Using Social Power to Enable Agents to Reason About Being Part of a Group

Cosmin Carabelea[1], Olivier Boissier[1], and Cristiano Castelfranchi[2]

[1] SMA-G2I-ENS Mines de St.Etienne,
158 Cours Fauriel, St.Etienne, F-42030, France
{carabelea, boissier}@emse.fr
[2] ISTC-CNR-Rome, Viale Marx 15, 00137, Rome, Italy
c.castelfranchi@istc.cnr.it

Abstract. One of the main challenges in multi-agent systems is the coordination of autonomous agents. In order to achieve this coordination, the agents are considered to be part of what we call a group (e.g., organization, institution, team, normative society, etc.). Our goal is to enable an agent to reason about the implications of being part of a group: what does it gain or lose, what are the constraints imposed on its behaviour. The theory of social power has been proposed as a paradigm to describe the agent's behaviour. In this paper we use this theory, we formalize it and we extend it to include group-related aspects. We then show how, using this theory, an agent is able to reason about the constraints imposed on its behaviour by the group, for example to decide whether it should enter or not a group.

1 Introduction

One of the main challenges in multi-agent systems is the coordination of autonomous agents. Many mechanisms have been proposed to achieve this coordination. Among them, of interest to our work, we can cite organizational design [17] and norms [1]. In this paper we will use the term *group* to denote a collection of agents coordinated by one or more of these mechanisms. Particular examples of types of groups are institutions [8], teams [16] or normative societies [1]. Generally speaking, a group contains several elements (e.g., agents, roles, norms) and several coordination and control mechanisms. These mechanisms are used, for example, to assign agents to roles [6], or for group's decision-making (e.g., choosing group's goals and distributing them among members) [16], to negotiate or to plan [12], to punish agents that violate norms [1], etc.

Besides designing coordination and control mechanisms for autonomous agents, another difficult problem in open systems is to enable agents to reason about these mechanisms, or, more generally, to reason about what does it mean to be part of a group. For example, when an agent enters a group, how does it understand what it will lose, what it will gain, what new constraints are imposed on its behaviour? *Our aim in this work is to enable an agent to reason about the implications of being part of a group.*

M.-P. Gleizes, A. Omicini, and F. Zambonelli (Eds.): ESAW 2004, LNAI 3451, pp. 166–177, 2005.

Several authors have proposed the notion of *social power* as a paradigm that can be used to explain the agents' behaviour (see for example [5] or [11]). One of the uses of this theory of social power is to enable the designer (or an observer) of a multi-agent system to analyze and predict the agents' behaviour in a group. Another interesting use of this theory is to enable the agents to reason about their (and others') powers when being part of the group. Thus, they will be able to understand the implications (i.e., the constraints they face) of being part of a group. However, in order to be used by artificial agents, the theory of social power needs to be formalized.

Based on the work in [4], in this paper we propose formal definitions for several forms of power, ranging from individual and social (Section 2) to group-related ones (Section 3). In Section 4 we describe how an agent can use this power theory to reason about what it means to be part of a group and to decide whether to be part of a group or not. Finally, we draw some concluding remarks and trace directions for future work.

2 Non-institutional Powers

The author of [5] argues that the notion of *power* is not intrinsically social and does not refer only to the theory of society or at least of social action: power can be related to individual. Using as basis the work described in [15], we use the basic notions of action, resource, plan and goal (noted respectively *a, r, pl* and *g*) to describe the behaviour of an agent. These notions will be used to define the *individual powers* of an agent, i.e., the powers it has without considering other agents. We will then use these powers to identify power relationships between agents, to address what we call *social powers*. Although all the following definitions are context dependent, to keep them simple, we will not use in the different formulae an additional parameter representing the context. However, we would like to stress out that an agent can have a power in a context and not have it in another.

Due to space reasons, we will not describe here all the predicates used in our definitions. Most of these predicates have self explanatory names, e.g., *believes*, *can_empower* (an agent can give a permission to another one), *can_commit_to* (an agent can form an intention or commitment to do something), etc. A more detailed description of all the used predicates can be found in [4].

