

On Achievable Goals and Feasible Plans in Open Multi-Agent Systems

Jaime Simão Sichman
University of São Paulo
Computer Engineering Department
Av. Prof. Luciano Gualberto, 158, travessa 3
05508-900, São Paulo, SP, Brazil
jaime@pcs.usp.br

Abstract

In this paper, we present a subjective representation of the notions of *feasible plans* and *achievable goals*, in order to model the decision mechanism of autonomous agents, immersed in an open multi-agent context. By open, we mean that agents may enter or leave the agency at any moment, without a centralized control. We believe that an agent who uses this model can better adapt himself to the changing conditions of the system, specifically to the fact that services may dynamically become available/not available.

1 Introduction

In [6, 11, 5], the future information processing environments are presented as being composed of huge heterogeneous networks of processing resources. These resources, autonomous and distributed, may consist of computers, huge applications and huge databases. Authors call these environments “societies of objects” or “electronic organizations”. Simply referring to such pro-

cessing resources as agents¹, a system composed of these agents will have the following characteristics: *decentralized design*, *openness*, and *local autonomy*. We call this kind of system an *open multi-agent system* (open MAS). In an open MAS, as services may become available or may disappear in running time, without a centralized control, agents have to cope with these changes by *representing and exploiting internally some properties of the other agents*. They will have to *adapt themselves* to these changes in the environment. By this expression, we mean specifically that as services may dynamically become available/not available, agents must therefore reason about each other, for instance in order to choose different partners with whom to work cooperatively. This choice depends on the available services at the agency level.

In this paper, we address one particular point of such *adaptation* procedure. We present both a model and an implementation of an agent's decision mechanism, specifically regarding the choice of a goal to be pursued and a plan that achieves this goal. This model is based on the notions of *feasible plans* and *achievable goals*, which are explicitly represented within the agents, as described next.

2 Social Reasoning Mechanism

We believe that an agent must have a *social reasoning mechanism* to cope properly with the scenario described above. We call social a mechanism that uses information about the others in order to infer some properties.

In [8, 9], a mechanism of this kind is presented, based on the notion of *social dependence* [3]. Briefly an agent is said to be *dependent* on another if this latter may facilitate/prevent him to achieve one of his goals. Agents use the information they have about each other to detect complementary expertise, and to form dynamically *coalitions* to achieve their goals. This information is stored in a *private* data structure called *external description*, composed of the agents' set of goals, plans, actions and resources. In [9], we have shown how an agent uses his external description to construct a *dependence network*, a single structure which contains all the agent's dependences. In particular, the leaves of a network represent the agents that are able to per-

¹This definition of agent is obviously vague and poor, and it is being used exclusively to enable us to explain some essential features of this work. A complete formal definition may be found in [8].

form some action/to control some resource, which the reasoning agent can not perform/control.

2.1 Agent Model

A social reasoning mechanism is an essential building block for really autonomous agents. In [8], we have proposed an agent model, inspired by the ASIC model [1], which exploits several aspects of this mechanism, like for instance dependence based coalition formation. A detailed description of this model is out of scope here, but we need to present some of its essential aspects, in order to enable the comprehension of the next sections:

- agents do not perform online planning. We consider that agents have pre-established plans² that achieve several goals, whether the agents do have these goals or not. The way these plans were generated are out of scope of this work. One may suppose, for instance, a case-based reasoning inference style. The important aspect to stress is that we do not intend to propose new planning techniques for a multi-agent scenario: we are interested to propose a model to describe how an agent chooses a pre-defined plan (and a goal to pursue) given that some services are not available all the time in an open MAS;
- agents first choose a goal to pursue, based on a worth function, and then a plan to achieve it. Eventually, if an agent can not execute the chosen plan alone, he will also choose the more susceptible partners to whom he should send a coalition proposal. This partner choice is based on the notion of dependence situation, described in [8, 10]. In this way, we do not use a strict game-theoretic approach³, since we consider that this approach is insufficient to characterize open MAS, as shown in the next sections. In particular, we believe that an explicit and subjective representation of achievable goals/feasible plans is essential to construct real autonomous agents.

²A more formal definition of plan may be found in [8]. For the purposes of this paper, a plan is considered as being composed of several actions, each of those possibly using several resources.

³We mean that the worth of goals and costs of plans are not compared.

2.2 Preliminary Definitions

In our formal model, the variables i, j denote agents, a denote actions and r denote resources.

