Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)
Sister Conferences Best Papers Track

Task Allocation on Networks with Execution Uncertainty (Extended Abstract)*

Yao Zhang, **Xiuzhen Zhang** and **Dengji Zhao**
Shanghai Engineering Research Center of Intelligent Vision and Imaging, ShanghaiTech University, Shanghai, China
{zhangyao1,zhangxzh1,zhaodj}@shanghaitech.edu.cn

Abstract

We study a single task allocation problem where each worker connects to some other workers to form a network and the task requester only connects to some of the workers. The goal is to design an allocation mechanism such that each worker is incentivized to invite her neighbours to join the allocation, although they are competing for the task. Moreover, the performance of each worker is uncertain, which is modelled as the quality level of her task execution. The literature has proposed solutions to tackle the uncertainty problem by paying them after verifying their execution. Here, we extend the problem to the network setting. We propose a new mechanism that guarantees that inviting more workers and reporting/performing according to her true ability is a dominant strategy for each worker. We believe that the new solution can be widely applied in the digital economy powered by social connections such as crowdsourcing.

1 Introduction

Task allocation is an important part of real-world applications such as crowdsourcing [Wu *et al.*, 2017; Goel *et al.*, 2014] and market supply [Dash *et al.*, 2007]. A common goal of task allocation is to find suitable workers to achieve a good performance at a low cost. Previous studies have made great progress in finding the best allocations under cases with a fixed number of workers. For example, the task requester seeks suitable workers in third-party platforms (e.g., Amazon Mechanical Turk) or holds a contest with attractive rewards [Chawla *et al.*, 2019]. Yet, such cases are less scalable due to the relatively fixed number of participants. Generally, we hope to involve more workers so that the task requester is capable of finding more suitable workers. Also, nowadays, people are connected with others via social networks. Therefore, a straightforward approach is to make full use of their connections such that we can involve more workers. The challenge remains such as workers are competitors for the task and they are unwilling to provide their connections.

More precisely, we consider a single-task allocation problem where the task is allocated to a single agent and the task performance of an agent is measured by the finished quality. Each agent has a cost to perform the task. Before conducting the tasks, agents are uncertain about their actual performance and only know their *probability distributions over the quality levels*, which is known as the execution uncertainty. Therefore, the requester also needs to take the execution uncertainty into accounts to allocate the task.

We propose the *PEV-based Diffusion Mechanism* to handle the above challenges one by one. Firstly, to solve the issue of agents unwilling to invite others, the proposed mechanism rewards them such that each agent will maximize her utility by inviting others. Then, the task requester is able to reach as many agents as possible. Secondly, to handle the execution uncertainty, the mechanism gives agents payoffs after they finished the task, which guarantees that agents will not misreport their abilities because their performance is verified. More importantly, previous studies focused on the uncertainty to finish a task but not they quality, introducing the probability of success (PoS) to describe the probability of an agent successfully completing the task [Porter *et al.*, 2008; Ramchurn *et al.*, 2009], e.g., 70% to fail and 30% to finish the task. In our setting, we define the probability of quality (PoQ), which represents the probability distribution on completion qualities, e.g., 30% to finish with a good quality, 20% to finish with a low quality, and 50% to totally fail.

1.1 Related Work

The social network is an effective medium to get access to more potential agents. Mechanism design in social networks has been widely utilized in auctions [Li *et al.*, 2020; Li *et al.*, 2022], answer querying [Tang *et al.*, 2011], social advertising [Li and Shiu, 2012] and influence maximization [Shi *et al.*, 2020]. An overview and prospect of all these topics on social networks is given by Zhao [2021; 2022]. In this paper, we are inspired by the idea of the Information Diffusion Mechanism [Li *et al.*, 2020; Li *et al.*, 2017; Li *et al.*, 2022], which is proposed to increase the seller' revenue in auctions via social networks. We show that the mechanism can be applied for the case where each agent can perform the task in the same quality with a probability of one. However, when the task requester is sensitive to agents' completion qualities and is uncertain about agents'

*The full version is initially published at PRIMA 2022.

performance, the Information Diffusion Mechanism cannot incentivize agents to invite each other. Though there has also been work studying task allocation problems on social networks [Jiang *et al.*, 2012; de Weerdt *et al.*, 2012], the network is known in advance.

