COOPERATION MECHANISM FOR A NETWORK GAME

Alon Grubshtein and Amnon Meisels
Dept. of Computer Science, Ben Gurion University of the Negev, Beer-Sheva, Israel

Keywords: Multi agent systems, Self interest, Cooperation, Cost of cooperation.

Abstract: Many real world Multi Agent Systems encompass a large population of self interested agents which are connected with one another in an intricate network. If one is willing to accept the common axioms of Game Theory one can assume that the population will arrange itself into an equilibrium state. The present position paper proposes to use a mediating cooperative distributed algorithm instead. A setting where agents have to choose one action out of two - download information or free-ride their neighbors' effort - has been studied recently. The present position paper proposes a method for constructing a Distributed Constraint Optimization Problem (DCOP) for a Network Game. The main result is that one can show that by cooperatively minimizing the constructed DCOP for a global solution all agents stand to gain at least as much as their equilibrium gain, and often more. This provides a mechanism for cooperation in a Network Game that is beneficial for all participating agents.

SCIENCE AND TECHNOLOGY PUBLICATIONS

1 INTRODUCTION

A key attribute of any Multi Agent System (MAS) is the level of collaboration that agents are expected to adopt. When agents share a common goal it is natural for agents to follow a fully cooperative protocol. A common goal can be the election of a leader, finding shortest routing paths, or searching for a globally optimal solution to a combinatorial problem (Meisels, 2007). When the involved parties (agents) have different and conflicting goals, competition and self interest are the natural behaviors one expects to find.

In its most basic form, a fully cooperative model involves a set of agents attempting to satisfy or optimize a common *global* objective. The most important aspect of such a model is that the actions taken by the agents do not bring into consideration the impact of the globally optimal solution on the state of the individual agent (Maheswaran et al., 2004; Grubshtein et al., 2010; Meisels, 2007).

In a fully competitive scenario it is common to assume that participants are only willing to take actions which improve (or do not worsen) their gains. In some situations agents can reach an equilibrium state from which no agents would care to deviate (Roughgarden, 2005; Meir et al., 2010). The efficiency of equilibria states with respect to the global objective has been studied intensively in the last decade (Roughgarden, 2005).

The central question that is at the focus of the present study is the interplay between the self-interests of agents and their cooperation towards a common goal. Specifically, towards the increase of some global gain. A rich and applicable framework for this study is the domain of Network Games (NGs) (Jackson, 2008; Galeotti et al., 2010). File sharing and ad hoc P2P networks are good examples of network games. The existence of a Bayesian Nash Equilibrium (BNE) for the NG introduced in Section 2 was shown by (Galeotti et al., 2010). The present paper proposes a method for securing cooperation among agents in a network game, by guaranteeing improved gains to each agent. This approach is related to former work on Cooperation Games (Grubshtein and Meisels, 2010). The proposed method is based on the construction of an Asymmetric Distributed Constraints Optimization Problem (ADCOP) (Grubshtein et al., 2010) from the given network game and cooperatively solving it.

Former attempts to introduce cooperation into competitive games fall into several categories. Strong empirical evidence to the benefits and emergence of cooperation in the iterated prisoners' dilemma tournament was reported in (Axelrod, 1984). Petcu et. al present a cooperative distributed search algorithm which faithfully implements the VCG mechanism for the problem of efficient social choice by self-interested agents in DCOP search (Petcu et al., 2008). An approach for achieving cooperation among competitive agents is described in (Monderer and Tennen-

Grubshtein A. and Meisels A..
COOPERATION MECHANISM FOR A NETWORK GAME.
DOI: 10.5220/0003276303360341
In *Proceedings of the 3rd International Conference on Agents and Artificial Intelligence* (ICAART-2011), pages 336-341
ISBN: 978-989-8425-41-6
Copyright © 2011 SCITEPRESS (Science and Technology Publications, Lda.)

holtz, 2009), which introduce mediators to achieve stronger equilibria among competitive agents.

The present study is unique in proposing a distributed search method that guarantees the improvement of the personal gain of *each* agent. It uses the definition of the *Cost of Cooperation* (Grubshtein and Meisels, 2010), which defines a measure of individual loss from cooperation. It is shown that for a known family of network games this cost is negative (see section 3). That is, a complex real world setting is presented where one can prove that cooperation produces better individual gains to ***all*** participating agents.

