child's reward $x(i+1,n)$. Violation of these properties leads to a failure in propagating the query under certain conditions. Formally, these properties are defined as follows.

Definition 8. *The Reward Mechanism M is $\rho-$**split**, if $x(i,n) \geq \rho x(i+1,n)$, for any $i \in P \setminus r$ and $0 < \rho < 1$.*

Impossibility Theorem

So far, we have defined the set of desirable properties that a reward mechanism should satisfy. Our new mechanisms are developed in response to the following impossibility results.

One of the impossibility results suggests that a mechanism cannot satisfy profitable opportunity (PO), Sybil-proof (SP), and collusion-proof (CP) simultaneously. For example, (Nath et al. 2012) sacrificed Sybil-proofness to ensure collusion-proofness, whereas (Zhang, Zhang, and Zhao 2021) proposed a Sybil-proof mechanism, which is not collusion-proof.

Theorem 1. *For $n \geq 3$, there is no Reward Mechanism that can achieve PO, SP, and CP simultaneously.*

Proof. The proof is similar to those of (Nath et al. 2012; Zhang, Zhang, and Zhao 2021). Assume there exists a mechanism that satisfies all three properties.

If the mechanism is Sybil-proof (Definition 5), considering for $m = 1$, we have

$$\begin{aligned} x(i,n) &\geq x(i,n+1) + x(i+1,n+1) \\ &\geq x(i,n+2) + 2x(i+1,n+2) + x(i+2,n+2) \end{aligned} \tag{3}$$

Meanwhile, if the mechanism is collusion-proof (Definition 6), we derive that

$$\begin{aligned} x(i,n) &\leq \sum_{k=0}^{m=2} x(i+k, n+m) \\ &= x(i,n+2) + x(i+1,n+2) + x(i+2,n+2) \end{aligned} \tag{4}$$

For both Eqs. 3 & 4 to hold, we need $x(i+1,n+2) \leq 0$, which violates violates to the condition of PO (Definition 4). $\square$

The following theorem explains how n affects the reward of the agent who solves the task for an SP (CP) mechanism.

Theorem 2. *If the Reward Mechanism is PO and SP (CP), then $x(n,n)$ is non-increasing (non-decreasing) in n.*

Proof. (**PO, SP and $x(n,n)$ is non-increasing**) We prove this by contradiction. Assume the mechanism is PO, SP, and $x(n,n)$ is increasing in n. Since the mechanism is SP (Definition 5), for any integer λ, we have

$$x(n,n) \geq \sum_{k=0}^{\lambda} x(n+k, n+\lambda).$$

Since the mechanism is PO (Definition 4), $x(i+k, n+\lambda) > 0$, for any $k \in [0,\lambda]$. Therefore, we have $x(n,n) > x(n+\lambda, n+\lambda)$, which contradicts to $x(n,n)$ is increasing in n.

(**PO, CP and $x(n,n)$ is non-decreasing**) This can be proven in a similar way, and details are provided in the full version. □

Tree Dependent Geometric Mechanism

Inspired by the work (Emek et al. 2011) in multi-level marketing, in this section, we generalize a geometric reward mechanism into our model. Our proposed mechanism, **Tree (Topology) Dependent Geometric Mechanism (TDGM)**, achieves Sybil-proof and collusion-proof, respectively.

Mechanism 1 (Tree Dependent Geometric Mechanism)**.** *Given the agents' report file θ and the corresponding task allocation path P, the reward policy of the TDGM is defined as*

$$x(i,n) = \alpha^{n-i}\beta, \tag{5}$$

for all $i \in P \setminus r$, $0 < \alpha < 1$, and $0 < \beta \leq \frac{1-\alpha}{1-\alpha^n} R_{max}$.

The understanding of TDGM is intuitive. Each agent is rewarded according to their contribution to the task. Contributions can be divided into two categories: inviting others or solving the issue. For agent n who solves the task is rewarded β (we will characterize β later). For ancestors of agent n, they receive a certain fraction α of the rewards of their children.