We will also use the predicates $is_g(i, g), is_a(i, a), is_r(i, r)$ and $is_p(i, p)$ to denote respectively the fact that the reasoning agent believes that agent i has goal g , can perform action a , can control resource r and has a plan p (to achieve some goal)⁴. In other words, these predicates are used to represent formally each element of the reasoning agent's external description⁵.

Some useful properties inferred by the social reasoning mechanism, which are specifically used for choosing a goal to be pursued and a plan to achieve it, are described next.

2.3 Available Actions and Resources

If we place ourselves as an external observer of an open MAS, we can notice that one can never assure that every action/resource needed in a plan (which achieves a certain goal) is always *available*, i.e., that each of the actions/resources needed to accomplish the plan can be currently executed/is currently controlled by at least one agent in the agency.

We can formally express these notions of available actions and resources as:

$$available_a(a) \Leftrightarrow \exists i is_a(i, a) \quad (1)$$

$$available_r(r) \Leftrightarrow \exists i is_r(i, r) \quad (2)$$

2.4 Feasible Plans

Using the definitions (1) and (2), we can express three different notions of feasible plans, similarly to the three notions of autonomy and dependence presented in [8]. Intuitively, a plan is *a-feasible* (action-feasible) if every action needed to perform it is available. Similarly, a plan is said to be *r-feasible* (resource-feasible) if every resource needed to execute it is available.

⁴The subscripts which appear in these predicates have nothing to do with its parameters: for instance, the symbol is_p is just a shortcut for is_plan .

⁵We would like to stress the fact that the external description is a *private* structure.

A plan which is both a-feasible and r-feasible is said to be *s-feasible* (social-feasible):

$$feasible_a(p) \Leftrightarrow \forall a (uses_a(p, a) \Rightarrow available_a(a)) \quad (3)$$

$$feasible_r(p) \Leftrightarrow \forall r (uses_r(p, r) \Rightarrow available_r(r)) \quad (4)$$

$$feasible_s(p) \Leftrightarrow feasible_a(p) \wedge feasible_r(p) \quad (5)$$

where the predicates $uses_a(p, a)$ and $uses_r(p, r)$ represent respectively that the action a /the resource r is needed to execute the plan p .

2.5 Achievable Goals

Similarly, one can never assure that a given agent's goal is always *achievable* in an open MAS. A goal is said to be achievable by a certain plan if this latter is feasible. Using the same framework of the last section, we can define three different notions for achievable goals: *a-achievable*, *r-achievable* and *s-achievable*:

$$achievable_a(g, p) \Leftrightarrow achieves_g(p, g) \wedge feasible_a(p) \quad (6)$$

$$achievable_r(g, p) \Leftrightarrow achieves_g(p, g) \wedge feasible_r(p) \quad (7)$$

$$achievable_s(g, p) \Leftrightarrow achievable_a(g, p) \wedge achievable_r(g, p) \quad (8)$$

where the predicate $achieves_g(p, g)$ means that the successful execution of plan p achieves goal g .

In the next sections, we show how an agent, immersed in an open MAS, can take into account these notions of *feasible plans* and *achievable goals* in order to choose a goal to be pursued and a plan to achieve it. We will restrict our analysis to available actions, as the extension to the resource case is rather simple.

3 Choice of Goals

In the general case, an agent may not necessarily be able to perform all the actions needed in a plan which achieves one of his goals. Therefore, when choosing a goal to pursue in a given moment, he needs to verify if the goal is

achievable or not:

$$\begin{aligned} \text{achievable}(i, g) \Leftrightarrow & \text{AUT}(i, g) \vee (\text{DEP}(i, g) \wedge \\ & \exists p (is_p(i, p) \wedge \text{achievable}_a(g, p)) \end{aligned} \quad (9)$$

An agent i ⁶ infers that a goal g is achievable if (i) he is autonomous for this goal — represented by the predicate $\text{AUT}(i, g)$ — or (ii) he is dependent on the others for this goal — represented by the predicate $\text{DEP}(i, g)$ — but he has a plan where all the actions can be performed by at least one agent in the agency⁷. It is shown in [8] that $\text{AUT}(i, g) \Rightarrow \text{achievable}(i, g)$, since the agent is able to perform all the actions needed in the plan by himself. We have decided, however, to explicitly differentiate the cases where a goal is achievable for an agent because he can perform all the needed actions alone from the case where there he needs to form a coalition with the others to achieve this goal.

