

## Levels of Delegation and Levels of Help for Agents with Adjustable Autonomy

Rino Falcone, Cristiano Castelfranchi

IP-CNR, Group of "Artificial Intelligence, Cognitive Modeling and Interaction"  
Viale Marx, 15 - 00137 ROMA - Italy  
E-mail: {cris, falcone}@pacs2.irmkant.rm.cnr.it

### Abstract

"Adjustable autonomy means dynamically adjusting the level of autonomy of an agent depending on the situation" [1]. Our claim is that in studying how to adjust the level of autonomy and how to arrive to a dynamic level of control, it could be useful an explicit theory of delegation able to specify different levels of autonomy.

In this paper, we present our model of levels of delegation and help. On such a basis, it is possible to analyze the adjustable autonomy of an agent both by considering the level of delegation allowed to the contractor by the client, and the possibility for the contractor itself to adjust its own autonomy by restricting or by expanding the received delegation. We consider also possible conflicts due to the initiative of the delegated agent (contractor) or to an inappropriate delegation by the client: conflicts due to the contractor's willingness to help the client better and more deeply (collaborative conflicts).

### Introduction

"Adjustable autonomy means dynamically adjusting the *level of autonomy* of an agent depending on the situation. For real-world teaming between humans and autonomous agents, the desired or optimal *level of control* may vary over time. Hence, effective autonomous agents will support adjustable autonomy." [1]

On the one side, in several situations the MA plan, the cooperation between user and agent (or between the delegating agent (*client*) and the delegated one (*contractor*)) require a strict collaboration, and a flow of the control between the partners [2], either in order to maintain the client trust [3] or in order to avoid breakdowns, failures and unsatisfactory solutions.

On the other side, mixed initiative is necessary since, in fact, it is necessary (but also dangerous) the initiative of the delegated agent. A real collaborator, not a simple tool [4], should be able to intelligently take care of client's interests. This implies the capability and the initiative either to be less autonomous than the client/user expected (in case of lack of competence, of conditions, or of knowledge), or to go beyond the user/client requests and expectations. In other words, 'the adjustable autonomy concept includes the ability.... for agents to adjust their own autonomy' [1] either by restricting or by expanding the received delegation.

This problem is particularly relevant also because delegation to agents is not only based on 'weak dependence' [5] i.e. on the possibility but not on the necessity to delegate. A client does not delegate only tasks that it could manage/perform by itself and it prefers to allocate to other agents. Software and autonomous agents will not be useful only for relieving human agents from boring and repetitive tasks; they will be mainly useful for situations where delegation and autonomy are necessary ('strong dependence', [6,7]) because the user/client will not have the local, decentralised and updated knowledge, or the expertise, or the just-in-time reactivity, or some physical skill that requires some local control-loop. Thus autonomy and initiative are not simply optional features for Agents, they are obligatory directions of study; however control cannot be completely lost and delegation cannot be complete, not only for reasons of confidence and trust, but for reasons of distribution of goals, of knowledge, of competence, and for an effective collaboration: humans needs autonomous artificial

collaborators, but artificial agents needs human collaboration and supervision.

Our claim is that in studying how to adjust the level of autonomy and how to arrive to a dynamic level of control, it could be useful an *explicit theory of delegation* (and of trust) able to specify different levels, and able to link to these the notion and the levels of autonomy. Thus, we propose in this paper our plan-based analysis of levels of delegation, of levels of help, of the notion of autonomy, and of possible conflicts due to collaborative initiative.

In this paper we will not develop - for reasons of space - this framework in a dynamic perspective. A dynamic perspective of the levels of delegation should specify:

- how the levels delegation or help can vary over time and in a reactive way (depending on the intermediate results of co-operation);
- how the levels of delegation, of trust, and of autonomy can be different for different aspects and dimensions of the same task.

We present here only the basic framework as a possible background for the theory of the Adjustable Autonomy to be developed and for this future work about its dynamics.