2 DOWNLOAD/FREE-RIDE NETWORK GAME

Consider a set of users wishing to download large amounts of information from a remote location. Users may receive the information from other users in their local neighborhood (if these neighbors have it) or directly download it from a central hub. The process of downloading data from the hub requires significant bandwidth and degrades performance (resource consumption, battery, etc). In contrast, sharing information with peers does not result in any significant impact on the users. Information exchange between users is limited to directly connected peers only (no transfer of information to second degree neighbors).

The entire interaction can be captured by a graph $G=(V,E)$ where each edge e_{ij} specifies a neighborhood relationship and a vertex $i\in V$ represents a user. x_i denotes the action (e.g., assignment) made by user i, and $x_N(i)$ is the joint action taken by all of i's neighbors. [1] The degree of i will be denoted by d_i.

While the network itself (the graph $G=(V,E)$) is not known to all agents, the total number of participants $n=|V|$ is known and so is the (fixed) probability p for an edge to exist between any two vertices. Such graphs $G=(n,p)$ are known as Poisson Random Graphs or an Erdös - Rényi network (Jackson, 2008). Figure 1 is an example of such a network with $n=8$ and $p=0.45$ (8 participants and 13 edges).

For simplicity, the gain from receiving information is unity and the cost of downloading it is c. Users are only aware of their immediate peers and are known to be self interested. That is, users decide to either download the relevant information from the hub (take action $\mathcal{D}$) or wait for one of their neighbors to download it (take action $\mathcal{F}$), and do so in a way which

[1]We specifically refrain from the common x_{-i} notation to represent a joint action by all players but i to emphasize that player i is only affected by the set of her neighbors.

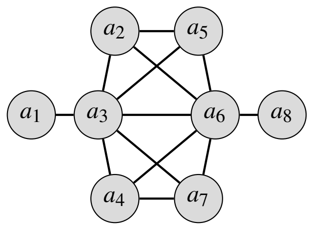

Figure 1: An interaction graph with $n=8$ and $p=0.45$.

maximizes their own utility function:

$$u_i(x_i,x_N(i))=\begin{cases}1-c & \text{if } x_i=\mathcal{F}\\ 1 & \exists j\in x_N(i) \text{ s.t } x_j=\mathcal{D}\\ 0 & \text{otherwise}\end{cases}$$

That is, if the user exerts effort and downloads the information her gain is $1-c$. If, on the other hand, she does not download but one of her peers does, her gain is 1. Finally, if neither her nor any of her peers download the information her gain is 0. We assume that the interaction between users is a one shot interaction (i.e. users are assumed to be at some decision crossroad).

2.1 The Game Theoretic Approach

As mentioned in section 1, the common approach taken when self interested entities are present, is to invoke *Game Theory* (Osborne and Rubinstein, 1994). By accepting the fundamental axioms of game theory one is able to predict possible outcomes of an interaction. These outcomes are standardly assumed to be equilibrium points (in terms of agents' actions) from which no agent is willing to deviate. The best known solution concept is the Nash Equilibrium (NE) point, but there are many other equilibrium points which fit different aspects of the interaction (cf. (Osborne and Rubinstein, 1994)). For example, in the above D/F network game the equilibrium concept used is the *Bayesian Nash Equilibrium* (BNE). The BNE captures the incomplete information of all participants (each knows only its own neighbors).

This is the approach taken in a recent work by Galeotti et. al (Galeotti et al., 2010) for a similar interaction. They describe our example of a network game and cast it to many other types of interactions: vaccinate or not, research new technology or wait for the competitors to do it, etc. In the domain of computer science one can think of remote rovers belonging to different agencies exploring the same area, of ad-hoc network participation, and of P2P networks.

Let us describe the network game that is at the focus of the present study in more detail. ¿From the point of view of each participant in the game (i.e. a user) there are two possible strategies: $\mathcal{D}$ (download),

$\mathcal{F}$ (free ride). Representing this game by a traditional n-dimensional matrix is clearly intractable and hence a *graphical game* representation is used. A graphical game (Kearns et al., 2001) is a succinct representation in which each vertex u represents a party and the set of edges emanating from it represents its relations to a subset of the vertices. These vertices are the only vertices which affect u. When the underlying graph of a game is the complete graph this representation is equivalent to the n-dimensional matrix representation.