Let us take as an example a scenario from a research laboratory. Let us suppose that agent ag_1 is a researcher that has 3 different goals to achieve: write an article about MAS (*write_mas_paper*), write another article about MAS and Social Simulation (*write_ss_mas_paper*), and review a paper about Object-Oriented Programming (*review_oop_paper*). Let us suppose that his plans for these goals are the following:

$$\begin{aligned} \pi_1 : \text{write_mas_paper}() & := \text{write_mas_section}(), \text{process_latex}(). \\ \pi_2 : \text{write_mas_paper}() & := \text{write_mas_section}(), \text{process_word}(). \\ \pi_3 : \text{write_ss_mas_paper}() & := \text{write_ss_section}(), \text{write_mas_section}(), \\ & \text{process_word}(). \\ \pi_4 : \text{review_oop_paper} & := \text{analyse_oop_paper}(), \text{write_report}(), \\ & \text{process_latex}(). \end{aligned}$$

Let us also consider that ag_1 knows \LaTeX , but is not very acquainted with Microsoft Word. He knows very well the MAS and the OOP domains, but is not exactly a specialist in the Social Simulation field.

⁶From this point on, we consider that the reasoning agent is i .

⁷A complete formal definition of the predicates AUT and DEP may be found in [8].

If the action *write_ss_section* can not be performed by any agent in his research laboratory, the social reasoning mechanism of *ag₁* would infer the following propositions: *achievable(ag₁, write_mas_paper)*, \neg *achievable(ag₁, write_ss_mas_paper)* and *achievable(ag₁, review_oop_paper)*.

Let $w(g)$ be a function which represents, for each agent, the worth of a goal g ⁸. Using a pure utilitarian approach, one can suppose that given two goals g and g' where $w(g) > w(g')$, an agent will choose to pursue g , since this goal has a greater worth value. We believe that a decision criterion based exclusively on this notion of worth is *insufficient* to characterize open societies. It does not seem reasonable that an agent chooses to pursue a goal with greatest worth value without analyzing if this goal is achievable or not. In fact, in a pure utilitarian approach, all the goals of an agent are considered to be achievable, as in [7]. The decision mechanism is basically activated in order to minimize harmful interference. Moreover, the very notion of goal is not usually represented explicitly within the agents. In our example, even if we suppose that $w(\textit{write_ss_mas_paper}) > w(\textit{write_mas_paper})$, agent *ag₁* will not choose to pursue goal *write_ss_mas_paper*, because his social reasoning mechanism has detected that this goal is not currently achievable.

3.1 Goal Decision Criterion

Let $M^G = (i, G, \textit{achievable}_{goal}, w_{goal}, \prec_{goal}, \textit{decision}_{goal})$ be a formal model which characterizes an agent's goal decision mechanism where:

- i denotes the agent to whom this model belongs;
- G denotes his set of goals;
- $\textit{achievable}_{goal} : G \mapsto B$ is a function which returns for each goal $g \in G$ a boolean value expressing if the goal is achievable or not;
- $w_{goal} : G \mapsto N$ is a function which returns for each goal $g \in G$ an integer value expressing its worth value;

⁸It is out of scope of this paper to propose such a worth function, we simply assume that there is one. The same argument holds for the costs of actions and resources, which are introduced in the next section.

- \prec_{goal} is a partial order relation defined as follows : If $g, g' \in G$ then $g' \prec_{goal} g$ iff: $achievable_{goal}(g') = false$ and $achievable_{goal}(g) = true$ or $achievable_{goal}(g') = achievable_{goal}(g) = true$ and $w_{goal}(g') < w_{goal}(g)$. As the relation \prec_{goal} is both transitive and anti-symmetrical, (G, \prec_{goal}) is a partial ordered set;
- $decision_{goal} : 2^G \mapsto G$ is a function which from the goal set G chooses a goal g to be pursued, as defined next.

Let us call w_{max} the highest worth value of all the agent's achievable goals. One may notice that this value does not necessarily correspond to the most important goal, because this latter may be non-achievable. Let G_{max} be the set of goals whose worth value is w_{max} and let $G_{max}^c = G \cap \overline{G_{max}}$. Clearly, we have:

$$\forall g \in G_{max} \quad \nexists g' \in G_{max}^c \quad g \prec_{goal} g' \wedge achievable_{goal}(g') = true$$

Using these definitions, we can define the goal decision function $decision_{goal}$ in the following way:

$$decision_{goal}(G) = \begin{cases} none & \text{if } G_{max} = \emptyset \\ random(G_{max}) & \text{in any other case} \end{cases}$$

As we do not impose that different goals must have different worth values, the choice of one goal among G_{max} is made randomly by the function $random(G_{max})$. In the case where G_{max} is an empty set, the agent must wait until at least one of his goals become achievable, for instance by the arrival of a new member in the agency.