In traditional task allocation problems, the performance and cost of each agent are private information. To incentivize each agent to report her type truthfully, the task scheduling mechanisms with verification were first proposed to take both agents' declarations and their actual performance into consideration [Nisan and Ronen, 2001; Conitzer and Vidali, 2014]. Later, to describe the execution uncertainty in more general settings, probability of success (PoS) is introduced to describe the probability of an agent successfully completing the tasks [Porter *et al.*, 2008; Ramchurn *et al.*, 2009; Zhao *et al.*, 2016].

2 The Model

Consider a social network represented by a graph $G = (V, E)$, where $V = \{s\} \cup N$ is the node set and E is the edge set. The task requester s has a single task to be performed and $N = \{1, 2, \cdots, n\}$ is the set of all other agents in the network. Each edge $\{i, j\} \in E$ indicates that agent i can directly communicate with agent j. For $i \in V$, let $r_i = \{j \in V \mid \{i, j\} \in E\}$ be the neighbour set of i. Given the task to be performed, let $Q \subset \mathbb{R}^+ \cup \{0\}$ be the set of all possible completion qualities. Let the discrete random variable Q_i be the completion quality of agent i and $q_i \in Q$ denote a realization of Q_i. Let f_i be the probability density function of Q_i, i.e., $P(Q_i = q_i) = f_i(q_i)$, which is called agent i's probability of quality (PoQ). There is also a fixed cost $c_i \geq 0$ for i to perform the task. Define $\theta_i = (f_i, c_i, r_i)$ as agent i's private type. Let Θ_i be the type space of agent i and $\theta = (\theta_1, \cdots, \theta_n)$ be the type profile of all agents.

Initially, only the task requester's neighbours r_s know the task. Hence, the task requester needs a mechanism to attract more participants, which is done by incentivizing agents to diffuse the task information to all their neighbours. Thus, each agent's action consists of reporting her PoQ, her cost to perform the task and inviting her neighbours, i.e., reporting her type. For agent $i \in N$, let $\theta'_i = (f'_i, c'_i, r'_i)$ be her report, where f'_i is a probability distribution over Q, $r'_i \subseteq r_i$ and $c'_i \geq 0$. Let $\theta' = (\theta'_1, \cdots, \theta'_n)$ be a report profile of all agents in N. Denote the graph constructed from θ' by $G(\theta') = (V, E(\theta'))$, where $E(\theta') = \{\{i, j\} \mid i \in V, j \in r'_i\}$. Let $I(\theta')$ be the set of all participants under θ', and $i \in I(\theta')$ holds if and only if there exists a path from s to i in the graph $G(\theta')$. Let Θ be the space of all possible type profiles.

Generally speaking, the mechanism consists of two steps. The task requester first announces a contract including a task allocation policy and a payoff policy and then assigns the task to an agent according to their declarations. After the task requester verifies the completion quality, she will give payoffs to agents according to the announced contract. We call such a mechanism the verified contract mechanism.

Definition 1 (Verified Contract Mechanism)**.** *A **verified contract mechanism** is defined by $\mathcal{M} = (\pi, p)$, where $\pi : \Theta \to \{0, 1\}^N$ and $p : \Theta \times Q \to \mathbb{R}^N$ are the allocation and payoff policies respectively. Given agents N and their report profile $\theta' \in \Theta$, set $\pi_i(\theta') = 0$ and $p_i(\theta', \cdot) = 0$ for all $i \notin I(\theta')$.*