Although the payoff structure is known to all participants, the degrees of neighbors (e.g., their connectedness) are unknown. Based on the global probability p for an edge, each participant can calculate the probability for a randomly selected neighbor to be of degree k:

$$Q(k;p) = \binom{n-2}{k-1} p^{(k-1)} (1-p)^{(n-k-1)}$$

Galeotti et. al use this information to find a threshold Bayesian Nash Equilibrium. They define a parameter t which is the smallest integer such that:

$$1 - \left[1 - \sum_{k=1}^{t} Q(k;p)\right]^t \geq 1 - c$$

In the unique BNE that they find for the network game, any participant of degree $k \leq t$ selects strategy $\mathcal{D}$ and any participant of degree $k > t$ selects $\mathcal{F}$.

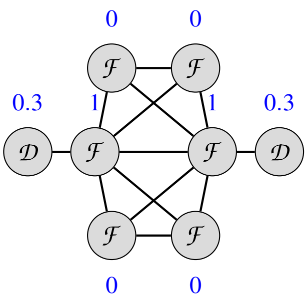

Figure 2: The strategies of all participants in the BNE and the corresponding gains.

For example, consider the graph $G = (8, 0.45)$ illustrated in Figure 1 and a cost value of $c = 0.7$. From the above equation we calculate the threshold $t = 2$. As a result, in the BNE, a_1 and a_8 will select strategy $\mathcal{D}$ while all others will assign $\mathcal{F}$. The BNE assignments and gains are depicted in Figure 2.

3 COOPERATIVE MECHANISM FOR SELF INTERESTED AGENTS

Despite the success game theory has had in many MAS problems and applications it suffers from two important drawbacks:

1. It often fails to predict human behavior (e.g., the iterated prisoners' dilemma (Axelrod, 1984)).
2. Results are often inefficient (i.e. equilibria points are not necessarily pareto efficient) (Nisan et al., 2007).

The former point induces interesting research (often related to psychology) which may produce new solution concepts (Halpern and Rong, 2010). The present paper focuses on the latter point. More specifically, we follow the ideas presented in (Grubshtein and Meisels, 2010) which define the Cost of Cooperation and Cooperation Games:

Definition 1. *An agent's Cost of Cooperation (CoC) with respect to a global objective function $f(x)$ is defined as the difference between the lowest gain that the agent can get in any equilibria of the underlying game (if any, otherwise zero) and the lowest gain an agent receives from a (cooperative) protocol's solution x that maximizes $f(x)$.*

Definition 2. *A game is a* Cooperation Game *(CG) if there exists a solution for which the CoC (with respect to some $f(x)$) of all agents is non positive.*

In other words, It is beneficial for all agents to cooperate in search for an optimal solution, rather then play their NE strategy (as competitive agents are expected to act) (Grubshtein and Meisels, 2010).

For a game that satisfies the above two definitions one can propose a mechanism which will provide the following guarantee: ***if all participants agree to follow the mechanism, they are each expected to gain at least as much as they would have gained had they played their NE strategy***.

The present study focuses on Network Games and on the D/F game in particular. The main result of the paper is a proof that the D/F game is a Cooperation Game (CG). This enables a clear mechanism that solves a generated Asymmetric DCOP (ADCOP) (Grubshtein et al., 2010) and can act as a natural choice of strategy for all participants of the game. An ADCOP is an extension to the DCOP paradigm which addresses individual agent's gains and captures game-like interactions between agents. An ADCOP can be viewed as a form of a cooperative graphical game.

The proposed framework for the mechanism makes the distinction between *users* and *agents*.

SCITEPRESS SCIENCE AND TECHNOLOGY PUBLICATIONS

When referring to *users* one refers to self interested entities seeking to maximize their gain. Each of the *users* employs an *agent* - an application entity which follows a cooperative protocol on the one hand, but is expected to faithfully represent its *user's* preferences. In the proposed setting agents represent the users, but are not under the users' direct control (i.e. users' entrust the agents with the ability to choose for them, but are unable to control or affect the agents' decisions).

An important feature of the proposed approach is its inherent cooperation. The desired solution it seeks is not necessarily a stable one in the game theoretic sense. Nonetheless, rational, self interested users would rather let their agents choose an action for them knowing in advance that agents' cooperation will yield results which are at least as good as they would get if they play selfishly.

