Power, Dependence and Stability in Multiagent Plans

Sviatoslav Brainov and Tuomas Sandholm

Washington University
Department of Computer Science
St. Louis, MO 63130
{brainov, sandholm}@cs.wustl.edu

Abstract

In this paper we present a decision-theoretic model of social power and social dependence that accounts for origins of different choices available in different situations. According to the model almost every group activity, whether it is cooperation or exploitation, has its origins in resolving some dependence or power relation. The model is intended for self-interested agents and explains power and dependence in terms of relations between agents' plans. It is a generalization of the dependence network model and accounts for situations of group dependence, i.e., situations in which an agent depends on a group or a group depends on an agent. The model is applied to the analysis of stability of multiagent plans. Stable dependence structures of multiagent plans are identified. Necessary and sufficient conditions for stability of joint plans are provided.

Introduction

Social relations between autonomous agents have been a subject of continuous interest in both multiagent systems and distributed AI (Castelfranchi 1990; Sichman et al. 1994; d'Iverno and Luck 1997). The set of all social relations between agents defines the social structure of society. Reasoning about the social structure enables agents to make consistent decisions about when and with whom to interact.

Since interaction is a distinctive feature of multiagent activity, agents need to know how to avoid accidental and harmful interactions and how to take advantage of beneficial interactions. Therefore, agents need a model of how to influence other agents' behavior. The ability of an agent to exert influence over another agent is captured by the notion of power. Almost every multiagent interaction has its origins in resolving some power relation. Power relations become relevant in both conflict and cooperative situations. In a situation where two or more agents have conflicting interests a decision has to be made as to whose interest shall prevail. In a cooperative situation, in order to achieve his goal, an agent relies on the power of another agent or a group of agents. A special kind of power situation is a dependence situation. In a dependence situation an agent submits himself to the power of another agent.

Social power and social dependence play a central role in game theory (Harsanyi 1962) and sociology (Coleman 1973). The notion of social power was introduced in AI in (Castelfranchi 1990). Later, this notion was employed in the definition of the social dependence network model and the implementation of different social reasoning mechanisms (Sichman et al. 1994). It was proved that the problem of determining whether cooperation is possible in a given social structure is NP-complete (d'Iverno, Luck and Wooldridge 1997).

Prior research on social dependence does not, however, exhibit credible causal mechanisms for dependence resolution or submission to someone's power. Usually social dependence is explained in terms of the theory of joint intentions and joint actions (Cohen and Levesque 1990; Levesque, Cohen and Nunes 1990). The theory of Cohen and Levesque is not intended for self-interested agents and, therefore, does not provide sufficient grounds for explaining self-motivated behavior. The notions of joint persistent goal and joint persistent action usually presuppose altruistic behavior (Brainov 1996). Achieving persistent and stable behavior of self-interested agents is a difficult task. Considerable research in multiagent systems has already focused on stability of multiagent agreements (Brainov 1994; Sandholm and Lesser 1995).

Another limitation of social dependence theory is that it accounts only for bilateral dependence situations, i.e., for situations involving two agents. In many situations of practical interest an agent might depend on another agent as well as on a group of agents.

In this paper we propose a decision-theoretic model for explaining social power and social dependence in multiagent plans. The model is intended for self-interested agents and explains power relations in terms of individual interest. It is a generalization of the dependence network model and accounts for situations of group dependence, i.e., situations in which an agent depends on a group or a group depends on an agent.

The model is applied to the analysis of stability of multiagent plans. It is shown that in order to be stable a joint plan has to balance the power of different agents.

This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994.

The paper is organized as follows. First, the notion of dependence is introduced and generalized to the case of more than two agents. Next, the resolution of dependence relations is discussed. Finally, the notion of dependence is applied in the analysis of stability of multiagent plans. The paper concludes by summarizing the results and providing directions for future research.

Unilateral Dependence

Let the set of all agents be denoted by N, N={1,2,…,n}. The n-tuple $(p_1,p_2,\ldots,p_n)$ is used to denote a joint plan of all agents for moving the world from the initial state s^1 to some final state s^r. Here p_i stands for the individual subplan of agent i. Each agent's subplan can be viewed as a sequence of actions: $p_i=(a_i^1,\ldots,a_i^{r-1})$ such that every n-tuple $[a_1^k,\ldots,a_n^k]$, possibly with some precedence constraints between pairs (a_i^k,a_j^k), moves the world from a state s^k to the state s^{k+1}. As usual, agents are utility maximizers and each agent i, $i\in N$, attempts to maximize his utility U_i in the final state. Let $p_{N\text{-}i}$ denote the plan of the group of agents N-{i}. In this notation we can write the plan $(p_1,p_2,\ldots,p_n)$ as $(p_{N\text{-}i}, p_i)$.