Theorem 3. *TDGM satisfies IC, PO, BB, and $\alpha-$split.*

Proof. (**IC**) We start our proof from IC. Considering the agent n who solves the task, if he does not provide the solution, he would be either in the solution path P or not. If he is still in the path P, his reward is $\alpha^{n'-i}\beta$, where $n' > n$. Since $\alpha < 1$ and $n' > n$, $\alpha^{n'-i}\beta < \beta$ (β is the reward if he solves the task). Moreover, if he is not in the path P, he receives nothing from the owner. As a result, agent n is worse off from misreporting. Similarly, for other agents $i \in P \setminus \{r, n\}$. If these agents do not spread the information successfully, they may not be in the path P, which leads to no reward for them.

(**PO & $\alpha-$split**) Since $0 < \alpha < 1$ and $0 \leq n-i \leq n$, $\alpha^{n-i} > 0$. Moreover, $0 < \beta \leq \frac{1-\alpha}{1-\alpha^n} R_{max}$, we have $x(i,n) > 0$.

(**BB**) Summarizing all the rewards, we derive that

$$\begin{aligned}
\sum_{i=1}^{n} x(i,n) &= \sum_{i=1}^{n} \alpha^{n-i}\beta \\
&= \alpha^n \beta \cdot \sum_{i=1}^{n} \alpha^{-i} \\
&= \alpha^n \beta \cdot \frac{\alpha^{-1}(1-\alpha^{-n})}{1-\alpha^{-1}} \\
&= \frac{1-\alpha^n}{1-\alpha}\beta
\end{aligned}$$

Since $\beta \leq \frac{1-\alpha}{1-\alpha^n} R_{max}$, we have $\sum_{i=1}^{n} x(i,n) \leq R_{max}$. □

Theorem 3 reveals that TDGM satisfies most basic properties discussed in Section .

As our next step, we aim to characterize β and analyze the mechanism for Sybil-proof and collusion-proof properties. Referring to Theorem 2, we note that the reward function of an agent who completes the task should be dependent on the length of the winning path, denoted by n. Therefore, we define β as a function of n and the budget, R_{max}, such that $\beta = \beta(n, R_{max})$.

Lemma 1. *TDGM is SP if $\beta(n, R_{max})$ follows*

$$\beta(n, R_{max}) - \beta(n+m, R_{max}) \frac{1-\alpha^{m+1}}{1-\alpha} \geq 0, \tag{6}$$

and it is CP if $\beta(n, R_{max})$ follows

$$\beta(n, R_{max}) - \beta(n+m, R_{max}) \frac{1-\alpha^{m+1}}{1-\alpha} \leq 0, \tag{7}$$

for any $m \in \mathbb{N}^+$. Furthermore, $0 < \beta(n, R_{max}) \leq \frac{1-\alpha}{1-\alpha^n} R_{max}$ in order to satisfy the Budget Balanced condition.

Proof. To check the SP and CP of TDGM, we expand $\sum_{k=0}^{m} x(i+k, n+m)$ and compare it with $x(i,n)$,

$$\begin{aligned}
\sum_{k=0}^{m} x(i+k, n+m) &= \sum_{k=0}^{m} \alpha^{n+m-i-k} \beta(n+m, R_{max}) \\
&= \alpha^{n+m-i} \beta(n+m, R_{max}) \sum_{k=0}^{m} \alpha^{-k} \\
&= \alpha^{n-i} \frac{1-\alpha^{m+1}}{1-\alpha} \beta(n+m, R_{max}).
\end{aligned}$$

If TDGM is SP (Definition 5), we need $x(i,n) \geq \sum_{k=0}^{m} x(i+k, n+m)$. Hence,

$$\begin{aligned}
\alpha^{n-i}\beta(n, R_{max}) - \alpha^{n-i} \frac{1-\alpha^{m+1}}{1-\alpha} \beta(n+m, R_{max}) &\geq 0 \\
\beta(n, R_{max}) - \frac{1-\alpha^{m+1}}{1-\alpha} \beta(n+m, R_{max}) &\geq 0
\end{aligned}$$

The condition of CP can be derived in a similar way. □

Lemma 1 characterizes the reward function of the agent who solves the task in order to satisfy SP and CP, respectively. Furthermore, it supplements the result of Theorem 2 such that no reward mechanism can achieve Sybil-proof and collusion-proof simultaneously.