Why is delegation so important and central in a theory of autonomy?

If we consider an agent delegated to take care of a given *task*, it has to choose from among different possible recipes (plans), or to adapt abstract or previous plans to suit new situations; it has to find additional (local and updated) information; it has to solve a problem (not just to execute a function, an action, or implement a recipe); sometimes it has to exploit its "expertise". In all these cases this agent takes care of the interests or goals of the delegating agent "remotely" i.e. far from it and without its monitoring and intervention (control), and autonomously. This requires what we will call an "open delegation": basically the delegation "to bring it about that ...". The agent is supposed to use its knowledge, its intelligence, its ability, and to exert some degree of discretion (in this paper we do not consider as part of the agent's autonomy that the agent itself could have its *own goals* to pursue, and the consequent possible conflicts).

Moreover, given that the knowledge of the delegating agent/user (client) concerning the domain and the helping agents is limited (possibly both incomplete and incorrect) the "delegated task" (the request or the elicited behaviour) might not to be so useful for the client itself. Either the expected behaviour is useful but cannot be executed, or it is useless or self-defeating, or dangerous for the client's other goals, or else there is a better way of satisfying the client's needs; and perhaps the helping agent is able to provide greater help with its knowledge and ability, going beyond the "literally" delegated task. We will call this kind

of help: "over-help" or "critical-help". To be really helpful this kind of agent must take the initiative of opposing (not for personal reasons/goals) the other's expectations or prescriptions, either proposing or directly executing a different action/plan. To do this it must be able to recognise and reason about the goals, plans and interests of the client, and to have/generate different solutions.

However, of course, there is a trade-off between pros and cons both in open delegation and in over(critical)-help: the more intelligent and autonomous the agent (able to solve problems, to choose between alternatives, to think rationally and to plan) the less quickly and passively "obedient" it is. The probability that the solution or behaviour provided does not correspond to what we expect and delegate exactly increases.

In addition, possible conflicts arise between a "client" delegating certain tasks to an agent, and the "contractor" or in general the agent adopting and/or satisfying those tasks; conflicts which are either due to the intelligence and the initiative of the delegated agent or to an inappropriate delegation by the client: we are interested here only in conflicts due to the agent's willingness to collaborate and to help the other better and more deeply: a sort of "collaborative conflict".

## Delegation and Adoption: Analyzing the Cooperation Theory

Delegation and adoption are two basic ingredients of any collaboration and organization [8]. In fact, the huge majority of DAI and MA, CSCW and negotiation systems [9], communication protocols, cooperative software agents [10], are based on the idea that *cooperation works through the allocation of some task (or sub-task) by a given agent (individual or complex) to another agent, via some "request" (offer, proposal, announcement, etc.) meeting some "commitment" (bid, help, contract, adoption, etc.)*.

In [11] we in fact described a theory of cooperation among agents by identifying the elementary mechanisms on which any collaboration must be founded.

Our research is based on three fundamental claims:

- i) only on the basis of a principled theory of cooperation will it be possible both to really understand the human cooperation and to design cooperation among artificial agents, among humans and artificial agents, among humans through artificial agents;
- ii) this theory must be founded on the main actions of *delegation* and *adoption*;
- iii) the analysis of the delegation/adoption theory must be based on the *plan model* of the action.

We will propose a definition of delegation and adoption, the identification of their various levels, the characterization of their basic principles and representations.

The aim of this analysis is to provide some instruments for characterizing high levels of agent's cooperativeness and autonomy.

### Delegation/Adoption theory

The notion of delegation is already explicitly present in the theory of MAS, of collaboration [8], and team-work. However, we have based our analysis on much more basic notions.

Informally, *in delegation an agent A needs or likes an action of another agent B and includes it in its own plan. In other words, A is trying to achieve some of its goals through B's actions; thus A has the goal that B performs a given action.*




fig.1

A is constructing an MA plan and B has a "part" (fig.1), a share in this plan: B's task (either a state-goal or an action-goal).