The following definition introduces the notion of individual dependence, i.e., an agent's dependence on another agent in some joint plan.

Definition 1. In a joint plan $(p_{N\text{-}j},p_j)$ agent i *depends* on agent j if there exists plan p^*_j such that for every plan $p'_{N\text{-}j}$ of the other agents:

$$U_i(p'_{N\text{-}j},p^*_j) < U_i(p_{N\text{-}j},p_j), \quad (1)$$
$$U_j(p_{N\text{-}j},p^*_j) \geq U_j(p_{N\text{-}j},p_j). \quad (2)$$

The plan p^*_j is said to be gainful for agent j at the expense of agent i. To indicate that the gainful plan depends on the initial plan $(p_{N\text{-}j},p_j)$, we denote it by $p_j^{gain}(i,p_{N\text{-}j},p_j)$.

Condition (1) says that by implementing his plan p^*_j, agent j lowers the utility of agent i. According to Condition (2) the plan p^*_j is at least as beneficial to agent j as the original plan p_j, provided that the other agents do not change their plans. The situation described by Definition 1 can be classified as one where agent j has social power (Harsanyi 1962). The plan p^*_j serves as the *means of power*, i.e., those specific actions by which agent j can influence agent i's behavior. According to Condition (2) the *cost of power* is zero or negative. That is, agent j does not incur any additional costs in order to exert any influence over agent i. Therefore, with each unilateral power situation we can associate two features: the means of power and the cost of power. The means of power is the particular plan by which the powerful agent can exert his power over the dependent agent. The cost of power is the cost of applying the means of power.

Condition (1) of Definition 1 holds for every plan $p'_{N\text{-}j}$ of the other agents. That is, no matter how other agents will help agent i, they cannot prevent agent j from harming agent i. Therefore, the power of agent j with respect to agent i is unconditional and does not depend on other agents. The notion of dependence introduced in Definition 1 generalizes traditional concepts of resource and action dependence (Sichman et al. 1994). Consider the following example. Agent i needs a resource r_j for achieving his goal, i.e., without r_j agent i is either incapable of achieving his goal or the achievement is too costly. The only agent who controls resource r_j is agent j. Suppose that agent j has committed himself to give r_j to agent i at some stage of the joint plan $(p_{N\text{-}j},p_j)$. If the commitment is not binding, agent j can deviate from the plan $(p_{N\text{-}j},p_j)$ by withholding r_j to himself. By so doing, agent j gains at the expense of agent i. Therefore, agent i depends on agent j. Thus, the notion of resource dependence is a special case of dependence. It is worth noting that the other agents in the environment cannot oppose agent j. Eventually, they could punish agent j for his behavior, but they cannot provide agent i with a resource equivalent or better than r_j.

The traditional notion of resource (action) dependence does not account for the case when a resource (action) can be provided by several agents. Suppose that in the previous example resource r_j can be provided by agent j or agent k. According to the traditional notion of resource dependence agent i depends on agent j for r_j. If agent j deviates from the plan $(p_{N\text{-}j},p_j)$ by withholding r_j to himself, agent i can still receive the resource from agent k. Therefore, in accordance with Definition 1, in this case agent i does not depend on agent j. Later we will classify this situation as conditional dependence of agent i on agent j with the tacit consent of agent k. That is, agent i might depend on agent j if agent k is backing up agent j.

Group Dependence

In this section the notion of dependence is generalized to the case where there are more than two agents. The existence of other agents in an environment gives rise to different types of group dependence. The next definition introduces the transitive closure of a dependence relation.

Definition 2. In a joint plan $p=(p_1,\ldots,p_n)$ agent i, $i\in N$, *depends transitively* on agent j, $j\in N$, if

(i) agent i depends on agent j, or
(ii) there exists agent k, $k\in N$, such that agent i depends on agent k and agent k depends transitively on agent j.

With the help of transitive dependence we are now in a position to define reciprocal dependence, i.e., a situation in which everybody depends on everybody directly or indirectly.

