

Modeling Travel Assistant Agents: a graded BDI approach

Ana Casali¹, Lluís Godo² and Carles Sierra²

¹ Depto. de Sistemas e Informàtica
Facultad de Cs. Exactas, Ingeniería y Agrimensura - UNR
Av Pellegrini 250, 2000 Rosario, Argentina.
acasali@fceia.unr.edu.ar

² Institut d'Investigació en Intel·ligència Artificial (IIIA) - CSIC
Campus Universitat Autònoma de Barcelona s/n
08193 Bellaterra, Catalunya, España.
{godo, sierra}@iiia.csic.es

Abstract. In this paper, we use a graded BDI agent model based on multi-context systems to specify an architecture for a Travel Assistant Agent that helps a tourist to choose holiday packages. We outline the theories of the different contexts and the bridge rules and illustrate the overall reasoning process of our model.

1 Introduction

Nowadays, an increasing number of multiagent systems (MAS) are being designed and implemented. Several theories and architectures have been proposed to give these systems a formal support. Among them, a well-known intentional formal approach is the BDI architecture proposed by Rao and Georgeff [12]. This model is based on the explicit representation of the agent's beliefs (B), its desires (D), and its intentions (I). Indeed, this architecture has evolved over time and it has been applied, to some extent, in several of the most significant multiagent applications developed up to now.

On the other hand, knowledge representation and reasoning under uncertainty is an important traditional AI research field. In the recent past, approximate reasoning models have been used to help knowledge based systems to be more flexible and useful for real applications. In the frame of multiagents systems, i.e. in a distributed platform of autonomous, proactive, reactive and social agents, we wonder how the ideas underlying approximate reasoning could be extended and applied to these systems to enhance their knowledge representation capabilities. Actually, most of agent architectures proposed do not account for uncertain or gradual information. There are a few works that partially address this issue and emphasize the importance of graded models. Notably, Parsons and Giorgini [11] consider in the BDI model the belief quantification by using Evidence Theory. They also set out the importance of quantifying degrees in desires and intentions, but this aspect is not addressed in their work.

Please use the following format when citing this chapter:

Casali, A., Godo, L., Sierra, C., 2006, in IFIP International Federation for Information Processing, Volume 217, Artificial Intelligence in Theory and Practice, ed. M. Bramer, (Boston: Springer), pp. 415–424.

We consider that making the BDI architecture more flexible, will allow us to design and develop agents potentially capable to have a better performance in uncertain and dynamic environments. Along this research line we are concerned with developing a general model for Graded BDI Agents, specifying an architecture able to deal with the environment uncertainty and with graded mental attitudes, see [2, 3] for first results. In these works, belief degrees represent to what extent the agent believes a formula is true. Degrees of positive or negative desires allow the agent to set different levels of preference or rejection respectively. Intention degrees give also a preference measure but, in this case, modeling the cost/benefit trade off of reaching an agent's goal. Then, agents having different kinds of behavior can be modeled on the basis of the representation and interaction of these three attitudes. The graded BDI model we have developed is based on the notion of *multi-context system* (MCS) introduced by Giunchiglia et.al. [5] in order to help in the design of complex logical systems. This framework allows the definition of different formal components and their interrelation. In our approach, we use separate contexts to represent each modality and formalize each context with the most appropriate logic apparatus. The interactions between the components are specified by using inter-unit rules, called *bridge rules*. This approach has been used previously to model agent architectures as in [10], as a framework where the different components of the architecture and their interactions can be neatly represented.

Recently, the Artificial Intelligence community has made a great effort in the development of recommender systems and intelligent agents to help users confronted with situations in which they have too many options to choose from. These systems assist users to explore and to filter out their preferences from a number of different possibilities, many of them coming from the Web. A complete taxonomy of recommender systems can be found in [9]. Between their potential applications, the tourist domain seems to be a good candidate as the offers of tourism products are in constant growth. In this paper, using the graded BDI agent model presented in [2] and extended in [3], we propose an specific architecture for a Travel Assistant Agent, a recommender agent that helps a tourist to choose holiday packages in Argentina. For this purpose, we present the necessary theories for the different contexts and some of the bridge rules. In particular, we introduce some changes in the previously proposed Intention context [2, 3]. In order to evaluate the intention degree of a formula, other variables are taken into account and a set of more flexible functions is defined. In [3] we introduced a Social Context to filter the agent's incoming information, considering the trust in other agents. In this paper, we also extend this Social Context in order to represent the trust needed to decide whether or not to delegate some plans in other agents. By means of this recommender agent, our aim is to illustrate how our model can be used to specify particular agents that operate with graded attitudes and also to illustrate the overall reasoning process of our model.