On the other hand: *in adoption an agent B has a goal since and for so long as it is the goal of another agent A, that is, B usually has the goal of performing an action since this action is included in the plan of A* (fig.2). So, also in this case B plays a part in A's plan (sometimes A has no plan at all but just a need, a goal).




fig. 2

In our model, *delegation and adoption are characterized in terms of the particular set of mental states (cognitive ingredients) of the agents involved in the interaction.* In fact, a delegation (or an adoption) is a set of agent's (agents') beliefs, goals, intentions, commitments, etc.: externally there may be no interaction between the agents, the delegation (adoption) being only in the mind of one of the agents (unilateral delegation/adoption) [11]. At this

basic level delegation (adoption) can be established also between a cognitive and a non cognitive agent.

We assume that *to delegate an action necessarily implies delegating some result of that action. Conversely, to delegate a goal state always implies the delegation of at least one action (possibly unknown to A) that produces such a goal state as result.* Thus, we consider the action/goal pair  $\tau=(\alpha,g)$  as the real object of delegation, and we will call it 'task'. Then by means of  $\tau$ , we will refer to the action ( $\alpha$ ), to its resulting world state ( $g$ ), or to both.

Delegation is generally a social action [6, 7], and also a meta-action, since its object is an action. We introduce an operator of delegation with four parameters:

Delegates(A B  $\tau$ ), where A,B are agents,  $\tau=(\alpha,g)$ .

This means that A delegates the task  $\tau$  to B.

In analogy with delegation we introduce the corresponding operator for adoption: Adopts(B A  $\tau$ ). This means that B adopts the task  $\tau$  for A.

### Weak Delegation

We call *weak delegation* the delegation based on exploitation, on the passive achievement by A of the task. In it there is no agreement, no request or even influence: A is just exploiting in its plan a fully autonomous action of B. In fact, A has only to recognize the possibility that B will realize  $\tau$  by itself and that this realization will be useful for A, which "passively" awaits the realization of  $\tau$ .

More precisely,

- a) The achievement of  $\tau$  (the execution of  $\alpha$  and its result  $g$ ) is a *goal* of A.
- b) A believes that there exists another agent B that has the *power of* [12] achieving  $\tau$ .
- c) A believes that B will achieve  $\tau$  in time.
- c-bis) A believes that B *intends* to achieve  $\tau$  in time (in the case that B is a cognitive agent).
- d) A *prefers*<sup>1</sup> to achieve  $\tau$  through B.
- e) The achievement of  $\tau$  through B is the goal of A.
- f) A has the goal (*relativized* to (e)) of not achieving  $\tau$  by itself.

We consider (a, b, c, and d) what the agent A views as a "*Potential for relying on*" the agent B, its *trust*; and (e and f) what A views as the "*Decision to rely on*" B. We

<sup>1</sup> This means that, either relative to the achievement of  $\tau$  or relative to a broader goal  $g'$  that includes the achievement of  $\tau$ , A believes to be dependent on B [15].

consider "Potential for relying on" and "Decision to rely on" as two constructs temporally and logically related to each other.

### Delegation-Adoption (Contract)

We will call *strict delegation*, delegation based on explicit agreement, on the active achievement by A of the task through an agreement with B. It is based on B's adopting A's task in response to A's request/order. We will call *strict adoption*, the adoption based on explicit agreement, on the active achievement by B of the task delegated/requested by A.

In *Strict Delegation*, the delegated agent knows that the delegating agent is relying on it and accepts the task; in *Strict Adoption*, the helped agent knows about the adoption and accepts it (very often both these acceptations are preceded by a process of negotiation between the agents). In other words, *Strict Delegation requires Strict Adoption, and viceversa*: they are two facets of a unitary social relation that we call "delegation-adoption" or "contract".