The distributed cooperative protocol for the agents which will provide the above guarantee to the users solves the following ADCOP:

- A set of Agents $A = \{a_1, a_2, \ldots, a_n\}$ - each holds a single variable and corresponds to a user.
- A set of Domains $D = \{D_1, D_2, \ldots, D_n\}$ for the variables held by the agents. Each domain consists of only two possible values: $\{\mathcal{D}, \mathcal{F}\}$.
- A set of asymmetric constraints C. Each constraint c_i is identified with a specific agent (e.g. agent i) and defined as a (d_i+1)-arity constraint acting over all of agent's i neighbors. The set of agents involved in constraint c_i is $A_{c_i} = a_i \cup x_N(i)$. That is, the ADCOP has n constraints and each agent contributes to $d_i + 1$ different constraints. The costs associated with each constraint that add up to the total cost of c_i for agent i are summarized in the table in Figure 3.

Costs of c_i	
n^3	if $\forall a_j \in A_{c_i}, x_j = \mathcal{F}$ (all agents of the constraint assign $\mathcal{F}$)
d_j	**for each** $a_j \in A_{c_i}$ with $d_j > t$ and $\forall a_k \in x_N(j), d_k > t$ that assign $x_j = \mathcal{D}$
$(n-1)n$	**for each** $a_j \in A_{c_i}$ with $d_j > t$ and $\exists a_k \in x_N(j), d_k \leq t$ that assign $x_j = \mathcal{D}$
1	**for each** $a_j \in A_{c_i}$ with $d_j \leq t$ which assigns $x_j = \mathcal{D}$

Figure 3: The constraint's costs for each combination of assignments.

Note: costs of the DCOP's solutions (full or partial) are completely unrelated to the agents' gains.

Let us now proceed to prove that finding a minimal assignment (an optimal solution) for the ADCOP results in gains which are at least as high as the actual gains resulting from the BNE, for all agents .

Let the optimal solution to the ADCOP be denoted with x. We use the term *null neighborhood* to define a situation in which an agent and all of its directly connected peers assign $\mathcal{F}$ as their action.

Lemma 1 (No null neighborhoods)**.** *In the optimal solution to the ADCOP there exists at least one agent in the local neighborhood of each agent which assigns $\mathcal{D}$ (specifically, the agent itself can assign $\mathcal{D}$).*

Proof. Assume by negation that there exists an agent a_i with a null neighborhood in x. This means that the cost of c_i (the constraint originating from a_i) is n^3. Let x' be a complete assignment which differs from x in the assignment of an agent $a_j \in c_i$. That is, $x'_j = \mathcal{D}$. As a result the previously null neighborhood of a_i in x is not part of x'.

One can compute an upper bound on the cost of x' by assigning this change to an agent a_j with $d_j = n-1$ (where $n-1 > t$) which has at least one neighbor $a_k \in x_N(j)$ of degree $d_k \leq t$ (i.e. a_j appears in all n constraints of the problem). The resulting cost of x':

$$\begin{aligned} COST(x') &= COST(x) - n^3 + (n-1)n^2 \\ &= COST(x) - n^2 \end{aligned}$$

which implies that $COST(x') < COST(x)$ in contradiction to the optimality of x. □

An important implication of Lemma 1 is that in the optimal solution to the ADCOP none of the agents receive a payoff of 0.

Lemma 2. *The gain of agents a_i with degree $d_i \leq t$ in x is at least as high as its gain in the BNE.*

Proof. In the BNE a_i's gain is exactly 1-c ($d_i \leq t$ and hence a_i assigns $\mathcal{D}$). The only possible lower gain is 0. However, this is the gain of an agent with a null neighborhood and hence (following Lemma 1) its gain must be at least 1-c when cooperating (either it will assign $\mathcal{D}$ or one of its peers will). □

Lemma 3. *The gain of agent a_i with degree $d_i > t$ in x is at least as high as its gain in the BNE.*

Proof. Due to its degree, a_i's BNE gain is never $1-c$. When all its neighbors $a_j \in x_N(i)$ are of degree $d_j > t$ its BNE gain is 0, and when at least one neighbor is of degree $d_j \leq t$ its BNE gain is 1. Since the ADCOP gain is always higher than 0 (Lemma 1) we only consider the case when a_i's gain is 1 in the BNE. This is only attainable when a_i has at least one neighbor a_j with degree $d_j \leq t$. Thus, to conclude our proof we have to show that a_i's ADCOP gain is not $1-c$.

By the problem description this gain can only occur when a_i assigns $\mathcal{D}$ in x. This assignment incurs a

SCITEPRESS SCIENCE AND TECHNOLOGY PUBLICATIONS

positive contribution on the ADCOP only when it prevents the existence of a null neighborhood for a_i, or when it prevents it for one (or more) of a_i's neighbors.