Given a verified contract mechanism and a report profile θ', $\pi_i(\theta') = 1$ means that the task is allocated to agent i, otherwise she will not perform the task. The actual completion quality under the allocation π is drawn from the true PoQ of the selected agent, denoted by q_π. Then $p_i(\theta', q_\pi)$ is the payoff given to agent i. We assume that the utilities of the task requester and the agents are quasi-linear, i.e., $u_s(\theta', q_\pi) = q_\pi - \sum_{i \in N} p_i(\theta', q_\pi)$ and $u_i(\pi(\theta'), p(\theta', q_\pi)) = p_i(\theta', q_\pi) - \pi_i(\theta')c_i$ for all $i \in N$. In the following, we define several properties concerned. The first one is the efficiency of the mechanism in terms of the expected social welfare.

Definition 2 (Efficiency)**.** *A verified contract mechanism is **efficient** if for all $\theta' \in \Theta$,*

$$\pi(\theta') \in \arg\max\nolimits_{\pi' \in \Pi} \mathbb{E}_{f'_{\pi'}} \left[q_{\pi'} - \sum\nolimits_{i \in N} \pi'_i c'_i \right]$$

where Π is the space of all feasible allocations, $f'_{\pi'}$ is the PoQ reported by the agent j such that $\pi'_j = 1$.

Another property is called incentive compatibility, which requires that for each agent i, reporting her type θ_i truthfully is a dominant strategy.

Definition 3 (Incentive Compatibility)**.** *A verified contract mechanism is **incentive compatible** (IC) if for all $i \in N$ and $\theta_i \in \Theta_i$, for all $\theta'_{-i} \in \Theta_{-i}$,*

$$\theta_i \in \arg\max_{\theta'_i \in \Theta_i} \mathbb{E}_{f_i} \left[u_i(\pi((\theta'_i, \theta'_{-i})), p((\theta'_i, \theta'_{-i}), q_{\pi((\theta'_i, \theta'_{-i}))})) \right],$$

where θ'_{-i} is the report profile of all agents without i, Θ_{-i} is the space of all possible θ'_{-i}.

The next property ensures that the expected utility of an agent is non-negative when she truthfully reports.

Definition 4 (Individual Rationality)**.** *A verified contract mechanism is **individually rational** (IR) if for all $i \in N$ and $\theta_i \in \Theta_i$, for all $\theta'_{-i} \in \Theta_{-i}$,*

$$\mathbb{E}_{f_i} \left[u_i(\pi((\theta_i, \theta'_{-i})), p((\theta_i, \theta'_{-i}), q_{\pi((\theta_i, \theta'_{-i}))})) \right] \geq 0.$$

The last desirable property is that the task requester should not suffer a deficit in expectation.

Definition 5 (Weakly Budget Balanced)**.** *A verified contract mechanism is **weakly budget balanced** (WBB) if for all $\theta' \in \Theta$,*

$$\mathbb{E}_{f_{\pi(\theta')}} \left[u_s(\theta', q_{\pi(\theta')}) \right] \geq 0,$$

where $f_{\pi(\theta')}$ is the true PoQ of the selected agent under θ'.

3 The Mechanism

3.1 Without Execution Uncertainty

We first consider the setting where the task performance does not require special skills. The task requester gets the same quality no matter which agent the task is assigned to, i.e., $Q = \{q\}$ and $f_i(q) = 1$ for all $i \in N$. However, agents may need different costs to perform the given task (e.g., time). Thus, to maximize both the task requester's utility and the social welfare, a mechanism should allocate the task to an agent who can perform the task with the least cost.

We apply the VCG mechanism, the task requester selects the agent who performs with the least cost among her neighbours, and the agent's payoff equals the decrease in others' utilities due to her participation. However, the VCG mechanism will cause deficits for the task requester[1].

To tackle the deficit issue in this setting, we can apply the Information Diffusion Mechanism (IDM) [Li *et al.*, 2017]. The IDM is proposed to incentivize information diffusion to sell a single item in a social network. We just need to change the goal to select the agent with the least cost. The mechanism relies on the concept of *critical agents*.