There is a delegation-adoption relationship between A and B for  $\tau$ , when:

1) there is a "Potential for request of contract" from A to B:

- From A's point of view:
  - a) The achievement of  $\tau$  (the execution of  $\alpha$  and its result  $g$ ) is a goal of A.
  - b) A believes that there exists another agent B that has the power of achieving  $\tau$ .
  - c) A prefers to achieve  $\tau$  through B.

- From B's point of view:
  - d) B believes that B has the power of achieving  $\tau$ .

2) After the "Agreement":

A series of mutual beliefs (MB) are true:

(MB A B (

- a) The achievement of  $\tau$  (the execution of  $\alpha$  and its result  $g$ ) is a goal of A.
  - b) A believes that there exists another agent B that has the power of achieving  $\tau$ .
  - c) A prefers to achieve  $\tau$  through B.
  - d) B believes that B has the power of achieving  $\tau$ .
  - e) A believes that B will achieve  $\tau$  in time.
  - f) The achievement of  $\tau$  through B is the goal of A.
  - g) A has the goal (relativized to (e)) of not achieving  $\tau$  by itself.
  - h) B is socially committed to A to achieve  $\tau$  for A.
  - i) B believes that the achievement of  $\tau$  is a goal of A.
  - l) B intends to achieve  $\tau$  for A.
- ).

### Delegation based on the specification of the task

An important dimension of the delegation/adoption problem concerns how the task is specified in the

delegation action; how this specification influences the contractor's autonomy, how different interpretations of the specification of the task (or different levels of granularity in the interpretation of the task specification) for client and contractor could produce misunderstanding and conflicts.

The object of delegation ( $\tau$ ) can be minimally specified (*open delegation*), completely specified (*close delegation*) or specified at any intermediate level.

Let us consider two extreme main cases:

#### - Pure Executive (Close) Delegation

From the client's (contractor's) point of view: when the delegating (delegated) agent believes it is delegating (adopting) a completely specified task (fig. 3): what A expects from B is just the execution of an (or more) elementary action(s) (what B believes A delegated to it is simply the execution of an (or more) elementary action(s)).




fig. 3

#### - Open Delegation

There are in fact two kinds of Open Delegation. One is relative to the delegated (sub)plan, the other is relative to the opportunity to act or not.

As for the first kind let us say that the client (contractor) believes it is delegating (adopting) an incompletely specified task (fig. 4): either A (B) is delegating (adopting) a complex or abstract action, or it is delegating (adopting) just a result (state of the world). The agent B can (or must) realize the delegated (adopted) task by exerting its autonomy.




fig. 4

As for the second kind, it is the case of commands like the following one: "If it is the case, then does  $\tau$ ".

This is a special case of *conditional delegation* like "If q then does  $\tau$ ". With q fully specified, we have a form of

delegation of control, even if it is not Open since everything is fully specified<sup>2</sup>.

In the *Open-conditional* form "If it is the case, then does  $\tau$ ", B is completely free to ascertain and decide when and how it is the case to do  $\tau$ . The condition is not specified at all and is up to the agent to specify it (fig. 5).




fig. 5

The importance of *open delegation* in collaboration theory should be examined.

On the one hand, we would like to stress that *open delegation* is not only due to client's preference (utility) or limited know-how or limited skills. Of course, when A is delegating  $\tau$  to B, it is dependent on B as for  $\tau$  [7]: it needs B's action for some of its goals (either some domain goals or goals like saving time, effort, resources, etc.). However, *open delegation* is fundamental because it is also due to A's ignorance about the world and its dynamics. In fact, frequently enough *it is not possible or convenient to fully specify*  $\tau$  because some local and updated knowledge is needed in order for that part of the plan to be successfully executed.

*Open delegation* is one of the bases of the *flexibility* of distributed and MA plans. To be radical, delegating actions to an autonomous agent always requires some level of "openness": the agent at least cannot *avoid monitoring and adapting* its own actions, during their execution.