If a_i has a null neighborhood which it would like to prevent, there must be a neighbor $a_j \in x_N(i)$ (with degree $d_j \leq t$) which assigns $\mathcal{F}$ in the ADCOP. In this case we define a new solution x', which differs from x in the assignments of a_i and a_j (i.e., $x'_i = \mathcal{F}$ and $x'_j = \mathcal{D}$). x''s cost is:

$$COST(x') = COST(x) - (d_i+1)(n-1)n + d_j$$

The highest cost x' can take will be when $d_i = 1$ and $d_j = n-1$ in which case:

$$\begin{aligned} COST(x') &= COST(x) - 2(n-1)n + (n-1) \\ &= COST(x) - (n-1)(2n-1) \end{aligned}$$

implying that a_i does not assign $\mathcal{D}$ to prevent a local null neighborhood in the optimal solution.

When a_i assigns $\mathcal{D}$ to prevent the existence of a null neighborhood for one of its neighbors, we classify the neighbor's type:

1. The neighbor a_j has at least one neighbor a_k with degree $d_k \leq t$. We generate the assignment x' in which $x_i = \mathcal{F}$ and $x_k = \mathcal{D}$. As before the cost of x' is lower than that of x since a_k contributes a single unit per each constraint it is involved in (at most n units) whereas the assignment change of a_i lowers the cost by at least $(n-1)n$ units in contradiction to the optimality of x.
2. The neighbor a_j has no neighbor a_k with degree $d_k \leq t$. Assigning $x_j = \mathcal{D}$ and $x_i = \mathcal{F}$ to a_i lowers the resulting cost since the contribution of a_j to the cost is bounded by n^2 (it incurs a cost of d_j to each of the $d_j + 1$ constraints it is involved in), whereas the gain from changing the assignment of a_i is $(d_i+1)(n-1)n$.

This means that even if all neighbors are of the second type (which always incur greater costs), and they are all connected to all other agents, the cost of the modified solution is:

$$COST(x') = COST(x) - (d_i+1)(n-1)n + d_i \cdot n(n-1)$$

Hence we conclude that that the gain of a_i will not be reduced from 1 in the BNE to $1-c$ (or 0) in the ADCOP's optimal solution. □

Theorem 1 (Gain guarantee)**.** *The optimal solution to the ADCOP described above results in gains which are at least as high as those achieved in the BNE.*

Proof. Directly follows from Lemmas 1-3. □

Figure 4 demonstrates the results of applying and solving the ADCOP on our previous example. Agents a_3, a_4, a_5 and a_6 assign $\mathcal{F}$ and do not incur any cost. Agent a_1 and a_8 incur a cost of 1 to c_1 and c_3 (to c_6 and c_8 in case of a_8) and agents a_2 and a_7 incur a cost of 3 (their degree) to the 4 constraints they are involved with (c_2, c_3, c_5, c_6 and c_3, c_4, c_6, c_7) resulting in a cost of 28.

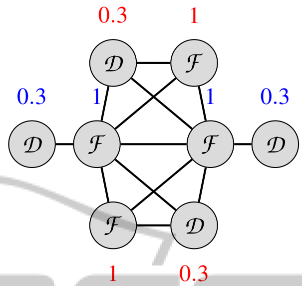

Figure 4: Strategies of all participants in the ADCOP and the corresponding gains. Red indicates higher gains than those received in the BNE.

SCITEPRESS SCIENCE AND TECHNOLOGY PUBLICATIONS

Note that the ADCOP's solution dictates an assignment in which 4 agents (namely a_2, a_5, a_6 and a_7) increase their gain. None of the agents gained less than their BNE gain depicted in Figure 2.

4 DISCUSSION

In a general combinatorial problem, maximizing the total sum of gains (also known as a "utilitarian" scheme, or the social choice) provides no guarantees to the individual participant. Depending on the nature of the problem, achieving the optimal social solution requires that some participants agree to an extremely low gain (Moulin, 1991). On the other hand, the equilibrium solution suggested by Game Theory for problems involving only self interested parties may be greatly improved when agents cooperate (e.g., the "prisoners' dilemma" (Axelrod, 1984; Monderer and Tennenholtz, 2009)).

The present paper proposes a cooperative, quality guaranteeing framework to overcome this problem. The proposed framework and mechanism are proven to provide strong guarantees for a general network game (the Download/Free-ride game) which can be cast into many diverse applications.