Definition 6. *Given a report profile θ', for agent $i, j \in I(\theta')$, j is one of i's* ***critical agent*** *if j exists in all simple paths from the task requester s to agent i in the graph $G(\theta')$. The set of i's all critical agents is denoted by $ct_i(\theta')$ and the sequence order of agents in $ct_i(\theta')$ is denoted by $(s, j_1, \cdots, j_m)$, where $j_m = i$ and $k < k'$ if and only if $j_k \in ct_{j_{k'}}$.*

Information Diffusion Mechanism (IDM)

INPUT: a report profile θ'.

1. Choose $j \in \arg\max_{i \in I(\theta')}\{q - c'_i\}$ with random tie breaking.
2. Set $\pi_i = 0, p_i = 0$ for all $i \notin I(\theta')$ and all $i \notin ct_j(\theta')$.
3. Compute $w_i = \sum_{k \neq i} \pi'_k(q - c'_k)$ for each $i \in ct_j(\theta')$, where $\pi' = \left\{\pi'_i(\hat{\theta}) = \mathbb{I}\left(i \in \arg\max_{k \in I(\hat{\theta})} q - c'_k\right)\right\}_{i \in N}$ and $\hat{\theta} = \left(nil, \theta'_{-i}\right)$.
4. Let the sequence order of agents in $ct_j(\theta')$ be $(s, i_1, \cdots, i_m)$, where $i_m = j$. For $k = 1 : m - 1$,

$$\pi_{i_k} = \begin{cases} 1 & \text{if } q - c'_{i_k} = w_{i_{k+1}} \text{ and } \sum_{l=1}^{k-1} \pi_{i_l} = 0 \\ 0 & \text{otherwise} \end{cases}$$

and $\pi_j = 1$ if $\sum_{k=1}^{m-1} \pi_{i_k} = 0$.
5. Suppose i_t be the chosen agent, i.e., $\pi_{i_t} = 1, 1 \leq t \leq m$. Then, the payoff of each agent $i_k \in ct_j(\theta')$ will be

$$p_{i_k} = \begin{cases} w_{i_{k+1}} - w_{i_k} & k < t \\ q - w_{i_t} & k = t \\ 0 & k > t. \end{cases}$$

OUTPUT: the allocation π and the payoff p.

Intuitively, the IDM allocates the task to the first critical agent with the least cost when the next agent in the sequence did not participate. The payoff of the chosen agent equals the least cost reported by other agents without her participation. The payoff of each critical agent before the chosen agent is determined by the difference between the maximum social welfare without next critical agent's participation and that without her participation. The properties of IDM will stay the same.

Theorem 1 ([Li *et al.*, 2017])**.** *The IDM is IR, IC and WBB.*

Failure in the Quality-Aware Setting. Consider another setting without execution uncertainty, where the task needs special skills to be performed, and then agents may perform the task with different qualities, i.e., $f_i(q_i) = 1$ for some $q_i \in Q$ and $i \in N$. In this case, the application of IDM is not IC anymore. The failure lies in no guarantee for incentive compatibility since the payoff to the selected agent is always related to her reported ability.

3.2 With Execution Uncertainty

Look back to our general model. To mitigate IDM'a failure, we propose the **Post Execution Verification-based Diffusion Mechanism** (PEV-based Diffusion Mechanism). The PEV-based Diffusion Mechanism chooses the agent to perform the task based on agents' reports, but pays her Given a report profile $\theta' \in \Theta$, define an allocation that maximizes the expected social welfare as

$$\pi^*(\theta') = \left\{\pi^*_i(\theta') = \mathbb{I}\left(i \in \arg\max_{k \in I(\theta')}\{\mathbb{E}_{f'_k}[Q_k] - c'_k\}\right)\right\}_{i \in N}.$$

PEV-based Diffusion Mechanism (PDM)

INPUT: a report profile θ'.