Interestingly, the Double Geometric Mechanism (DGM) (Zhang, Zhang, and Zhao 2021) which is Sybil-proof, is a sub-class of TDGM, and the δ-Geometric Mechanism (δ-GEOM) (Nath et al. 2012) which is collusion-proof also belongs to a TDGM with a certain condition. The proof can be found in the full version.

Corollary 1. *DGM (Zhang, Zhang, and Zhao 2021) and δ-GEOM (Nath et al. 2012) belong to a family of TDGM.*

Proposition 1. *Sybil-proof-TDGM is not collusion-proof, and collusion-proof-TDGM is not Sybil-proof.*

Proof. Assume the task solver creates m Sybil attacks, if the mechanism is SP and CP simultaneously, then

$$x(i,n) = \sum_{k=0}^{m} x(i+k, n+m)$$

$$\alpha^{n-i}\beta(n, R_{max}) = \sum_{k=0}^{m} \alpha^{n+m-i-k}\beta(n+m, R_{max})$$

$$\beta(n, R_{max}) = \beta(n+m, R_{max})\frac{1-\alpha^{m+1}}{1-\alpha}$$

$$\frac{\beta(n, R_{max})}{\beta(n+m, R_{max})} = \frac{1-\alpha^{m+1}}{1-\alpha}$$

(**SP-TDGM is non-CP**) Since TDGM is SP (Lemma 1), we have

$$\frac{\beta(n, R_{max})}{\beta(n+m, R_{max})} \geq \frac{1-\alpha^{m+1}}{1-\alpha}.$$

Note that $\frac{1-\alpha^{m+1}}{1-\alpha} > 1$ for $m > 0$, and it is increasing in m. Hence, $x(i,n) > \sum_{k=0}^{m} x(i+k, n+m)$ for any $m > 0$.

As a result, SP-TDGM is not CP.

(**CP-TDGM is non-SP**) Since TDGM is CP, by Theorem 2, we have $\beta(n, R_{max}) \leq \beta(n+m, R_{max})$. Hence,

$$\frac{\beta(n, R_{max})}{\beta(n+m, R_{max})} \leq 1 < \frac{1-\alpha^{m+1}}{1-\alpha},$$

for any integer $m > 1$ and $0 < \alpha < 1$.

Therefore, $x(i,n) < \sum_{k=0}^{m} x(i+k, n+m)$ for any $m > 0$.

As a result, CP-TDGM is not SP. ☐

Proposition 1 demonstrates that agents benefit from forming a coalition group under SP-TDGM, while all agents gain more rewards from applying Sybil attacks under CP-TDGM.

Generalized Contribution Reward Mechanism

Given the impossibility results in Section , a reward mechanism cannot achieve PO, BB, SP, and CP simultaneously. In Section , we introduce a family of geometric mechanisms which satisfy SP and CP, respectively. In particular, SP-TDGM is not CP, and CP-TDGM is not SP.

However, the ability to form a collusion group with a size larger than three is limited since an agent not only needs to make a deal with their parents and children but also with their ancestors and descendants. Therefore, we propose a more practical mechanism that can simultaneously achieve Sybil-proof and collusion-proof to some extent.

Mechanism 2 (Generalized Contribution Reward Mechanism)**.** *Given the agents' report file θ and the corresponding task allocation path P, the reward policy of GCRM is defined as*

$$x(i,n) = \frac{\alpha^{n-i}}{(1+\alpha)^i}\beta, \tag{8}$$

for all $i \in P \setminus r$, $0 < \alpha < 1$ and $\beta = R_{max}$.

In GCRM, each agent is rewarded according to his contribution to the solution, $(\frac{1}{\alpha(1+\alpha)})^i$. If $0 < \alpha \leq \frac{\sqrt{5}-1}{2}$, the task solver is considered to contribute the most and also rewarded the most, while if $\frac{\sqrt{5}-1}{2} \leq \alpha < 1$, the agent who is closest to the task owner is considered to contribute the most and receives the highest reward.