Moreover, the distributed character of the MA plans derives from *open delegation*.

As we saw, A can delegate to B either an entire plan or some part of it (*partial delegation*). The combination of *partial delegation* (where the contractor can ignore the other parts of the plan) and *open delegation* (where the client can ignore the sub-plan chosen and developed by the contractor) creates the possibility that A and B collaborate

<sup>2</sup> In fact, notice that also in this case B has some discretion about performing or not its task, since it is delegated to autonomously test the condition for its action. Consider also that performing the test might be a complex action, requiring an autonomous plan and discretion about it. In this case also this is *Open delegation*

in a plan that they do not share and that *nobody* knows fully: that is to say a truly *distributed plan* [4, 6]. However, for each part of the plan there will be at least one agent that knows it.

The object of the delegation can be a practical or *domain action* as well as a *meta-action* (searching, planning, choosing, problem solving, and so on).

When A is open delegating to B some domain action, it is necessarily also delegating to B some meta-action: at least searching for a plan, applying it, and sometimes deciding between alternative solutions. We call B's *discretion* concerning  $\tau$  the fact that some decision about  $\tau$  is delegated to B.

### Delegation of the Control

The control (or check up) is an action aimed at ascertaining whether another action has been successfully executed (or if a given state of the world has been realized or maintained).

*Controlling an action means verifying that its relevant results hold* (including the execution of the action itself).

Plans typically contain control actions of some of their actions. When the client is delegating a given object-action, what about its control actions?

Considering, for the sake of simplicity, that the control action is executed by a single agent, when  $\text{Delegates}(A \ B \ \tau)$  there are at least four possibilities:

- i) A delegates the control to B: the client does not (directly) verify the success of the delegated action to the contractor;
- ii) A delegates the control to a third agent;
- iii) A gives up the control: nobody is delegated to control the success of  $\alpha$ ;
- iv) A maintains the control for itself.

Each of these possibilities could be explicit or implicit in the delegation of the object-action, in the roles of the agents (if they are part of a social structure), in the preceding interactions between the client and contractor, etc.

### Delegation of Initiative

The notion of initiative is often ambiguous and too broad. Agents should be characterised at least by two forms of initiative:

- *Motu proprio*: Agents do not execute the required action/plan under the direct and immediate command of their client or user; they take care of and also "decide" the appropriate moment for the execution of the task (which can be ignored by the client or user). Their behaviour is not fired or elicited by the user or the request, but by the agent's autonomous relation with its environment. It takes

the initiative of executing the task where and when appropriate, also depending on its internal state (this can be considered also as another aspect of “discretion”).

- *Spontaneous interaction or pro-active help*: Agents can act “for us” (or for the client) also without any request or beyond the request. The agent might spontaneously help the other or over-help it (doing more or better than requested) or spontaneously starting an interaction for example by offering some help. They may anticipate the user/client request and even desires.

**Levels of Adoption**

In order for the adoption be an effective help (*deep cooperation*), the contractor should consider/foresee the client's plan (in which the delegated task is inserted), its goals and interests and, on the basis of the circumstances, deeply-understand/improve/preserve the requested help. In this way it is possible to classify the contractor's adoption at the various levels:

Literal help

*The contractor adopts exactly what has been delegated by the client (see fig. 6).*




fig. 6

Overhelp

*The contractor goes beyond what has been delegated by the client without changing the client's plan (see fig. 7).*




fig. 7

Critical help

*The contractor achieves the relevant results of the requested plan/action, but modifies the plan/action (see fig. 8).*




fig. 8

Critical overhelp

*The contractor implements an overhelp and in addition modifies/changes the plan/action (see fig. 9).*




fig. 9

Hyper-critical help

*The contractor adopts goals or interests of the client that the client itself did not take into account: by doing so, the contractor neither performs the delegated action/plan nor totally achieves the results that were delegated (see fig. 10).*




fig. 10

**Conflicts Due to the Level of Adoption of the Contractor**

Given our characterisation of delegation and adoption, we can derive a series of conflicts arising between the two agents when there is a mismatch between the intended delegation and the intended adoption. These mismatches are neither due to simple misunderstandings of A's request/expectation nor to B's offer, nor to a wrong or incomplete plan/intention recognition of B. We are mainly

interested in *collaborative conflicts* which come from B's intention to help A beyond its request or delegation and to exploit its own knowledge and intelligence (reasoning, problem solving, planning, and decision-making skills) for A [13].