This paper is organized as follows. In Section 2, the Graded BDI agent model is introduced. In Section 3 we specify the Travel Assistant Agent and

in its different subsections, its contexts and some of the main bridge rules are described. Finally, in Section 4 we present some conclusions and future work.

2 Graded BDI agent model

The architecture proposed is inspired by the work of Parsons et.al. [10] about multi-context BDI agents. The MCS specification of an agent contains three basic components: units or contexts, logics, and bridge rules, which channel the propagation of consequences among theories. Thus, an agent is defined as a group of interconnected units: $\langle \{C_i\}_{i \in I}, \Delta_{br} \rangle$, where each context $C_i \in \{C_i\}_{i \in I}$ is the tuple $C_i = \langle L_i, A_i, \Delta_i \rangle$ where L_i , A_i and Δ_i are the language, axioms, and inference rules respectively. When a theory $T_i \in L_i$ is associated with each unit, the specification of a particular agent is complete. The deduction mechanism of these systems is based on two kinds of inference rules, internal rules Δ_i , and bridge rules Δ_{br} , which allow to embed formulae into a context whenever the conditions of the bridge rule are satisfied. In our model, we have *mental* contexts to represent beliefs (BC), desires (DC), intentions (IC), and a social context (SC) which represents the trust in other agents. We also consider two *functional* contexts: for Planning (PC) and Communication (CC). In summary, the BDI agent model is defined as: $A_g = (\{BC, DC, IC, SC, PC, CC\}, \Delta_{br})$.

The overall behavior of the system will depend of the logic representation of each intentional notion in the different contexts and the bridge rules. Figure 1 shows the graded BDI agent proposed with the different contexts and some of the bridge rules relating them.

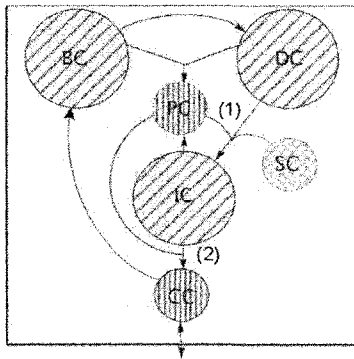


Fig. 1. Multicontext model of a graded BDI agent

In order to represent and reason about graded notions of beliefs, desires and intentions, we use a modal many-valued approach [7] where uncertainty reasoning is dealt with by defining suitable modal theories over suitable many-valued logics. For instance, let us consider a Belief context where belief degrees are to be modeled as probabilities. Then, for each classical formula φ , we consider

a modal formula $B\varphi$ which is interpreted as “ φ is probable”. This modal formula $B\varphi$ is then a *fuzzy* formula which may be more or less true, depending on the probability of φ . In particular, we can take as truth-value of $B\varphi$ precisely the probability of φ . Moreover, using a many-valued logic, we can express the governing axioms of probability theory as logical axioms involving modal formulae. Then, the many-valued logic machinery can be used to reason about the modal formulae $B\varphi$, which faithfully respect the uncertainty model chosen to represent the degrees of belief. To set up an adequate axiomatization for our belief context logic we need to combine axioms for the crisp formulae, axioms of Lukasiewicz logic for modal formulae, and additional axioms for B-modal formulae according to the probabilistic semantics of the B operator. The same many-valued logic approach is used to represent and reason under graded attitudes in the other mental contexts. The formalization of the adequate logics for the different contexts are described in [2, 3].

3 A Travel Assistant Agent

We have designed a Travel Assistant Agent (T-Agent) as an example of recommender agent using our graded BDI agent model. The T-Agent will be in charge of looking for different holidays plans in Argentina, in order to satisfy the desires of a tourist. The plan the T-Agent is expected to offer must be the best choice among the tourist packages supplied by a set of operators. The T-Agent will decide which plan to recommend, taking into account the interests of the tourist, the expected satisfaction of the preferences by the plan, its cost and the trust in the plan supplier. A schematic view of the T-agent and its interaction with the different tourist operators and the user is illustrated in Figure 2.