In this sense, α may be considered as a parameter to control the contribution between information propagation and task solution. Moreover, the reward of each agent is normalized by a factor α^n, which depends on the length of the path P.

The following theorem shows that GCRM satisfies the basic properties IC, PO, and BB without any restrictions. The proof is similar to that of Theorem 3 and is provided in the full version.

Theorem 4. *GCRM satisfies IC, PO, and BB.*

We then analyze the condition of α for GCRM to satisfy the property of ρ-split.

Lemma 2. *GCRM is ρ-split with $\alpha \in (0, \frac{\sqrt{5}-1}{2})$, where $\rho = \alpha(1+\alpha) \in (0, 1]$.*

Proof. According to the reward function of GCRM (Eq. 8), we have

$$\frac{x(i,n)}{x(i+1,n)} = \frac{\frac{\alpha^{n-1}}{(1+\alpha)^i}}{\frac{\alpha^{n-i-1}}{(1+\alpha)^{i+1}}} = (1+\alpha)\alpha = \rho$$

Since $0 < \alpha < \frac{\sqrt{5}-1}{2}$, by simple calculation, we derive that $0 < \rho \leq 1$, which follows the definition. ☐

So far, we have shown that GCRM is IC, PO, BB, and ρ-split under a certain condition. Following then, we study the performance of GCRM with regards to Sybil-proof and collusion-proof.

Lemma 3. *GCRM is λ^*-Sybil proof, where integer $\lambda^* > 1$. Furthermore,*

- *λ^* increases with α,*
- *agents maximize their reward by $\lceil \lambda' \rfloor$ Sybil attacks, where $\lambda' = \frac{\log(\frac{-\log(1+\alpha)}{\alpha(1+\alpha)(\log(\alpha))})}{\log(\alpha)\log(1+\alpha)}$, ($\lceil \cdot \rfloor$ denotes for the nearest integer function.)*
- *the reward more than $\lceil \lambda' \rfloor$ Sybil attacks is at most twice as the original one.*

Proof. We start by proving the mechanism is λ^*-SP, where $\lambda^* > 1$. According to the reward function of GCRM (Eq.8), we have

$$\begin{aligned}\sum_{k=0}^{\lambda} x(i+k, n+\lambda) &= \sum_{k=0}^{\lambda} \frac{\alpha^{n+\lambda-i-k}}{(1+\alpha)^{i+k}} \\ &= \frac{\alpha^{n+\lambda-i}}{(1+\alpha)^i} \sum_{k=0}^{\lambda} \frac{1}{(\alpha(1+\alpha))^k} \\ &= x(i,n)\alpha^\lambda \sum_{k=0}^{\lambda} \frac{1}{(\alpha(1+\alpha))^k} \\ &= x(i,n)\frac{1-\alpha^{\lambda+1}(1+\alpha)^{\lambda+1}}{(1+\alpha)^\lambda - \alpha(1+\alpha)^{\lambda+1}}\end{aligned}$$

For convenience, hereafter, we denote $f(\alpha, \lambda) = \frac{1-\alpha^{\lambda+1}(1+\alpha)^{\lambda+1}}{(1+\alpha)^\lambda - \alpha(1+\alpha)^{\lambda+1}}$.

By substituting $\lambda = 1$ into $f(\alpha, \lambda)$, we derive that

$$\begin{aligned}f(\alpha, 1) &= \frac{1-\alpha^2(1+\alpha)^2}{(1+\alpha)^1 - \alpha(1+\alpha)^2} \\ &= \frac{(1-\alpha(1+\alpha))(1+\alpha(1+\alpha))}{(1+\alpha)(1-\alpha(1+\alpha))} \\ &= \frac{1}{1+\alpha} + \alpha\end{aligned}$$

Since $\alpha \in (0,1)$, $\frac{1}{1+\alpha} + \alpha \in (1, \frac{3}{2})$. As a result, $x(i,n) < x(i,n+1) + x(i+1,n+1)$ is always true for GCRM, which is never a 1-SP mechanism.