- *Conflicts due to the contractor's over-help, critical help, critical over-help, hyper-critical help*

In any case of over, critical and hypercritical adoption there is apparently a conflict, since A has the goal that B does  $\alpha$ , while B is doing or intends to do something different for A. Normally these conflicts can be quickly solved, for two reasons. First, B's intention is to help A, it's a collaborative intention; second, normally B is "entitled"<sup>3</sup> by A (either explicitly or implicitly) to provide this deeper help, and A is expecting this initiative and autonomy. Thus, normally there is no real conflict since A is *ready to accept* B's collaborative initiative. However, sometimes these cases trigger serious conflicts which have to be negotiated. This is specially true in organizations and among different roles.

Leaving aside possible cases of misunderstanding between client and contractor about A's request/expectation or B's offer (or to a wrong plan/intention by B), we can distinguish the reasons for conflict (i.e., A is against B's initiative) into two main classes:

i) *Trouble for A's goals*

B can jeopardize the goal achievement of A: this is possible for at least two reasons:

i1) *Lack of coordination*

A plan is composed of many actions (assigned to several agents, when there is a partial delegation), so a unilateral initiative on the part of one agent to change that plan without reconsidering the general plan might be fatal (because of interference) or lead to a waste of resources, time, etc. (because of redundancy).

i2) *Disagreement about action results*

A knows or believes that the action executed by B does not bring about the results expected or believed by B itself.

ii) *Role and Status*

In this case the conflict is relative to the entitlement of B by A to take the initiative of changing the delegated task. For reasons of power, job, subordination, role B while doing such a sub/over/critical help is going beyond what it is permitted to do (according to A).

This important aspect concerning conflicts extends beyond the plan-based analysis of delegation we are illustrating here.

---

<sup>3</sup> We will say that B is *entitled* by A to  $\tau$  through the delegation *Delegates(A B  $\tau$ )*, when there is common (to A and B) knowledge that A is committed not to oppose, not to be astonished, etc., if B pursues  $\tau$  [15].

## Levels of Autonomy

On the basis of the previous analysis of the delegation and adoption dimensions, it is possible to identify various levels and types of agency.

*Types and levels of delegation characterize the autonomy of the delegated agent.*

There are at least two meanings of "autonomy": one is equal to *self-sufficiency*, not being dependent on others for our own goals [7]; on this side the less dependent B is on A regarding the resources necessary for the task, the more autonomous B is of A regarding that task.

The other meaning is related to action and goals and to their levels. One could distinguish between *performance or executive autonomy* (the agent is not allowed to decide anything but the execution of the entire delegated plan [6]: in our terms, given an agent A and a plan  $\alpha$ ,  $\alpha$  is completely specified in the delegation itself); *planning autonomy* (the agent is allowed to plan by itself, to choose its own plan to obtain the goal delegated, for example  $\alpha$  is not completely specified in the delegation itself); *goal autonomy* (the agent is allowed to have/find goals). Here we ignore the autonomous goals of the delegated agent, so we can characterise different degrees of autonomy in delegation as follows.

The autonomy of the contractor vis-à-vis the client increases along various dimensions:

- the more *open the delegation* (the less specified the task is), or
- the more *control actions* given up or delegated to B by A, or
- the more delegated decisions (*discretion*),

the more autonomous B is of A regarding that task.

## Concluding Remarks

As we said, we presented here only our basic framework as a possible background for the theory of the Adjustable Autonomy to be developed and for the future work about its dynamics.