2.1 Individual Powers

An agent can have several types of individual powers: executional powers (*can_do*), deontic powers (*entitled_to*) and full (total) powers (*power*). In what follows we will use as examples only actions and goals. Nevertheless, the notions we introduce can be easily extended for resources or plans.

2.1.1 Executional and Deontic Powers: *Can_do* and *Entitled_to*

We note by *can_do* the executional power an agent has. For example, we say that an agent *can do* an action, if it has all the needed resources, both external and internal (know-how, skills) to execute that action. Or we say that an agent *can do* a goal if it has a plan (or it can obtain one) to achieve the goal and it *can do* all the subgoals, actions or resources needed by that plan:

$$\text{can\_do}(X, r) =^{d} \text{has}(X, r)$$
$$\text{can\_do}(X, a) =^{d} \text{know\_how}(X, a) \wedge (\forall r\ \text{needs}(a, r) \Rightarrow \text{can\_do}(X, r))$$
$$\text{can\_do}(X, pl_g) =^{d} \forall \alpha \in pl_g\ \text{can\_do}(X, \alpha),\ \alpha \text{ - a resource, action or goal in the plan}$$
$$\text{can\_do}(X, g) =^{d} \exists\ pl_g\ \text{achieves}(pl_g, g) \wedge \text{can\_do}(X, pl_g) \quad (1)$$

Besides the executional power, an agent can have a deontic power to execute an action, to achieve a goal, etc. We note this by *entitled_to*. Its meaning is that the agent has all the necessary permissions to execute the action and use the necessary resources. Or, in the case of a goal, the agent is *entitled to* achieve the goal if it has the permissions to pursue the goal and it is entitled to do all the elements of a plan that achieves the goal. The *entitled_to* definitions for resources, actions and goals are similar with the ones above and they are not given here due to space reasons.

Let's take for example an agent A that has the resource *printer*, the only resource needed by the action *print*, and the agent knows how to print documents. We can thus say that *can_do(A, print)*. Even if A knows how to print a document, one must also specify if A is forbidden (or allowed) to use the printer or to print using it. For example, we can say that A has all the rights to print a document using that printer: *entitled_to(A, print)*.

2.1.2 Full Powers: *Power*

If an agent has both the executional (*can_do*) and deontic (*entitled_to*) powers, then we say that it has the *power_of* executing/achieving the action/goal:

$$\text{power\_of}(X, a/g) =^{d} \text{can\_do}(X, a/g) \wedge \text{entitled\_to}(X, a/g) \quad (2)$$

The above definitions are not sufficient to fully describe a power of an agent. If we consider the agent A from the previous example, it can print and it has the permission for it, so it has the power of doing so: *power_of(A, print)*. A first condition for a cognitive agent to really have a power is to be aware that it has that power (otherwise it cannot use it). Thus, in order to have a full power, an agent must be aware it has that power and it also must be able to form an intention to (or to commit to) use its power. If an agent cannot commit whenever it wants to use its power, it does not really have the full power. This is the reason why we introduce the notion of *having access to a power*, which allows us to define what having a full power (noted *power*) means:

$$\text{access\_to}(X, p) =^{d} \text{believes}(X, \text{power\_of}(X,p)) \wedge \text{can\_commit\_to}(X, p) \quad (3)$$

$$\textbf{power(X, p)} =^{d} \text{power\_of}(X,p) \wedge \text{access\_to}(X, p) \quad (4)$$

2.2 Social Powers: Dependence, Power over and Influencing Power

The notions introduced in the previous section are taken from the point of view of an agent and do not take into consideration the other agents in the system. Very often an agent X lacks the *power_of* executing an action or achieving a goal. There are many reasons for this, among them being that it lacks some resources or the knowledge how to do it, or even the permissions to do it. If there is another agent, Y, that can provide X with the thing it lacks, we say that X depends on Y for the execution of the action or the achievement of the goal and we write *depends_on(X, Y, a/g)*. For example, in the case of a goal and a lack of executional power:

$$\text{depends\_on}(X, Y, g) =^{d} \forall pl_g \ \text{achieves}(pl_g, g) \Rightarrow \exists \alpha \in pl_g \ \neg \text{can\_do}(X, \alpha) \wedge \text{can\_do}(Y, \alpha) \quad (5)$$