In the proposed framework, agents acting on behalf of users adopt a cooperative behavior (cooperatively solve an ADCOP) to achieve a solution which pareto dominates the Bayesian Nash Equilibrium. Thus, users are always expected to gain at least as much as in the BNE. One of the benefits of this approach is that it justifies the use of cooperative algorithms as mechanisms for instances where the Nash Equilibrium is not readily known.

The state of a Network Game in which agents do not fully know the whole network is natural to a distributed scenario in which information is often local. The analogous situation in ADCOPs is one in which agents only communicate with their neighbors during search and are not aware of more remote agents. The corresponding family of distributed search algorithms (local search algorithms) are not guaranteed to find optimal results but may be scaled to large populations of agents (Grubshtein et al., 2010; Maheswaran et al., 2004; Zhang et al., 2005).

It is not clear that our previous guarantee can be satisfied in a setting where agents employ a local search algorithm: Consider the interaction of Figure 1 and a set of agents participating in a stochastic search (Zhang et al., 2005). The initial assignment of all agents is $\mathcal{F}$. As a result, all agents consider a change of assignment to $\mathcal{D}$ in the next round. Specifically, Agents A_3 and A_8 may change their assignments while the rest of the agents stochastically avoid any change. The resulting solution is a local minima from which *agents* will not deviate (similar to an equilibria from which *users* will not deviate). In this converged solution our guarantee is violated - agent A_3's gain is reduced from 1 to $1-c$. Nonetheless, a simple manipulation to distributed hill climbing algorithms such as MGM (Maheswaran et al., 2004) can result in solutions which provide the desired guarantee.

The framework and methods proposed in the present position paper form a mechanism that enables the self driven desires and goals of users to be solved by a cooperative system of computerized agents. We believe that this is a natural mechanism for many user applications which interact with their environment and with other users. In such settings, cooperation between agents is the natural action only if it can provide a strong and realistic guarantee regarding the expected gain to each user. A guaranteed negative CoC provides a suitable incentive for cooperation - securing a gain which is at least as high as the worst possible gain attained by the user.

ACKNOWLEDGEMENTS

This work was supported by the Lynn and William Frankel center for computer science and the Paul Ivanier Center for Robotics.

REFERENCES

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books, New York.

Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., and Yariv, L. (2010). Network games. *Review of Economic Studies*, 77(1):218–244.

Grubshtein, A. and Meisels, A. (2010). Cost of cooperation for scheduling meetings. *Computer Science and Information Systems (ComSIS)*, 7(3):551–567.

Grubshtein, A., Zivan, R., Meisels, A., and Grinshpoun, T. (2010). Local search for distributed asymmetric optimization. In *Proc. of the 9th Intern. Conf. on Autonomous Agents & Multi-Agent Systems (AAMAS'10)*, pages 1015–1022, Toronto, Canada.

Halpern, J. Y. and Rong, N. (2010). Cooperative equilibrium (extended abstract). In *Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems(AAMAS 2010)*, Toronto, Canada.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.

Kearns, M. J., Littman, M. L., and Singh, S. P. (2001). Graphical models for game theory. In *UAI '01: Proc. of the 17th Conf. in Uncertainty in Artificial Intelligence*, pages 253–260.

Maheswaran, R. T., Pearce, J. P., and Tambe, M. (2004). Distributed algorithms for DCOP: A graphical-game-based approach. In *Proc. Parallel and Distributed Computing Systems PDCS)*, pages 432–439.

Meir, R., Polukarov, M., Rosenschein, J. S., and Jennings, N. (2010). Convergence to equilibria of plurality voting. In *The Twenty-Fourth National Conference on Artificial Intelligence (AAAI'10)*.

Meisels, A. (2007). *Distributed Search by Constrained Agents: Algorithms, Performance, Communication*. Springer Verlag.

Monderer, D. and Tennenholtz, M. (2009). Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195.

Moulin, H. (1991). *Axioms of Cooperative Decision Making*. Cambridge University Press.

Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press.

Osborne, M. and Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press.

Petcu, A., Faltings, B., and Parkes, D. (2008). M-DPOP: Faithful distributed implementation of efficient social choice problems. *Journal of AI Research (JAIR)*, 32:705–755.

Roughgarden, T. (2005). *Selfish Routing and the Price of Anarchy*. The MIT Press.

Zhang, W., Xing, Z., Wang, G., and Wittenburg, L. (2005). Distributed stochastic search and distributed breakout: properties, comparishon and applications to constraints optimization problems in sensor networks. *Artificial Intelligence*, 161:1-2:55–88.