Due to space limitations, the rest of the proof is provided in full version. □

Lemma 4. *GCRM is (λ^*+1)-Sybil-proof and λ^*-collusion proof.*

Proof. As a part of proof of Lemma 3, $f(\alpha, \lambda)$ decreases with λ if $\lambda > \lambda'$. In addition, λ^* is the maximum number of profitable Sybil attacks, hence, $(\lambda^* + 1)$ is the **smallest** integer that $f(\alpha, \lambda^* + 1) \leq 1$.

Then, we derive that

$$\begin{aligned}\sum_{k=0}^{\lambda^*-1} x(i+k, n+\lambda^*-1) &= f(\alpha, \lambda^*)x(i,n) \\ &\geq x(i,n).\end{aligned}$$

By Definition 6, the mechanism is λ^*-CP. □

Intuitively, under GCRM, agents always benefit from at least 1 Sybil attack. Moreover, as α increases, the maximum number of profitable Sybil attacks and the minimum size of the profitable group collusion increase. In contrast to CP-TDGM, under GCRM, there exists an optimal number (λ') of Sybil attacks to maximize the reward of each agent, and the new reward is at most twice the reward without Sybil attacks.

Despite GCRM can neither prevent Sybil attacks nor collusion, the total reward is upper bounded and decreases with a sufficiently large number of agents. Hence, the total reward never exceeds the budget of the questioner.

Lemma 5. *The total reward of GCRM maximized with $\lceil n' \rfloor$ agents, where $n' = \frac{\log(\frac{-\log(1+\alpha)}{\log(\alpha)})}{\log(\alpha)+\log(1+\alpha)}$.*

As a result, agents have to trade off between manipulating multiple identities and cooperating with other agents. For example, creating too many Sybil attacks reduces the reward. In the meanwhile, forming a large coalition group is also impractical.

Empirical Evaluations

In this section, we begin by empirically evaluating the performance of GCRM in terms of Sybil-proofness and collusion-proofness. Next, we compare GCRM with TDGM.

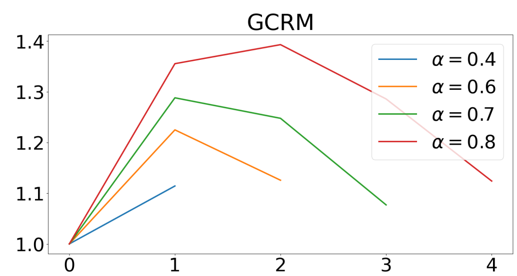

Figure 1: The ratio of new reward to original one after several Sybil attacks. The line end at λ^*. Note that it is unprofitable for $(\lambda^* + 1)$ Sybil attacks.

Figure 1 empirically evaluates how λ^* behaves and how reward improves after Sybil attacks with different α values under GCRM. It also complements the result of Lemma 3 by demonstrating that there exists a λ' value such that agents maximize their reward by creating $\lceil \lambda' \rfloor$ Sybil attacks, and both λ^* and $\lceil \lambda' \rfloor$ increase as α increases.

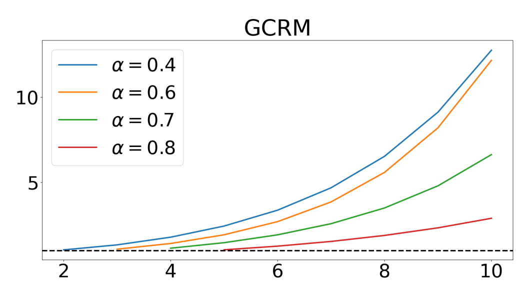

Figure 2: The ratio of new reward to the original one with different collusion size. Black dashed line represents $reward_{new} = reward_{original}$.

Similarly, Figure 2 presents an empirical evaluation of how the individual reward of an agent varies with different collusion sizes and α under GCRM. The results demonstrate that, as α increases, the minimum profitable collusion size also increases.