Definition 3. A joint plan $p=(p_1,\ldots,p_n)$ is based on *reciprocal dependence* if for every two agents i and j, $i,j\in N$, $j\neq j$, agent i depends transitively on agent j in the plan p.

Dependence structure of a joint plan can be represented by a dependence graph (Brainov 1994; Sichman et al. 1994). The dependence graph is constructed in the following way. The agents are represented by nodes and if I and J are two nodes corresponding to agents i and j respectively, these two nodes are joined by an arc pointing toward J if and only if agent i depends on agent j. It is evident that a joint plan is based on reciprocal dependence if and only if the dependence graph associated with it is strongly connected.

The following definition introduces the notion of group dependence. In this paper we consider two types of group dependence. The first type refers to the case where a group of agents depends on a single agent. The second type occurs when an agent depends on a group

Definition 4. In a joint plan (p_{N-j},p_j) a group of agents S, $S \subseteq N-j$, *depends* on agent j if there exists a plan p^*_j such that for every plan p'_{N-j} of the other agents:

$$U_k(p'_{N-j},p^*_j) < U_k(p_{N-j},p_j) \quad \text{for every } k \in S$$
$$U_j(p_{N-j},p^*_j) \geq U_j(p_{N-j},p_j).$$

The plan p^*_j is said to be gainful for agent j at the expense of group S. To indicate that the gainful plan depends on the initial plan (p_{N-j},p_j), we denote it by $p_j^{gain}(S,p_{N-j},p_j)$.

Definition 4 requires that plan p^*_j be common for all agents of the group S. If every agent in S depended on j for a different plan, then he would be able to resolve his dependence separately and without the help of the other members of the group. In this case there would not be sufficient grounds to regard agents in S as a group. Therefore, if every agent in the set of agents S depends on agent j, we cannot conclude that S depends on j. That is, individual dependence does not imply group dependence. The opposite statement, however, is true.

Proposition 1. If in a joint plan (p_{N-j},p_j) a group of agents S, $S \subseteq N-j$, depends on agent j, then every member of the group depends on agent j.

Since the proofs of most of the propositions require considerable technical preparation and space, in this version of the paper the proofs are omitted. The complete proofs are provided in (Brainov, 1998).

The notion of dependence formalizes power situations in which the ability of an agent to influence other agents is unconditional and does not depend on the will or abilities of third parties. Many situations of practical importance, however, exhibit constrained power. In such situations several agents have power over an agent or group of agents. That is, the power of each powerful agent is constrained by the other powerful agents. Since the actions of the powerful agents interfere, the final effect on the dependent agent is determined by a strategic interaction between the powerful agents.

The next definition introduces the notion of conditional dependence. In a situation of conditional dependence, a powerful agent needs the consent of all other powerful agents in order to influence the dependent agent.

Definition 5. In a joint plan (p_{N-S-j},p_S, p_j) agent i *depends conditionally* on agent j with the *tacit consent* of a group S, $i \notin S$, if there exists p^*_j such that for every p'_{N-S-j}:

$$U_i(p'_{N-S-j},p_S,p^*_j) < U_i(p_{N-S-j},p_S,p_j)$$
$$U_j(p_{N-S-j},p_S,p^*_j) \geq U_j(p_{N-S-j},p_S,p_j).$$

The plan p^*_j is called conditionally gainful for agent j and is denoted by $p_j^{gain}(i,p_{N-S-j},p_S,p_j/p_S)$.

In Definition 5 the powerful agents are agent j and the group S. Agent j can gain at the expense of agent i if the group S adheres to the plan p_S. By adhering to p_S the group S does not oppose agent j. Definition 5 says that if we disregard the existence of group S, agent i depends on agent j. Therefore, the notion of conditional dependence is more general and subsumes the notion of dependence.

Consider the previous example where the resource r_j needed by agent i can be provided by agent j or agent k. In this case, in order to apply their power, agents j and k depend on each other. If agent j withholds his resource r_j, then agent k can provide agent i with the same resource. Therefore, in order to harm agent i, agent j needs the consent of agent k.

Proposition 2. If in a joint plan (p_{N-j},p_j) agent i depends on agent j, then agent i depends conditionally on agent j with the tacit consent of every group S, $S \subseteq N-i-j$.