Or, in the case where the power X lacks is a deontic power:

$$\text{depends\_on}(X, Y, g) =^{d} \neg \text{entitled\_to}(X, g) \Rightarrow \text{can\_empower}(Y, X, g) \quad (6)$$

In the previous example, the agent A has the power of printing: *power_of(A, print)*. Let's assume there is another agent, B, that has the goal of printing a document D, but it does not know how. If A is the only agent having the power of printing, then B depends on A to achieve its goal: *depends_on(B, A, print document D)*.

Due to space reasons, we will not enter into details here; the dependence relationship is deeply analyzed in [15], together with dependencies towards several agents (OR- and AND-dependence) or the mutual or reciprocal dependences. As seen above, our approach extends the existing theory of dependence by taking into account the deontic aspect: an agent can depend on another due to the lack of deontic power.

We would like to stress out that the dependence relationship is objective, i.e., it exists even if the agents are not aware of it. However, if X believes that it depends on Y for an action/goal, then we say that Y has a *power over* X for that action/goal:

$$\text{power\_over}(Y, X, a/g) =^{d} \text{believes}(X, \text{depends\_on}(X, Y, a/g)) \quad (7)$$

The last power relationship we introduce here is the *influencing power* an agent has over another one. From the social point of view (i.e., not institutional and not normative), an agent has influencing power over another agent, if the other agent is aware that it depends on the first for one of its goals:

$$\text{infl\_power}(Y,X,p) =^{\mathbf{d}} \text{power\_of}(X, p) \wedge (\exists g \ \text{goal}(X,g) \wedge \text{power\_over}(Y, X, g)) \quad (8)$$

In this definition, *p* is a power of X that is not necessarily related to the action or the goal for which X depends on Y. Using the above example, imagine that B also has the power of negotiating objects in e-markets: *power_of(B, negotiate)*. Because B depends on A for its goal to print – *depends_on(B, A, print)* – and both agents are aware of this situation, A acquires the power of influencing B for the power related to negotiation: *infl_power(A, B, power_of(B, negotiate))*.

As we will discuss in Section 3, there are other ways in which an agent can obtain a power of influencing another agent. In the following we will describe how an agent can use its power of influencing over another agent to acquire new powers.

2.3 Putting at the Disposal of

There are many operations the agents can execute in order to modify their powers and/or the powers of another agent, like empowerment, putting at the disposal of, requesting power, etc. (see [5] and [4]). We will focus here on the two operations related to the influencing power: the *putting of a power at the disposal of another agent* and the *request for a putting at the disposal*. But first, we will introduce what we call an indirect power.

The full power of an agent, presented in the previous sections is a *direct* power of an agent. The agent has the power of executing an action (or to satisfy a goal) and it also has *access to* its power. We say that an agent has an *indirect power* if it has access to the power of another agent. The predicate used bellow, *transferred_to* denotes the fact that an agent has transferred the access to its own power to another agent. The indirect power means that an agent cannot use it directly, but if it commits to use it, the other agent will do it:

$$\text{ind\_power}(Y, p) =^{d} \exists X\!: \text{power\_of}(X, p) \land \text{transferred\_to}(X, Y, p) \quad (9)$$

When an agent X that has a full power gives to another agent Y access to its power, we say that X *puts at the disposal of* Y its power. Thus, Y is able to use this power whenever it wants to:

$$\text{power}(X, p) \land \text{puts\_at\_disposal\_of}(X, Y, p) \Rightarrow \text{ind\_power}(Y, p) \quad (10)$$

We can note that in this definition, X does not offer its power p to Y (e.g., a resource), but the access to the power p (e.g., the access to the resource). In other words, Y has the control over the power of X: it can decide whenever it wants to use (indirectly) the power.