Let us in particular stress how in this future development a merging between the theory of trust and the theory of delegation and help levels will be necessary.

To delegate a task or a role, in fact, a given *degree of trust* is necessary, based on some evaluation and agent modelling [14]. However, trust is not only the global final result of the evaluation, that determines if one delegates or not that task to that agent. It should be articulated along several dimensions. Our claim is that it is precisely an articulated trust which determines the level and kind of delegation. For example, -as we said- open delegation presupposes that the client has a model of the agent 'intellectual' competence and that specifically trusts its

ability to find or build a plan, to choose among different possible plans, to intend a given plan and to persist in its execution, to assess or provide the necessary resources and conditions, and to successfully control its own performance. If the client does not delegate all these things, and takes back to himself either the control, or the planning, or the discretion, etc. it is precisely because he trusts some capabilities of the agents but not the others. He has *different degree of trust relative to different aspects* of the complex task, and he delegates more or less on one aspect or on the other depending on this *differential reliance*.

To model this process -that we believe to be the basis of Adjustable Autonomy- we claim that it is necessary to merge the theory of levels of delegation [11] with the theory of the degrees of trust [3].

### References

- [1] Call for Papers, 1999 *AAAI Spring Symposium on Agents With Adjustable Autonomy*, March 22-24, 1999, Stanford University.
- [2] Cesta, A., D'Aloisi, D. Mixed-Initiative Issues in an Agent-Based Meeting Scheduler. *User Modeling and User-Adapted Interaction*, 1998. To appear.
- [3] Castelfranchi C., Falcone R., Principles of trust for MAS: cognitive anatomy, social importance, and quantification, *Proceedings of the International Conference of Multi-Agent Systems (ICMAS'98)*, pp. 72-79, Paris, July, 1998.
- [4] B. Grosz, Collaborative Systems. *AI Magazine* (summer 1996) 67-85.
- [5] Jennings. N.R. 1993. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 3, 223-50.
- [6] R. Conte & C. Castelfranchi, *Cognitive and Social Action* (UCL Press, London, 1995).
- [7] Sichman, J, R. Conte, C. Castelfranchi, Y. Demazeau. A social reasoning mechanism based on dependence networks. In *Proceedings of the 11th ECAI*, 1994.
- [8] A. Haddadi, *Communication and Cooperation in Agent Systems* (the Springer Press, 1996).
- [9] Rosenschein, J.S. and Zlotkin, G. *Rules of encounters Designing conventions for automated negotiation among computers*. Cambridge, MA: MIT Press. 1994.
- [10] Werner, E., *Cooperating agents: A unified theory of communication and social structure*. In L.Gasser and M.N.Huhns, editors, *Distributed Artificial Intelligence: Volume II*. Morgan Kaufmann Publishers, 1990.
- [11] Castelfranchi, C., Falcone, R. (1998), *Towards a Theory of Delegation for Agent-based Systems, Robotics and Autonomous Systems*, Special issue on Multi-Agent Rationality, Elsevier Editor. Vol. 24, pp. 141-157.
- [12] C. Castelfranchi, *Social Power: a missed point in DAI, MA and HCI*. In *Decentralized AI*. Y. Demazeau & J.P.Mueller (eds) (Elsevier, Amsterdam 1991) 49-62.
- [13] J. Chu-Carroll, S. Carberry, Conflict detection and resolution in collaborative planning, *IJCAI-95 Workshop on Agent Theories, Architectures, and Languages*, Montreal, Canada (1995).
- [14] Falcone R. & Castelfranchi C., Levels of help, levels of delegation and agent modeling. *AAAI-96 Agent Modeling Workshop*, 4 August 1996.
- [15] C. Castelfranchi, *Commitment: from intentions to groups and organizations*. In *Proceedings of ICMAS'96*, S.Francisco, (AAAI-MIT Press 1996).