Proposition 2 is in accordance with the intuition that if an agent has unconstrained power over another agent, then the consent of the other agents in the environment is unnecessary.

Definition 6. Agent j can *gain without harming* agent i in the plan (p_{N-j},p_j) if there exist p^*_j and p^*_i such that for every p'_{N-i-j}:

$$U_i(p'_{N-i-j},p^*_i, p^*_j) \geq U_i(p_{N-j},p_j)$$
$$U_j(p'_{N-i-j},p^*_i,p^*_j) > U_j(p_{N-j},p_j).$$

According to this definition agent i has a reply to the plan p^*_j of agent j. Therefore, after agent j has changed his plan, agent i needs only to adjust his activity in accordance with the change. For this adjustment, agent i does not rely on the other agents in the environment. It is evident that if a joint plan is Pareto optimal, then no agent can gain without harming the other agents. The opposite statement is not true. That is, a joint plan might not be Pareto optimal even if no agent can gain without harming the other agents. A joint plan is Pareto optimal if there does not exist another plan which is as good or better for all the agents and strictly better for at least one agent.

The following definition introduces a second type of group dependence, viz., the case when an agent depends on a group of agents.

Definition 7. In a joint plan (p_{N-j},p_j) agent i *depends on a group* of agents S, $i \notin S$, if

(i) (group requirement) agent i depends conditionally on every agent k, $k \in S$, with the tacit consent of the group S-k;

(ii) (minimality requirement) for every agent k, $k \in S$, every conditionally gainful plan $p_k^{gain}(i,p_{N-S},p_{S-k},p_k/p_{S-k})$ and every plan p'_{N-S}, it holds that every agent m, $m \in S-k$, can gain without harming agent i in the plan $(p'_{N-S},p_{S-k},p_k^{gain}(i,p_{N-S},p_{S-k},p_k/p_{S-k}))$

Definition 7 is justified by the intuition that only by acting together can members of the group S help or harm agent i. Every member of the group acting separately cannot exert any power over agent i. The first condition of

Definition 7 says that agent k needs the tacit consent of the rest of the group S in order to gain at the expense of agent i. The second condition states that the group S is the minimal group (in the sense of set theoretic inclusion) that has any power over agent i. That is, every member of S contributes to the power of the group, i.e., there are no dummy members.

The following example helps to clarify Definition 7. Suppose that there are three agents i, j, and k operating in the blocks world. The goal of agent i is to move block A from Position 1 to Position 2. Block A is heavy and only agents j and k acting together can pick it up. Suppose further that in the joint plan (p_i,p_j,p_k) agents j and k have committed unilaterally to move block A from Position 1 to Position 2. It is clear that in the plan (p_i,p_j,p_k) agent i depends on the group of agents j and k. Both j and k, however, can deviate unilaterally and gain at the expense of agent i. Suppose that agent j has already deviated from his commitment. Since block A will not be moved to Position 2, damage has been inflicted upon agent i. After the deviation of agent j, it does not matter to agent i whether agent k continues to adhere to his commitment. Since agent j has already deviated, agent k can also deviate from his commitment and gain without harming agent i. Thus, in the case of group dependence the first deviating agent can inflict damage upon the dependent agent. Agents deviating afterwards can gain without harming the dependent agent.

After the deviation of agent j, agent i might still depend on agent k for things other than moving block A. That is, agent k can still gain at the expense of agent i, but this gain has nothing to do with the group activity and particularly with agent j.

The following proposition follows immediately from Definition 7.

Proposition 3. If in a joint plan (p_{N-i},p_i) agent i depends on the group of agents S, $i \notin S$, then agent i conditionally depends on every member of the group S with the tacit consent of the rest of the group.

Proposition 3 can be regarded as the inverse of Proposition 1. In general, it is not true that if agent i depends on a group S, then agent i depends on every member of the group. In this case, one can only conclude that agent i depends *conditionally* on every member of the group.

If agent i depends on every member of a group, Definition 7 does not allow us to conclude that agent i depends on the group of these agents. Thus, we cannot mechanically group agents on account of their relationship to agent i. Agent i can depend on different agents for different purposes and these agents can even conflict with one another. To summarize, according to Definition 7, group activity is a necessary requirement for group formation.

Consider again the example with the three agents i, j and k operating in the blocks world. Suppose now that the goal of agent i is to move blocks B and C. Block B can be moved only by agent j and block C only by agent k. In this case agent i depends on agent j as well as on agent k. However, agent i does not depend on the group of agents j and k.