1. Choose $j \in \arg\max_{i \in N}\{\mathbb{E}_{f'_i}[Q_i] - c'_i\}$ with random tie breaking. Set $\pi_i = 0, p_i = 0$ for all $i \notin ct_j(\theta')$.
2. For each agent $i \in ct_j(\theta')$, compute $w_i = \sum_{k \neq i} \pi'_k \left(\mathbb{E}_{f'_k}[Q_k] - c'_k\right)$ where $\pi' = \pi^*((nil, \theta'_{-i}))$.
3. Let the sequence order of agents in $ct_j(\theta')$ be $(s, i_1, i_2, \ldots, i_m)$, where $i_m = j$. For $k = 1 : m - 1$,

$$\pi_{i_k} = \begin{cases} 1 & \text{if } \mathbb{E}_{f'_{i_k}}[Q_{i_k}] - c'_{i_k} = w_{i_{k+1}} \\ & \text{and } \sum_{l=1}^{k-1} \pi_{i_l} = 0 \\ 0 & \text{otherwise} \end{cases}$$

and $\pi_j = 1$ if $\sum_{k=1}^{m-1} \pi_{i_k} = 0$.
4. Suppose $\pi_{i_t} = 1$ and i_t performs the quality q_π.
5. The payoff of each agent $i_k \in ct_j(\theta')$ is defined as

$$p_{i_k} = \begin{cases} w_{i_{k+1}} - w_{i_k} & k < t \\ q_\pi - w_{i_t} & k = t \\ 0 & k > t \end{cases}$$

OUTPUT: the allocation π and the payoff p.

Intuitively, the PDM allocates the task to the agent with the highest expected welfare when the next critical agent in the sequence did not participate. The payoff of the chosen agent is determined by her actual execution quality. Each critical agent before the chosen agent gets a payoff based on the increase of the expected social welfare due to her participation when the next one did not participate.

[1]An example is illustrated in the full version of this paper in the proceeding of PRIMA 2022.

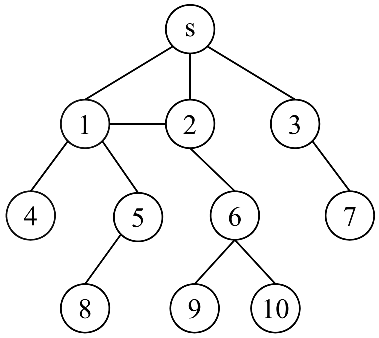

Figure 1: An example for the PEV-based Diffusion Mechanism

1	$f_1(2) = .5, f_1(3) = .5$	$c_1 = 0.5$
2	$f_2(1) = 1$	$c_2 = 0.2$
3	$f_3(5) = 1$	$c_3 = 1$
4	$f_4(3) = 1$	$c_4 = 1$
5	$f_5(4) = .4, f_5(6) = .6$	$c_5 = 1.6$
6	$f_6(3) = .3, f_6(4) = .6, f_6(7) = .1$	$c_6 = 0.9$
7	$f_7(6) = .5, f_7(8) = .5$	$c_7 = 4.2$
8	$f_8(1) = .2, f_8(3) = .8$	$c_8 = 0$
9	$f_9(8) = .8, f_9(10) = .2$	$c_9 = 1$
10	$f_{10}(4) = .5, f_{10}(5) = .3, f_{10}(6) = .2$	$c_{10} = 0.2$

Table 1: Types of the agents in Figure 1.

Example 1. *Consider the network in Figure 1, and all agents' types are listed in the Table 1. Thus, PDM chooses agent* 9 *to perform the task. Suppose that the actual execution quality of agent* 9 *is* $q_\pi = 8$*. Then, the payoffs given to agents* 2, 6 *and* 9 *are* $p_2 = w_6 - w_2 = 0$, $p_6 = w_9 - w_6 = 0.5$ *and* $p_9 = q_\pi - w_9 = 3.5$*. Finally, the utility of* s *is* $u_s = 4$.