We use DGM (Zhang, Zhang, and Zhao 2021) and δ-GEOM (Nath et al. 2012) to represent SP-TDGM and CP-TDGM, respectively. However, α has a different meaning for DGM, δ-GEOM, and GCRM. To keep consistency, we restrict the parameter ρ to ensure that all three mechanisms are ρ-split ($\alpha_{DGM} = \frac{\rho}{1+\rho}$, $\delta = \rho$, and $\alpha_{GCRM} = \frac{\sqrt{1+4\rho}-1}{2}$).

We begin with analyzing the performance of δ-GEOM and GCRM on Sybil-proof, which is graphically shown in Figure 3.

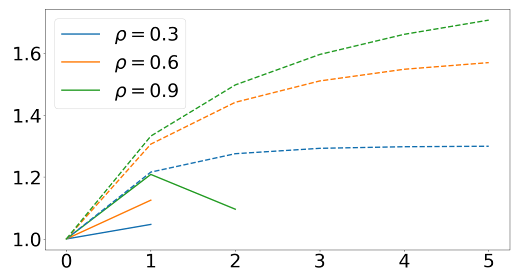

Figure 3: The ratio of new reward to original one after several Sybil attacks. Dashed lines represent δ-GEOM. Solid lines represent GCRM.

As we can see in Figure 3, agents always benefit from manipulating multiple identities under δ-GEOM. In addition, the more Sybil attacks they create, the more reward they gain. However, under GCRM, there exists a λ' such that agents maximize their reward by $\lceil \lambda' \rfloor$ Sybil attacks, and the number of profitable Sybil attacks is upper bounded. Furthermore, agents under GCRM receive less reward from Sybil attacks than under δ-GEOM.

Next, we evaluate the performance of DGM and GCRM on collusion-proof in Figure 4. As shown in Figure 4, GCRM performs better than DGM on collusion-proof. Regardless of the value of ρ, agents under GCRM receive less reward for forming a coalition group than agents under DGM. Furthermore, both Figures 3 & 4 supplement the results of Lemma 3, α (α is proportional to ρ) increases the minimum collusion size requirement to improve the reward.

As we proved in Theorems 3 & 4, all three mechanisms are budget balanced. These patterns are depicted in Figure 5. .

Conclusion

This paper studies a reward mechanism for a single solution task in a social network. We propose two classes of reward mechanisms to achieve desirable properties, such as PO, IC, BB, SP, and CP. Nevertheless, our impossibility results suggest that a reward mechanism cannot achieve all properties simultaneously. In particular, different questioners may consider different properties to be necessary.

Tree Dependent Geometric Mechanism (TDGM) is a classic geometric mechanism that can achieve Sybil-proof

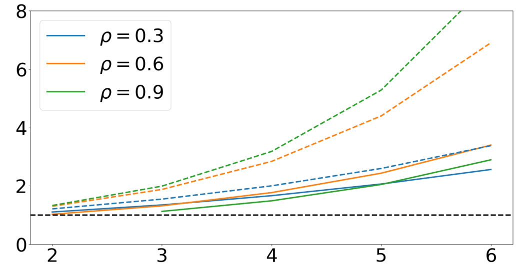

Figure 4: The ratio of new reward to the original one with different collusion size. Dashed lines represent DGM. Solid lines represent GCRM. Black dashed line represents $reward_{new} = reward_{original}$.

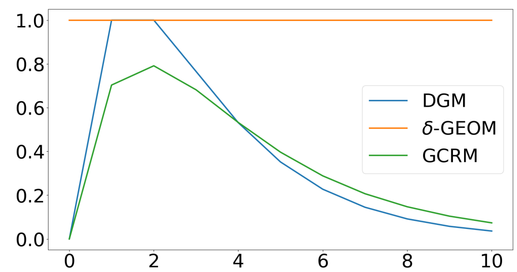

Figure 5: The ratio of total reward to budget with different size n and $\rho = 0.6$.

and collusion-proof, respectively. The second mechanism, the Generalized Contribution Reward Mechanism (GCRM), is a more flexible mechanism that sacrifices SP to achieve the approximate versions of SP and CP simultaneously. Specifically, GCRM is λ-SP and λ-CP. In other words, the optimal number of profitable Sybil attacks is limited and known to the questioner. The questioner adjusts the parameter α to trade-off between Sybil-proof and collusion-proof in order to improve the efficiency and the cost of obtaining the solution.