Resolution of Dependence

Definition 1 accounts for the case when a powerful agent can benefit by harming another agent. In many situations an agent can still have power over another agent but at some cost. The following definition formalizes the case when the power is associated with a cost.

Definition 8. In a joint plan (p_i,p_j) agent j has *power* over agent i if there exists plan p^*_j such that for every plan p'_i of the other agents:

$$U_i(p'_i, p^*_j) < U_i(p_i,p_j)$$

The cost of power $c_j^{power}(i, p_i,p_j)$ is defined by:

$$c_j^{power}(i, p_i,p_j) = U_j(p_i,p_j) - U_j(p_i,p^*_j).$$

The damage $d_i^{depend}(j, p_i,p_j)$ of agent i and is defined by:

$$d_i^{depend}(j, p_i,p_j) = U_i(p_i,p_j) - U_i(p_i, p^*_j),$$

Definition 8 is a generalization of Definition 1. That is, if agent i depends on agent j, then agent j has power over agent i. In Definition 8 we do not require the plan p^*_j be gainful for agent j. It is sufficient that agent j can harm agent i and thereby can influence his behavior. Consider the following example. Agent *m*, $m \in N$, is executing his plan and according to it he is going to use some resource r_s. The resource r_s is shared between him and agent *n*, $n \in N$. Agent *n* is idle. He anticipates that agent *m* will use r_s and occupies it before agent *m* accomplishes his current task. By doing so agent *n* bears some negligible costs c_n. Suppose that the resource r_s can be used only by one agent at a time and by occupying it, agent *n* does not deplete it. If the goal of agent *m* is valuable enough to him, agent *m* will be willing to pay agent *n* some compensation for the right to use the resource first. If the compensation exceeds the cost c_n, agent *n* will benefit from occupying the resource r_s. In this example agent *n* exerts some power over agent *m*. The source of the power is agent *n*'s ability to use the resource r_s at any time.

Since self-interested agents are utility maximizers they should try to resolve each power relation. That is, the powerful agent should attempt to receive some compensation for not exerting his power and the dependent agent should be willing to offer compensation to avoid harm. The amount and the form of compensation should be determined e.g. by a negotiation between the powerful and the dependent agent.

Consider the unilateral power situation when agent j has power over agent i in the joint plan (p_i,p_j). Such a power situation can be resolved if:

$$c_j^{power}(i,p_i,p_j) < d_i^{depend}(j,p_i,p_j) \quad (3)$$

That is, the situation can be resolved, if the cost of damaging is less than the damage. Agent i is willing to pay agent j at most $d_i^{depend}(j,p_i,p_j)$ and agent j is willing to receive at least $c_j^{power}(i,p_i,p_j)$. Therefore, if inequality (3) holds, then a mutually acceptable side payment exists.

Definition 9. The *amount of power* $A_{ji}(p_i,p_j)$ of agent j over agent i in a joint plan (p_i,p_j) is defined by:

$$A_{ji}(p_i,p_j) = d_i^{depend}(j,p_i,p_j)- c_j^{power}(i,p_i,p_j).$$

That is, the amount of power that agent j wields is the difference between the damage and the cost of power. It follows from (3) that a unilateral power situation can be resolved if the amount of power is strictly positive.

Consider the bilateral power situation where in the joint plan (p_i,p_j) agent i has power over agent j and agent j has power over agent i. If

$$A_{ji}(p_i,p_j) \neq A_{ij}(p_i,p_j),$$

than the power is balanced, so no side payment is necessary. Otherwise, the situation can be resolved by having the less powerful agent pay the more powerful agent some compensation.