An interesting question is whether X still keeps its direct power of p or it loses it when putting it at the disposal of Y. The nature of the *access_to* predicate makes that X loses its access to its power:

$$\begin{aligned}&\text{power\_of}(X, a/g) \land \text{access\_to}(X, a/g) \land \text{puts\_at\_disposal\_of}(X, Y, a/g) \Rightarrow \\ &\quad \text{power\_of}(X, a/g) \land \neg\text{access\_to}(X, a/g) \land \text{transferred\_to}(X, Y, a/g)\end{aligned} \quad (11)$$

If an agent has the power of influencing another agent for a power and it wants that power, it requests the second agent to put at its disposal the power:

$$\begin{aligned}&\text{power}(X, p) \land \text{infl\_power}(Y, X, p) \land \text{requests\_power}(Y, X, p) \Rightarrow \\ &\quad \text{puts\_at\_disposal\_of}(X, Y, p)\end{aligned} \quad (12)$$

Thus, by using the Formula 11, the *request_power* operation has the effect of obtaining an indirect power. Using the previous example, due to a dependence, the agent A has the power of influencing B for B's power of negotiating: *infl_power(A, B, power_of(B, negotiate))*. If A requests B to put at its disposal its power to negotiate, A will obtain an indirect power to negotiate: whenever A decides to, B will negotiate on its behalf – *ind_power (A, negotiate)*. This example illustrates how, by reasoning upon the dependencies existing between them, the agents in a multi-agent system are able to acquire new powers. The next section will show how the fact that they belong to a group enables the agents to obtain even more powers.

While we are aware of the fact that there are autonomous agents that can refuse a *request_power*, in this paper we will not focus on this aspect of the power theory. The interested reader is invited to [4] for more details on how agent autonomy can be defined by using the power theory.

3 Agents' Powers in a Group

The previous section described the various forms of powers an agent have by itself (individual powers) and the power relationships that appear because of dependencies between agents. All the definitions above were given without taking into account institutional aspects, i.e., the fact that the agents belong to a group. In the following we will describe what are the institutional aspects we take into account and we will show how agent's powers are influenced by them.

3.1 What Is a Group Made of ?

There are many approaches on how to coordinate autonomous agents. The top-down ones usually specify an explicit organizational structure that is known to all agents (e.g., the MOISE+ model [9]), while the bottom-up ones try to ensure the coordination without using an organizational structure known by all agents (e.g., GPGP/TAEMS [12]). In this paper we call *group* a collection of agents coordinated by one or more coordination mechanisms. We use this term to denote the common things in the notions of institution [8], organizational structure [3] or [17], team [16], normative society [1], etc. A group contains several elements (e.g., roles, norms) and several coordinations and control mechanisms (e.g., assigning agents to roles [6], punishing agents that violate norms [1], etc.).

In this paper we focus on how the group's elements influence the agents' powers, i.e., what new forms of power appear because the agents belong to a group. Based on the group models proposed in the literature (institutions, normative societies, etc.), we consider a group formed by three types of elements. A group contains a set of *roles*, a set of *authority* relationships between the roles (i.e, a hierarchy of roles) and a set of deontic specifications of the roles' expected behaviour (i.e., *norms*). In the following we will show how this model, although simple, allows us to model the powers that appear in a group. We will then show that the power relationships we introduced can be extended to other coordination elements, such as the *contract* [2] (or its weaker form, the *commitment*) [10].

We would like to point out that in this paper we do not propose formal definitions for the notions we use to define a group (norms, roles, etc.). Defining these notions is a very active research area in the multi-agent community (see for example [1], [6] or [17]). The aim of this paper is not to replace this related work, but to complement it by defining agents' powers using these notions.

3.2 Power Relations in a Group

In Section 2 we introduced the notion of *power over* that describes the power an agent has over another. The source of this power over is the goal (action) dependence of the second agent on the first (Formulae 5 and 6). However, if the agents are part of a group, there are two other sources for the power of an agent over another: the authority relationships and the norms.