4 Choice of Plans

Once an agent has chosen a goal to be pursued, he must select a plan to achieve this goal. In the general case, an agent may have several plans to achieve some goal. Using the same arguments presented in the earlier section, before choosing a plan, an agent must ensure that it is feasible:

$$feasible(i, p) \Leftrightarrow is_p(i, p) \wedge feasible_a(p) \quad (10)$$

where the predicate $is_p(i, p)$, as stated before, denotes the fact that the agent i has generated the plan p ⁹.

Let us take the same example of the previous section, and suppose that agent ag_1 has chosen to pursue goal *write_mas_paper*. He has two alternative plans to achieve this goal, namely π_1 and π_2 . Let us also consider that there is no expert in Microsoft Word in the laboratory and that ag_1 knows this fact. The social reasoning mechanism of agent ag_1 infers the following propositions: $feasible(ag_1, \pi_1)$ and $\neg feasible(ag_1, \pi_2)$.

Let $c(a)$ and $c(r)$ be two functions which represent respectively the cost of an action/resource, and $n(a)$ and $n(r)$ represent the number of times a given action/resource appears in a plan p . As a first approach, we can say that the cost of a plan is the sum of the cost of all actions and resources needed to accomplish it¹⁰:

$$c(p) \triangleq \sum_{a_k \in A_p} n(a_k) * c(a_k) + \sum_{r_k \in R_p} n(r_k) * c(r_k) \quad (11)$$

where $A_p = \{a_k \mid uses_a(p, a_k)\}$, $R_p = \{r_k \mid uses_r(p, r_k)\}$. If we adopt a pure utilitarian approach, we can suppose that given two plans p and p' such as $c(p) < c(p')$, both of them achieving the same goal, an agent will choose to execute plan p , because this plan is less costly than p' . Like in the previous section, we believe that a decision criterion based exclusively on this notion of cost is also *insufficient* to model open societies. As in the goal case, it does not seem reasonable that an agent chooses a less expensive plan which is not feasible. In our example, even if we suppose that $c(\pi_1) > c(\pi_2)$, agent ag_1 will choose to execute plan π_1 , because his social reasoning mechanism has detected that the other plan is not currently feasible.

4.1 Plan Decision Criterion

Let $M^P = (i, g, P, feasible_{plan}, c_{plan}, \prec_{plan}, decision_{plan})$ be a formal model which characterizes an agent's plan decision mechanism where:

⁹In some cases, as shown in [8, 8], an agent may use the plans he believes his possible partners have in order to infer some useful properties, like his dependence situations.

¹⁰A more rigorous description would take into account the partner who may execute the action/release the resource, since different partners may attribute different costs to perform actions/release resources.

- i denotes the agent to whom this model belongs;
- g denotes the goal being pursued;
- P represents the agent's set of plans which achieve g ;
- $feasible_{plan} : P \mapsto B$ is a function which returns for each plan $p \in P$ a boolean value expressing if the plan is feasible or not;
- $c_{plan} : P \mapsto N$ is a function which returns for each plan $p \in P$ an integer value expressing its cost;
- \prec_{plan} is a partial order relation defined as follows : If $p, p' \in P$ then $p' \prec_{plan} p$ iff: $feasible_{plan}(p') = false$ and $feasible_{plan}(p) = true$ or $feasible_{plan}(p') = feasible_{plan}(p) = true$ and $c_{plan}(p') < c_{plan}(p)$. As the relation \prec_{plan} is both transitive and anti-symmetrical, (P, \prec_{plan}) is a partial ordered set;
- $decision_{plan} : 2^P \mapsto P$ is a function which from the plan set P chooses a plan p to be executed, as defined next.

Let us call c_{min} the smallest cost value of all the agent's feasible plans achieving g . Once more, one can notice that this value does not necessarily correspond to the least expensive plan, because this latter may be non-feasible. Let P_{min} be the set of plans whose cost is c_{min} and let $P_{min}^c = P \cap \overline{P_{min}}$. Clearly, we have:

$$\forall p \in P_{min} \quad \exists p' \in P_{min}^c \quad p' \prec_{plan} p \wedge feasible_{plan}(p') = true$$

Using these definitions, we can define the plan decision function $decision_{plan}$ in the following way:

$$decision_{plan}(P) = random(P_{min})$$

As we do not impose that different plans must have different costs, the choice of one plan among P_{min} is made randomly by the function $random(P_{min})$. According to our definition of achievable goal (9), we can assure that given an achievable goal g there is at least one feasible plan p to achieve it.