Throughout this paper we assume that every dependence can be resolved permanently. More formally, if agent i depends on agent j in the joint plan (p_{N-j},p_j) and $p_j^{gain}(i,p_{N-j},p_j)$ is a gainful plan for agent j, then in every joint plan $(p'_{N-j}, p_j^{gain}(i,p_{N-j},p_j))$ agent i does not depend any more on agent j. Thus, we eliminate the case of total dependence. Total dependence occurs when, after resolving a particular dependence, the dependent agent or group of agents find themselves in the same dependent position. Suppose that agent i depends on agent j in the joint plan (p_{N-j},p_j). Agent j can apply his gainful plan $p_j^{gain}(i,p_{N-j},p_j)$ thereby harming agent i. Let p^*_{N-j} be the reaction of all other agents to the plan $p_j^{gain}(i,p_{N-j},p_j)$. That is, agents arrive at the plan $(p^*_{N-j}, p_j^{gain}(i,p_{N-j},p_j))$. If in this plan agent i depends again on agent j, then there exists a second gainful plan for agent j: $p_j^{gain}(i,p^*_{N-j}, p_j^{gain}(i,p_{N-j},p_j))$. After agent j has applied the new gainful plan, agent i finds himself again dependent on agent j, etc., ad infinitum. Thus, we obtain a sequence of gainful plans for agent j: $p_1,p_2,p_3,\ldots$. Each plan in this sequence results from applying the previous one and gives rise to the next one. Therefore, the power of agent j over agent i is total. Dependence based on total power cannot always be resolved. It can be proved, however, that if for every sequence of gainful plans $p_1,p_2,p_3,\ldots$, the sequence of losses incurred by agent i is convergent and its limit belongs to some set bounded from above, then the total dependence can be resolved (Brainov 1998).

Stability of Multiagent Plans

In this section we move to the case where agents cannot make side payments. Therefore, the only way for a dependent agent to resolve a dependence relation is by changing his plan. One agent's change of plan (even before execution) can cause other agents to change their plans which induces further changes by others, etc. Therefore, the question of stability of a multiagent plan becomes crucial.

In this section we analyze the dependence structure of multiagent plans and provide necessary and sufficient conditions for stability. In order to approach the problem of stability of multiagent plans we need some preliminary notions, in particular the notions of individual and coalitional stability. A joint plan is individually stable if every agent's plan is a best response to the plans of other agents. Formally,

Definition 10. A joint plan (p_{N-i},p_i) is *individually stable* if there exists no agent i, $i\in N$, and a plan p^*_i such that:

$$U_i(p_{N-i},p^*_i) > U_i(p_{N-i},p_i).$$

Thus individual stability eliminates incentives for unilateral deviations. Group deviations are captured by the notion of coalitional stability. A joint plan is coalitionally stable if no coalition, taking the actions of all other agents as fixed, can deviate in a way that benefits all its members.

Definition 11. (Aumann 1959) A joint plan p* is *coalitionally stable* if for every group of agents S, $S\subseteq N$, and every plan p_S of the group S there exists an agent i, $i\in S$, such that:

$$U_i(p^*)>U_i(p_S,p^*_{N-S}).$$

That is, every attempt to deviate is opposed by at least one member of the deviating group. This definition of coalitional stability is known as Strong Nash equilibrium.

Hereafter we assume that agents are able to revise their plans dynamically based on the actions of the other agents. We suppose that the actions of all agents are observable. We also assume that the cost of achieving agents' goals increases with the time when replanning occurs. In the context of this paper replanning is usually a result of other agents' actions. The later an agent realizes that the other agents will not help him or will obstruct his current plan, the greater is the number of his previous actions that usually become obsolete. That is, the later an agent replans, the higher are his costs.

The following proposition provides necessary conditions for stability of 2-agent plans.

Proposition 4. If a plan (p_i,p_j) is Pareto optimal, agent i depends on agent j, and agent j depends on agent i, then the plan (p_i,p_j) is individually and coalitionally stable.

Proposition 4 says that if in a bilateral power situation the power of the agents is balanced, no agent or group of agents is willing to change its current plan. If a bilateral situation has only one powerful agent, it might be unstable. In such a situation the powerful agent might gain at the expense of the dependent agent. According to the next proposition, if a bilateral situation does not have powerful agents, it is stable.

Proposition 5. If a plan (p_i,p_j) is Pareto optimal, agent i does not depend on agent j, and agent j does not depend on agent i, then the plan (p_i,p_j) is individually and coalitionally stable.

Proposition 4 can be generalized to the case of more than two agents. In this case, the reciprocal dependence is a balanced power structure. In such a structure everybody depends on everybody directly or transitively. Therefore, every attempt to gain at the expense of somebody else will be opposed by the rest of the agents.

Proposition 6. Every Pareto optimal joint plan which is based on reciprocal dependence is individually and coalitionally stable.

The next definition accounts for multiagent plans that contain at least one dependence relation. Special attention has to be given to such plans, since the existence of dependence is a potential source of instability.