3.2.1 Power over Due to *Authority over*

An agent can have power over another because it has authority over that agent. In other words, if there are two roles and there is an authority relation between them, the authority relation will exist between any agents playing the roles and this gives an agent a power over another:

$$\text{power\_over}(Y, X, p) =^{d} \exists\, R_1, R_2\text{: plays}(X, R_1) \wedge \text{plays}(Y, R_2) \wedge \text{authority\_over}(R_1, R_2, p) \tag{13}$$

If both agents are aware of the roles they play and of the relations between their roles, then Y has also the power of influencing X (Formula 8). If it requests this power *p* (Formula 12) and X puts it at Y's disposal, then Y acquires the indirect power regarding X's power *p*. Thus, whenever Y wants, X will use its power p.

For example, in the case of a soccer team, there are two roles: *coach* and *player*. Some agents can play only using their right foot, some using their left, and some using both. We model this by saying that an agent has the *power_of(A, play ball with left/right foot)*. We also consider that an agent should be able to play a given move (e.g., pass to another player, run to a given position, receive ball and shoot) – *power_of(A, play move)*. The hierarchical structure in this group specifies that the role coach has authority over the role player for the power of a player to play a move: *authority_over(coach, player, power_of(player, play move))*. If an agent B plays the role coach and an agent A plays the role player, then *power_over(B, A, power_of(A, play move)*. Using the reasoning above, B acquires the indirect power for *power_of(A, play move*) and thus, whenever B (coach) wants, A (player) plays a move.

We would like to point out that the *authority_over* relation (and thus the *power_over* one) depends on the power *p*: a role can have authority over another role for a power, but not have it for another. Using the example above, the authority relation does not say anything about the power of an agent to play the ball with its right/left foot. Thus, the agent playing the role coach cannot decide whenever it wants with what foot a player will play the ball.

3.2.2 Power over Due to *Norms*

Another source for the power over relation is the existence of a norm. If an agent plays a role and there is a norm targeting that role and concerning a power p, then any agent that plays that role has its behaviour regarding p constrained by the norm:

$$\text{power\_over}(\textit{group}, X, p) =^{d} \exists R\ \text{belongs\_to}(R, \textit{group}) \wedge \text{plays}(X, R) \wedge \text{has\_norm\_on}(\textit{group}, R, p) \tag{14}$$

The same reasoning concerning the influencing power and the requesting of a power can be done for this form of the power_over. Thus, the group can acquire an indirect power regarding *p*: whenever the group wants, the agent X uses its power p. Using the example above, let's say there is an obligation in the soccer team that all players use the same equipment. This norm targets the role *player* and the power_of a player to wear a given equipment, and thus we have *power_over(soccer team, A, power_of(A, wear equipment))*, where A is an agent playing the role *player*. If the group requests this power from the agent and the agent puts it at the group's disposal, then whenever the group decides that A wears a given equipment, A wears.

It is out of the scope of this paper to talk about agent autonomy, more precisely about norm-autonomy [7]. However, we would like to point out that if an agent refuses the requests_power operation, it does not put at the disposal of the group its power_of. Thus, the decision on what equipment the agent will wear will not be made by the group, but by the agent.

This definition of power_over differs from the others in the sense that it is not another agent which has a power over an agent, but it is the group that has this power. Often, the group *empowers* an agent to act as its representative and thus this agent acquires the power over, but this is just a transfer of power: it is the group that has real power over. A good example of this situation is given in [5]: an agent does not surrender herself to a policeman, but to the institution represented by the policeman.