Theorem 2. *The PDM is IR, IC, and WBB*[2].

3.3 A General Class

Li *et al.* [2022] gave a methodology to extend IDM to a generalized class of diffusion auction called Critical Diffusion Mechanisms (CDM). The key idea is choosing a different set of competitors when allocating the item. By doing this, it may give more chances to win to critical agents, and may also improve the utility of the seller.

Inspired by their idea, we can also extend PDM generally. For the allocation policy of PDM, agents who are critical parents of the best worker have priorities of winning the task. Particularly, an agent wins the task if she is the first agent among these parents whose has the best performance (i.e., the expected quality minus the cost) when her critical children are not considered. The key idea here is that a set of competitors are removed from the network when determining whether an agent can perform the task. Intuitively, the way of removing agents can be various as long as it will not affect incentive compatibility. Also, the allocation efficiency and the payoffs to the critical agents of the selected agent may vary with the removing methods. Therefore, we can have the following class of mechanisms called PEV-based Critical Diffusion Mechanism (PCDM), which is parameterized by a set selection function α_i for all $i \in ct_j(\theta')$, where j is the best worker under θ'.

[2]A complete proof can be found in the full version of this paper in the proceeding of PRIMA 2022.

PEV-based Critical Diffusion Mechanism (PCDM)

INPUT: a report profile θ'.

---

1. Choose $j \in \arg\max_{i \in N}\{\mathbb{E}_{f'_i}[Q_i] - c'_i\}$ with random tie breaking. Set $\pi_i = 0, p_i = 0$ for all $i \notin ct_j(\theta')$.
2. For each agent $i \in ct_j(\theta')$, compute
   (1) $w_i = \sum_{k \neq i} \pi'_k \left(\mathbb{E}_{f'_k}[Q_k] - c'_k\right)$, where $\pi' = \pi^*((nil, \theta'_{-i}))$;
   (2) $\tilde{w}_i = \sum_{k \notin \alpha_i} \pi''_k \left(\mathbb{E}_{f'_k}[Q_k] - c'_k\right)$, where $\pi'' = \pi^*((nil, \theta'_{-\alpha_i}))$.
3. Let the sequence order of agents in $ct_j(\theta')$ be $(s, i_1, i_2, \ldots, i_m)$, where $i_m = j$. For $k = 1 : m-1$,

$$\pi_{i_k} = \begin{cases} 1 & \text{if } \mathbb{E}_{f'_{i_k}}[Q_{i_k}] - c'_{i_k} = \tilde{w}_{i_k} \\ & \text{ and } \sum_{l=1}^{k-1} \pi_{i_l} = 0 \\ 0 & \text{otherwise} \end{cases}$$

and $\pi_j = 1$ if $\sum_{k=1}^{m-1} \pi_{i_k} = 0$.
4. Suppose $\pi_{i_t} = 1$ and i_t performs the quality q_π.
5. The payoff of each agent $i_k \in ct_j(\theta')$ is defined as

$$p_{i_k} = \begin{cases} \tilde{w}_{i_k} - w_{i_k} & k < t \\ q_\pi - w_{i_t} & k = t \\ 0 & k > t \end{cases}$$

OUTPUT: the allocation π and the payoff p.

Theorem 3. *The PCDM is IR, IC and WBB if for all* $i \in ct_j(\theta')$*, it is satisfied that* α_i

- *contains the next agent in* $ct_j(\theta')$ *after* i;
- *is independent from the reports of* i*'s critical children;*
- *and is monotonically increasing with* r'_i.

Notice that PCDMs cannot always guarantee efficiency since the task may not be allocated to the best worker. Actually, no mechanism simultaneously satisfies IR, IC, WBB and efficiency.

Proposition 1. *There is no verified contract mechanism satisfying IR, IC, WBB and efficiency in the social networks.*

Acknowledgements

This work is supported by Science and Technology Commission of Shanghai Municipality (No. 22ZR1442200).