Despite the fact that our research provides some interesting insights into reward mechanisms in a single task allocation, numerous aspects remain to be explored. It is an interesting problem to consider referral costs, such that different agents have different costs to invite others. However, as mentioned in (Rochet and Choné 1998), mechanism design problems with multi-dimensional type distributions are challenging.

Acknowledgments

This work was supported by the National Key Research and Development Program of China under Grant No. 2020AAA0103401.

References

Babaioff, M.; Dobzinski, S.; Oren, S.; and Zohar, A. 2012. On bitcoin and red balloons. In *Proceedings of the 13th*

ACM conference on electronic commerce, 56–73.

Brünjes, L.; Kiayias, A.; Koutsoupias, E.; and Stouka, A.-P. 2020. Reward sharing schemes for stake pools. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, 256–275. IEEE.

Cebrian, M.; Coviello, L.; Vattani, A.; and Voulgaris, P. 2012. Finding red balloons with split contracts: robustness to individuals' selfishness. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 775–788.

Che, Y.-K.; and Kim, J. 2006. Robustly collusion-proof implementation. *Econometrica*, 74(4): 1063–1107.

Chen, W.; Wang, Y.; Yu, D.; and Zhang, L. 2013. Sybil-proof mechanisms in query incentive networks. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 197–214.

Douceur, J. R. 2002. The sybil attack. In *International workshop on peer-to-peer systems*, 251–260. Springer.

Douceur, J. R.; and Moscibroda, T. 2007. Lottery trees: motivational deployment of networked systems. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, 121–132.

Drucker, F. A.; and Fleischer, L. K. 2012. Simpler sybil-proof mechanisms for multi-level marketing. In *Proceedings of the 13th ACM conference on Electronic commerce*, 441–458.

Emek, Y.; Karidi, R.; Tennenholtz, M.; and Zohar, A. 2011. Mechanisms for multi-level marketing. In *Proceedings of the 12th ACM conference on Electronic commerce*, 209–218.

Golle, P.; Leyton-Brown, K.; Mironov, I.; and Lillibridge, M. 2001. Incentives for sharing in peer-to-peer networks. In *International workshop on electronic commerce*, 75–87. Springer.

Kleinberg, J.; and Raghavan, P. 2005. Query incentive networks. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, 132–141. IEEE.

Laffont, J.-J.; and Martimort, D. 1997. Collusion under asymmetric information. *Econometrica: Journal of the Econometric Society*, 875–911.

Lv, Y.; and Moscibroda, T. 2015. Incentive networks. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Lv, Y.; and Moscibroda, T. 2016. Fair and resilient incentive tree mechanisms. *Distributed Computing*, 29(1): 1–16.

Marshall, R. C.; and Marx, L. M. 2007. Bidder collusion. *Journal of Economic Theory*, 133(1): 374–402.

Nath, S.; Dayama, P.; Garg, D.; Narahari, Y.; and Zou, J. 2012. Mechanism design for time critical and cost critical task execution via crowdsourcing. In *International Workshop on Internet and Network Economics*, 212–226. Springer.

Pickard, G.; Pan, W.; Rahwan, I.; Cebrian, M.; Crane, R.; Madan, A.; and Pentland, A. 2011. Time-critical social mobilization. *Science*, 334(6055): 509–512.

Rahwan, T.; Naroditskiy, V.; Michalak, T.; Wooldridge, M.; and Jennings, N. R. 2014. Towards a Fair Allocation of Rewards in Multi-Level Marketing. *arXiv preprint arXiv:1404.0542*.

Rochet, J.-C.; and Choné, P. 1998. Ironing, sweeping, and multidimensional screening. *Econometrica*, 783–826.

Shapley, L. S. 1952. *A Value for N-Person Games*. Santa Monica, CA: RAND Corporation.

Zhang, Y.; Zhang, X.; and Zhao, D. 2021. Sybil-proof answer querying mechanism. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 422–428.