3.3 Other Sources for the Power over

In this paper we chose to describe a group in terms of roles, authority relations and norms. Other terms have been used as well to define agent institutions or societies. For example, the authors of [2] use the notion of *contract* as a coordination element in a group. We can extend the above definitions by defining the power of an agent over another due to a contract between them:

$$\text{power\_over}(Y, X, p) =^{d} \exists C\text{: contract}(X, Y, p, C) \tag{15}$$

This is to say that Y has the power over X for p because X signed a contract with Y regarding p. It is obvious that this is just a particular form of the norm-based definition of power_over (Formula 14). Y has this power over because both agents are part of a group in which there is a norm that says that all contracts must be fulfilled. However, in some situations, it can be useful to model this situation directly using the notion of contract and not by passing through the intermediary of norms.

The notion of contract can be weakened to the one of *social commitment* [10][14]. We then say that an agent has power over another because the second is committed towards the first for a power:

$$\text{power\_over}(Y, X, p) =^{d} \text{committed\_to}(X, Y, p) \tag{16}$$

If an agent X commits to an agent Y for a power p and Y is aware of this, then Y has the power of influencing X for p (Formula 8). It is likely that Y will try to use this influencing power and to request X to put at its disposal the power p (Formula 11). It is still an open question for us whether to use the above definition only when the agents are part of a group or even if there is no group with norms that enforce the fulfillment of commitments. If there is no such norm, there is nothing that ensures that X will fulfill its commitment and thus will put at the disposal of Y its power p.

4 Reasoning About Entering a Group Using the Power Theory

In this paper we are interested in how to enable agents to reason about the existing relationships with other agents and with the society in which they evolve. A similar attempt has been done in [13], where the authors propose a model of normative multi-agent systems and identify the normative constraints imposed to agents. While we generally agree with their approach, we believe their model lacks the capacity of capturing the dynamicity of these constraints, e.g., a norm that is not always active. We believe that the power theory is a suitable tool to model these constraints and that this theory can be tailored to different organization models. In this section we will show how an agent can use the notions introduced in the previous sections to reason about what it means to be part of a group. Such an agent will then be able to decide whether to enter or not a group.

Even if an agent is not part of a group, it still interacts with other agents. It is to be expected that the agent is not self-sufficient, i.e., its powers do not allow it to achieve all its goals. The agent depends on other agents to achieve its goals, so the other agents have power over this agent (see for example the agents A and D in Figure 1). If the others are aware of the existing dependencies, they will have the power of influencing the agent. Thus, there are constraints imposed on the behaviour of the agent because it lacks some of the powers it needs.

When an agent is part of a group, it is also subject of several *power_over* relations. If it plays a role in that group and the group specifies a hierarchy of roles, it is to be expected that some agents playing the appropriate roles will have power over this agent. For example, as Figure 1 shows, agent A belongs to a group G where it plays the role R2. Since the agent B plays the role R1 and there is an *authority_over* relation from R1 to R2, A can compute that B has *power_over* A (cf. Formula 13). For the same reason, A has *power_over* C. Moreover, since the group is a normative one, i.e., the roles' behaviour is regulated by norms (see link *has_norm_on* from group to roles), the group has *power_over* the agents because of the norms (cf. Formula 14). As argued in the previous section, these *power_over* relations are in fact relations of influencing power. Thus, the fact that an agent belongs to a group and plays a role in that group greatly limits the liberty of the agent: other agents (or the group) will have the power of influencing this agent, hence imposing constraints on its behaviour.

In both of situations described above, when an agent does and does not belong to a group, the agent's behaviour is constrained by other agents. However, in the second case, if the agent is not part of the group, the constraints on its behaviour are due to the insufficiency of its powers. When it is part of a group, to these constraints are added those due to the fact that it belongs to a group and thus other agents have the power of influencing it. If those agents request its powers and the agent puts them at their disposal, it will lose access to its powers.

So why should an agent enter a group then, if this means that it loses powers? The answer is that by entering a group, an agent also gains power. The most direct form of power gained is because an agent entering a group usually receives resources (e.g., a payment for playing a role) and *permissions* to use resources of the group (e.g., the permission to use a printer). These powers gained by the agent are represented in

Figure 1 by the *power_of* arrow towards some resources. An agent loses some powers because it plays a role in a hierarchy of roles and thus other agents will have authority over it. But the agent might have authority over other agents, so it will gain access to their powers (e.g., relationship between A and C in Figure 1).