5 Experimental Results

The goal and plan decision mechanisms described in the previous sections were implemented within the agent model of the DEPINT system [8], conceived as a software tool which illustrates several uses of the social reasoning mechanism, like adaptative choices of goals, plans and partners, dependence-based coalition formation and autonomous belief revision.

The system was build as an application layer of the MAGMA platform, an environment for the development of multi-agent systems, which is being currently developed at LEIBNIZ Institute, Grenoble [4, 2].

We have simulated our research laboratory scenario using the DEPINT system. The results concerning agent ag_1 's goal choice are the following:

```
===== Reasoning about goals ...
```

```
My dependence network is:
```

```
  ag1
  <ag1>
  ----- write_mas_paper (20)
  |----- write_mas_paper:=write_mas_section(), process_latex().
  |         |----- A-AUTONOMOUS
  |         |         |-----
  |         |         |
  |         |         | write_mas_paper:=write_mas_section(), process_word().
  |         |         |----- process_word
  |         |         |----- UNKNOWN
  |         |         |-----
  |         |
  |         | review_oop_paper (10)
  |         |----- review_oop_paper:=analyse_oop_paper(), write_report(),
  |         |         | process_latex().
  |         |         |----- A-AUTONOMOUS
  |         |         |-----
  |         |
  |         | write_ss_mas_paper (30)
  |         |----- write_ss_mas_paper:=write_ss_section(), write_mas_section(),
  |         |         | process_word().
  |         |         |----- write_ss_section
  |         |         |***** UNKNOWN
  |         |         |         |-----
  |         |         |         | process_word
  |         |         |***** UNKNOWN
  |         |         |-----
```

My current list of possible goals is :

```
write_mas_paper(20)    achievable
review_oop_paper(10)  achievable
write_ss_mas_paper(30) non achievable
```

==== Deciding about goals ...

The goal selected is : write_mas_paper (20)

In order to select a goal, ag_1 uses his dependence network, which was constructed based on the information stored in his external description. Within this network, non-available actions are represented as possible partners — leaves — with the label UNKNOWN. As we would expect, agent ag_1 chooses to pursue goal *write_mas_paper*, which is the goal with biggest worth value among those which are achievable.

Once selected the goal *write_mas_paper*, agent ag_1 decides which plan to execute. The simulation results are the following:

==== Reasoning about plans ...

My dependence network is:

```
ag1
<ag1>
----- write_mas_paper (20)
  |----- write_mas_paper:=write_mas_section(), process_latex().
    |----- A-AUTONOMOUS
    |          |-----
    |          |
    |          | write_mas_paper:=write_mas_section(), process_word().
    |----- process_word
              |----- UNKNOWN
              |-----
```

My current list of possible plans is:

```
write_mas_paper:=write_mas_section(), process_latex().(30) feasible

write_mas_paper:=write_mas_section(), process_word().(20) non-feasible
```

==== Deciding about plans ...

```
The plan selected is :  
write_mas_paper:=write_mas_section(), process_latex(). (30)
```

As we have expected, the first plan π_1 is chosen, because the other plan π_2 is not currently feasible. In the example, we have supposed that the actions *write_mas_section* and *process_word* has a cost of 10, and the action *process_latex* a cost of 20.

In our scenario, agent ag_1 can execute all the actions needed in the selected plan alone. This situation is represented by the leaf with label **A_AUTONOMOUS**. If this was not the case, ag_1 would have to decide to which possible partner a proposal of coalition should be sent. A decision criteria based on dependence situations, implemented within the agents of the DEPINT system, may be found in [8, 10].

Let us suppose that later on, an agent which is able to perform the actions *process_word* and *write_ss_section* arrives. Within the DEPINT system, whenever an agent enters the agency, an interaction protocol is developed, in order to enable the agents to update their information about the members of the agency¹¹. If agent ag_1 were to choose a goal to achieve after the arrival of this new agent, he would choose then to pursue the goal *write_ss_mas_paper*, since this goal has become achievable. Therefore, our model is well-suited to be used in an open MAS context.