References

[Chawla *et al.*, 2019] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. *Games and Economic Behavior*, 113:80–96, 2019.

[Conitzer and Vidali, 2014] Vincent Conitzer and Angelina Vidali. Mechanism design for scheduling with uncertain execution time. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI'14, page 623–629. AAAI Press, 2014.

[Dash *et al.*, 2007] Rajdeep K Dash, Perukrishnen Vytelingum, Alex Rogers, Esther David, and Nicholas R Jennings. Market-based task allocation mechanisms for limited-capacity suppliers. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 37(3):391–405, 2007.

[de Weerdt *et al.*, 2012] Mathijs M de Weerdt, Yingqian Zhang, and Tomas Klos. Multiagent task allocation in social networks. *Autonomous Agents and Multi-Agent Systems*, 25(1):46–86, 2012.

[Goel *et al.*, 2014] Gagan Goel, Afshin Nikzad, and Adish Singla. Allocating tasks to workers with matching constraints: truthful mechanisms for crowdsourcing markets. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 279–280, 2014.

[Jiang *et al.*, 2012] Yichuan Jiang, Yifeng Zhou, and Wanyuan Wang. Task allocation for undependable multiagent systems in social networks. *IEEE Transactions on Parallel and Distributed Systems*, 24(8):1671–1681, 2012.

[Li and Shiu, 2012] Yung-Ming Li and Ya-Lin Shiu. A diffusion mechanism for social advertising over microblogs. *Decision Support Systems*, 54(1):9–22, 2012.

[Li *et al.*, 2017] Bin Li, Dong Hao, Dengji Zhao, and Tao Zhou. Mechanism design in social networks. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 586–592. AAAI Press, 2017.

[Li *et al.*, 2020] Bin Li, Dong Hao, and Dengji Zhao. Incentive-compatible diffusion auctions. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 231–237. ijcai.org, 2020.

[Li *et al.*, 2022] Bin Li, Dong Hao, Hui Gao, and Dengji Zhao. Diffusion auction design. *Artificial Intelligence*, 303:103631, 2022.

[Nisan and Ronen, 2001] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic behavior*, 35(1-2):166–196, 2001.

[Porter *et al.*, 2008] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.

[Ramchurn *et al.*, 2009] Sarvapali D Ramchurn, Claudio Mezzetti, Andrea Giovannucci, Juan A Rodriguez-Aguilar, Rajdeep K Dash, and Nicholas R Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research*, 35:119–159, 2009.

[Shi *et al.*, 2020] Haomin Shi, Yao Zhang, Zilin Si, Letong Wang, and Dengji Zhao. Maximal information propagation with budgets. In *ECAI 2020 - 24th European Conference on Artificial Intelligence*, volume 325, pages 211–218. IOS Press, 2020.

[Tang *et al.*, 2011] John C Tang, Manuel Cebrian, Nicklaus A Giacobe, Hyun-Woo Kim, Taemie Kim, and Douglas" Beaker" Wickert. Reflecting on the darpa red balloon challenge. *Communications of the ACM*, 54(4):78–85, 2011.

[Wu *et al.*, 2017] Xiaocan Wu, Danlei Huang, Yu-E Sun, Xiaofei Bu, Yu Xin, and He Huang. An efficient allocation mechanism for crowdsourcing tasks with minimum execution time. In *International Conference on Intelligent Computing*, pages 156–167. Springer, 2017.

[Zhao *et al.*, 2016] Dengji Zhao, Sarvapali D. Ramchurn, and Nicholas R. Jennings. Fault tolerant mechanism design for general task allocation. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 323–331. ACM, 2016.

[Zhao, 2021] Dengji Zhao. Mechanism design powered by social interactions. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 63–67, 2021.

[Zhao, 2022] Dengji Zhao. Mechanism design powered by social interactions: a call to arms. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 5831–5835, 2022.