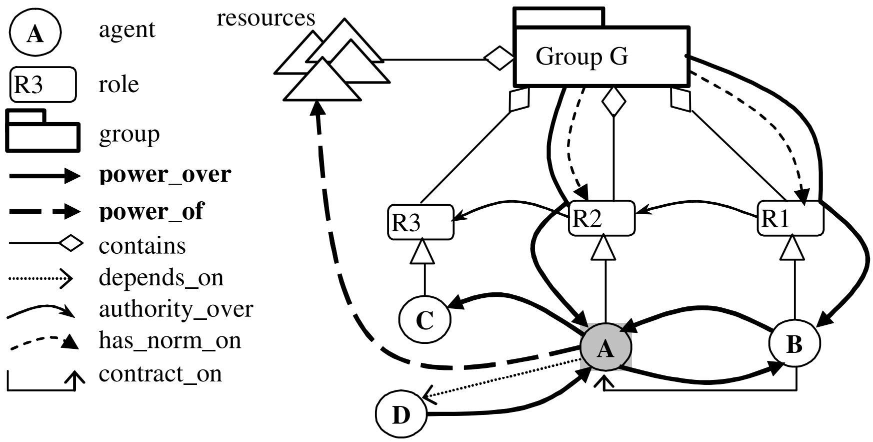

Fig. 1. Several *power_over* relations from the point of view of agent A in a group G

However, the main reason for an agent to enter a group is because of the security the group provides. The agent's behaviour will be constrained by norms, but so will be the behaviour of the other agents too. Even if the agent loses some powers because of the norms, the same powers will be lost by the other agents too. For example, no agent will be able to cheat the others. In Figure 1, this situation is represented by the power that A has over B because of a contract between B and A. Because B's behaviour is regulated by the norms too, B will have to obey the contract it has signed with A (or to disobey and suffer the consequences) since these norms impose such a behaviour.

It is clear for us that an agent facing the decision whether to enter or not a group must be able to do more than just recognize the group's structure or norms. It must be able to understand how the fact that it will be part of the group will hinder or facilitate the achievement of its goals. By using the power theory and reasoning about the power relationships that will appear within the group, the agent will be able to decide whether it is worth or not to enter the group, i.e., whether the power it loses are compensated by the ones it gains or not.

It is worth mentioning that the power theory can be used by the group to reason whether to allow or not an agent to enter a group (or to play a role). If the group has goals and delegates some of them (or subgoals of them) to a role, the agent playing the role should have at least the power of achieving those goals. Thus, reasoning about the powers of an agent, a group computes the utility of that agent for the group.

5 Conclusions

Using the social power theory, a designer (or an observer) of a multi-agent system can analyze and predict the agents' behaviour in a group. Moreover, the same theory can be used by the agents themselves to reason about the constraints imposed to their behaviour when being part of a group. However, in order to enable the agents to use it, one must formalize this theory. In this paper we proposed formal definitions for several forms of power, ranging from individual and social to group-related ones. In order to validate our approach, in the near future we intend to describe a multi-agent application in terms of powers and to enable the agents within the application to reason about their powers. Towards this aim, we started to implement in Prolog a power-based reasoning mechanism. The MOISE+ model [9] allows the specification of organizational structures and it will be used to specify the group-related aspects.

Placing agents together in a group (institution, organization, normative society) does not solve all the problems of coordinating autonomous agents. One must also provide solutions to various problems, such as assigning agents to roles in a group [6], making decisions in a group (e.g., choosing group's goals, assigning subgoals to members) [12][16], punishing norm violation [1], etc. Our power-based model does not try to substitute these solutions, but to provide a complementarity to them. For example, if the group is a normative one (i.e., there are norms constraining the agents' behaviour), norm enforcing mechanisms must be present (e.g., penalties, lowering reputations, etc.). Our model enables agent to be aware of how the norms limit their powers and thus of the constraints imposed on their behaviour by these norms. However, if an autonomous agent decides to disobey a norm, it is the norm enforcing mechanisms that is used to punish it.