6 Conclusion and Further Work

In this paper, we have presented a model and an implementation of an agent's decision mechanism, specifically concerning the choice of a goal to be pursued and a plan to achieve it. This mechanism, designed to be used in an open MAS context, is based on the notions of *feasible plans* and *achievable goals*. With this mechanism, one can better model the decision procedure of really autonomous agents, immersed in an open multi-agent system, which have to adapt themselves to changes in these environments, specifically due to the arrival/exit of agency members.

To our knowledge, this work is the first attempt to explicitly represent in a subjective way the notions of feasible plans and achievable goals within cognitive agents.

¹¹A detailed description of this protocol may be found in [8].

One may argue that using a decision theoretic approach, we can represent unavailable actions/resources by ascribing them a very high cost. It seems obvious that such a model would not allow an agent to differentiate between an action/resource that is effectively very expensive from one that is not currently available in the agency. From a pure utilitarian approach, this lost of expressive power is not so critical, because the final result is the same: agents will not chose these plans. However, if we take instead a more mentalistic approach, this kind of information is essential, and must be explicitly represented within the agents.

Another important aspect is that the additional algorithmic cost of using this model is not important, since the social reasoning mechanism already ensures that agents maintain their dependence networks updated.

As a further work, we intend to verify if the activation order of the several decisions steps (goal, plan and partner) has a significant effect on the agent behavior. As an example, an agent may want to choose a goal whose worth value is smaller than a second one¹² if he believes that he could find a more cooperative partner for the second one than for the first one.

7 Acknowledgments

We would like to thank previous referees as well as Anna Rillo (EPUSP) and Yves Demazeau (LEIBNIZ/IMAG) for their comments which helped improve this paper. The work described in this paper was partially developed during the author's PhD program at LIFIA/IMAG, Grenoble, France, when he was supported by FAPESP (Brazil), grant number 91/1943-5. Currently, the author is partially supported by CNPq (Brazil), grant number 301041/95-4.

References

- [1] Olivier Boissier and Yves Demazeau. ASIC: An architecture for social and individual control and its application to computer vision. In Yves Demazeau, Jean-Pierre Müller, and John Perram, editors, *Pre-proceedings of the 6th MAAMAW*, pages 107–118, Odense, Denmark, August 1994.

¹²Supposing that both of them are achievable.

- [2] Eleri Cardozo, Jaime Simão Sichman, and Yves Demazeau. Using the active object model to implement multi-agent systems. In *Proceedings of the 5th IEEE International Conference on Tools with Artificial Intelligence*, pages 70–77, Boston, USA, November 1993.
- [3] Cristiano Castelfranchi, Maria Micelli, and Amedeo Cesta. Dependence relations among autonomous agents. In Eric Werner and Yves Demazeau, editors, *Decentralized A. I. 3*, pages 215–227. Elsevier Science Publishers B. V., Amsterdam, NL, 1992.
- [4] Yves Demazeau. From interactions to collective behaviour in agent-based systems. In *Pre-proceedings of the invited lectures of the 1st European Conference on Cognitive Science*, St. Malo, France, March 1995.
- [5] Carl E. Hewitt. Some requirements for mobile distributed telecomputing architecture. In Cristiano Castelfranchi and Eric Werner, editors, *Artificial Social Systems*, volume 830 of *LNAI series*, pages 259–270. Springer-Verlag, Berlin, DE, 1993.
- [6] John R. Nicol, Thomas Wilkes, and Frank A. Manola. Object orientation in heterogeneous distributed computing systems. *IEEE Computer*, pages 57–67, June 1993.
- [7] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge, MA, 1994.
- [8] Jaime Simão Sichman. *Du Raisonnement Social Chez les Agents: Une Approche Fondée sur la Théorie de la Dépendance*. Thèse de Doctorat, Institut National Polytechnique de Grenoble, Grenoble, France, 1995.
- [9] Jaime Simão Sichman, Rosaria Conte, Yves Demazeau, and Cristiano Castelfranchi. A social reasoning mechanism based on dependence networks. In Tony Cohn, editor, *Proceedings of the 11th ECAI*, pages 188–192, Amsterdam, The Netherlands, August 1994. John Wiley & Sons Ltd.
- [10] Jaime Simão Sichman and Yves Demazeau. Exploiting social reasoning to enhance adaptation in open multi-agent systems. In Jacques Wainer and Ariadne Carvalho, editors, *Advances in AI*, volume 991 of *LNAI series*, pages 253–263, Berlin, DE, 1995. Springer-Verlag.
- [11] Mario Tokoro. The society of objects. Technical report SCSL-TR-93-018, Sony Computer Science Laboratory Inc., Tokyo, Japan, December 1993.