Definition 12. A joint plan $(p_1,\ldots,p_n)$ is *based on dependence* if there exist agents i and j, such that agent i depends on agent j.

Proposition 7 provides necessary and sufficient conditions for individual stability of a multiagent plan based on dependence. Surprisingly, the only dependence structure that can guarantee individual stability is reciprocal dependence.

Proposition 7. Every joint plan $(p_1,\ldots,p_n)$ which is based on dependence is individually stable if and only if it is based on reciprocal dependence.

Propositions 8 and 9 refer to the case of group dependence. According to Proposition 8, if in a joint plan an agent depends on a stable group and the group depends on that agent, then the plan is stable.

Proposition 8. If in a Pareto optimal joint plan (p_S,p_j):
(i) p_S is based on reciprocal dependence,
(ii) group S depends on agent j,
(iii) agent j depends on the group S,
then the plan (p_S,p_j) is individually and coalitionally stable.

The dependence between agent j and the group S can be thought of as a reciprocal dependence between an agent and a group of agents. Proposition 9 provides a stronger result than Proposition 8. In Proposition 9 the dependence of agent j on the group S is replaced by an ordinary dependence.

Proposition 9. If in a Pareto optimal joint plan (p_S,p_j):
(i) p_S is based on reciprocal dependence,
(ii) group S depends on agent j,
(iii) agent j depends on some agent i, $i\in S$,
then the plan (p_S,p_j) is individually and coalitionally stable.

Proposition 9 says that if power is well balanced in every group of agents and between groups, then the joint plan is stable.

Conclusions

In this paper a decision-theoretic model of social dependence and social power was presented. The model enables self-interested agents to recognize the means and the amount of their power and to influence other agents' behavior. The model was applied to the analysis of stability of multiagent plans. It was shown how dependence can guarantee stability.

In contrast to game theory and sociology where power is defined as general ability to influence someone's behavior, in our approach power is related to a particular multiagent plan. Thus, the notion of dependence captures all the possibilities for strategic interaction available in a particular environment. In addition, we differentiate between the notion of power potential and the exercising of power.

The model proposed in the paper forms a basis for agents to perform deliberative actions in order to increase the amount and to enlarge the means of their power. That is, to transform their power potential into existing power.

Acknowledgements

The authors would like to thank Prof. V. Khoroshevsky from the Computer Center of Russian Academy of Sciences for his supervision of the research reported in this paper.

References

Aumann, R. 1959. Acceptable Points in General Cooperative N-persons Games. In *Contributions to the Theory of Games IV*. Princeton University Press.

Brainov, S. B. 1994. Deviation-Proof Plans in Open Multiagent Environments. In *Proceedings of the European Conference on Artificial Intelligence*, 274-278.

Brainov, S. B. 1996. Altruistic Cooperation Between Self-Interested Agents. In *Proceedings of the European Conference on Artificial Intelligence*, 519-523.

Brainov, S. B. 1998. Models of Stable Cooperation in Multiagent Systems. Ph.D. diss., Computer Center of Russian Academy of Sciences, Moscow.

Castelfranchi, C. 1990. Social Power. In Demazeau Y. and Muller J.-P. eds. *Decentralized AI*, 49-62. Elsevier.

Cohen, P. R., Levesque, H. J. 1990. Intention is Choice with Commitment. *Artificial Intelligence*, 42:213-261.

Coleman, J. 1973. *The Mathematics of Collective Action*. Aldine Publishing Company.

D'Iverno, M., Luck, M., Wooldridge, M. 1997. Cooperation Structures. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 600-605.

Harsanyi, J. 1962. Measurement of Social Power, Opportunity Costs, and the Theory of Two-Person Bargaining Games. *Behavioral Science* 7:67-80.

Levesque, H. J., Cohen, P. R., Nunes, J. H. T. 1990. On Acting Together. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 94-99. AAAI Press.

Sandholm, T. W., Lesser, V. R. 1995. Equilibrium Analysis of the Possibilities of Unenforced Exchange in Multiagent Systems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 694-701.

Sichman, J., Demazeau, Y., Conte, R., Castelfranchi, C. 1994. A Social Reasoning Mechanism Based on Dependence Networks, In *Proceedings of the European Conference on Artificial Intelligence*, 274-278. John Wiley&Sons.