As pointed out in [4], there is a direct relationship between autonomy and power. We intend to investigate this relationship and to propose formal definitions for several types of agent autonomy, such as social- or norm-autonomy. Using these definitions, the designer of a multi-agent system will be able to describe the system in terms of power and to detect what degrees of autonomy are necessary for the agents within.

Institutional empowerment is an important operation with power we did not discuss in this paper. As future work, we would like to define this operation and to study the various forms of empowerment that can appear within a group.

Acknowledgements

The authors would like to thank the Rhône-Alpes Region for financing an internship of the first author at ISTC/CNR-Rome, via a EURODOC scholarship. The Ph.D. study of the first author is partially funded by the French Embassy in Romania.

References

1. Boella, G., Lesmo, L.: Norms and Cooperation: Two Sides of Social Rationality. In Hexmoor, H., and Falcone, R., (eds.): Agent Autonomy, Kluwer (2002).
2. Boella, G., van der Torre, L.: Contracts as legal institutions in organizations of autonomous agents. In Proceedings of AAMAS'04, New York (2004).

3. Boella, G., van der Torre, L.: Organizations as socially constructed agents in the agent oriented paradigm. This volume (2005).
4. Carabelea, C., Castelfranchi, C., Boissier, O.: Autonomy and social power. In Proceedings of JFSMA04, Paris (2004), 195-208.
5. Castelfranchi, C.: A micro and macro definition of power. In ProtoSociology – An International Journal of Interdisciplinary Research, 18-19 (2002), 208-268.
6. Dastani, M.M., Dignum, V., Dignum, F.: Role-Assignment in Open Agent Societies. In Proceedings of AAMAS'03, Melbourne, ACM Press (2003).
7. Dignum, F.: Autonomous Agents with Norms. In AI and Law Journal, 7 (1999), 69-79.
8. Esteva, M., Padget, J., Sierra, C.: Formalizing a language for institutions and norms. In Intelligent Agents VIII, LNAI 2333, Springer-Verlag, Berlin (2002), 348-366.
9. Hubner, J., Sichman, J.S., Boissier, O.: A model for the structural, functional, and deontic specification of organizations in multiagent systems. In Proceedings of 16th Brazilian Symposium on Artificial Intelligence, Porto de Galinhas, Brazil (2002), 118-128.
10. Jennings, N.R.: Commitments and Conventions: The Foundation of Coordination in Multi-Agent Systems. The Knowledge Engineering Review, 8 (3), (1993), 223-250.
11. Jones, A.J.J., Sergot, M.: A Formal Characterisation of Institutionalised Power. Journal of IGPL, 3, (1996), 427-443.
12. Lesser, V., *et al.*: Evolution of the GPGP/TAEMS Domain-Independent Coordination Framework. International Journal of Autonomous Agents and Multi-Agent Systems, 9(1) (2004), 87-143.
13. Lopez y Lopez, F., Luck, M.: A Model of Normative Multi-agent Systems and Dynamic Relationships. In Linderman, G., *et al.* (eds.): Regulated Agent-Based Social Systems, LNAI 2934, Springer-Verlag, Berlin (2004), 259-280.
14. Pasquier, Ph., Flores, R., Chaib-draa, B.: Modeling flexible social commitments and their enforcement. This volume (2005).
15. Sichman, J., Demazeau, Y., Conte, R., Castelfranchi, C.: A Social Reasoning Mechanism Based on Dependence Networks. In Proceedings of ECAI'94, Amsterdam, (1994).
16. Tambe, M.: Towards Flexible Teamwork. Journal of AI Research, 7 (1997), 83-124.
17. Vazquez-Salceda, J., Dignum, F.: Modelling Electronic Organizations. In Marik, V., *et al.* (eds.): MAS and applications III, LNAI 2691, Springer-Verlag, Berlin (2003), 584-593.