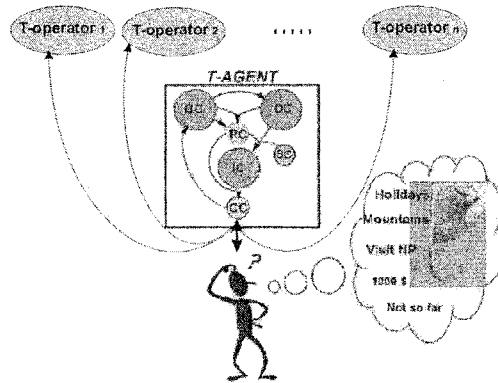


Fig. 2. The multiagent environment of the T-Agent

In the following subsections we outline the particular characteristics of the different contexts, specifying the necessary theories to complete the multicontext specification of the T-Agent.

3.1 Belief Context

The purpose of this context is to model the agent's beliefs about the environment. In order to represent beliefs, we use modal many-valued formulae, following the above mentioned logical framework and considering probability theory as its uncertainty model. In order to define the base (crisp) language, we extend a propositional language L to represent actions, taking advantage of Dynamic logic [2]. These actions, the environment transformations they cause, and their associated cost must be part of any situated agent's belief set. The propositional language L is thus extended to L_D , by adding to it action modalities of the form $[\alpha]$ where α is an action or plan. The interpretation of $[\alpha]\varphi$ is "after the execution of α , φ is true". We define a modal language BC over the language L_D to reason about the belief on crisp propositions. To do so, we extend the language L_D with a (fuzzy) unary modal operator B . If $\varphi \in L_D$, the intended meaning of $B\varphi$ is that " φ is probable". Then the B-modal formulae are built from elementary modal formulae $B\varphi$, and truth constants, using the connectives of Lukasiewicz many-valued logic (\rightarrow_L , $\&$). In this logic, modal formulae of the type $\bar{r} \rightarrow_L B\varphi$ express that the probability of φ is at least r and will be denoted as $(B\varphi, r)$.

The theory for the BC of the T-Agent contains:

- General knowledge about the tourism and Argentinian regions and destinations, the geographic characteristic of each region, activities allowed in each place, among others. We structure this knowledge inspired by existing tourism ontologies.
- Information about the tourist plans that the different operators provide. The plans are tourist packages and include the supplier, the cost and description of itinerary. They are structured as follows:

$package ::= (ID, Operator, Cost, [travel_1, stay_1, \dots, travel_n, stay_n, travel_{n+1}])$

where $travel_i$ is a description of the travel characteristics (e.g. type of transportation, travel length, etc.) and $stay_i$ includes destination, number of days, type of accommodation and activities. Each $travel_i$ and $stay_i$ is considered as atomic sub-plans of a set Π_0 , amenable to satisfy desires. Packages P are therefore modeled as composed plans, $\alpha_P \in \Pi$, alternating travel and stay sub-plans.

- Beliefs about how possible desires D (e.g. going to a mountain place or making rafting) are satisfied after executing different plans $\alpha \in \Pi$. Following the model presented, the truth-value of $B([\alpha]D)$ is the probability of having D after following plan α . For instance, the formula $(B[Atuel7/rafting, 0.9)$ expresses that the probability of satisfying the goal of making rafting as a consequence of the execution of the plan $Atuel7$ is greater than 0.9.

If a package P is composed by a number of subplans $\alpha_i \in \Pi_0$, that is $\alpha_P = \alpha_1; \dots; \alpha_n$, the truth-value r of $B([\alpha_P]D)$ will depend on the probabilities r_i of having D after the execution of the sub-plan α_i . Depending on the user's preferences of having the satisfaction of his desire in all

the sub-plans, in at least one of them, in most of them, ..., we include the following axiom in this context to model these possible preferences: $(B([\alpha_1]D), r_1) \wedge \dots \wedge (B([\alpha_n]D), r_n) \rightarrow (B([\alpha_P]D), \oplus_{i=1,n} r_i)$, where \oplus is an appropriate aggregation operator.

3.2 Desire Context

In this context, we represent the agent's desires. Desires represent the agent's *ideal* preferences regardless of the agent's current perception of the environment and regardless of the cost involved in actually achieving them. Inspired by the works on bipolarity representation of preferences by Benferhat et.al. [1], we suggest to formalize agent's desires also as positive and negative. Positive desires represent what the agent would like to be the case. Negative desires correspond to what the agent rejects or does not want to occur. Both, positive and negative desires can be graded. As for the *BC* language, the language *DC* is defined as an extension of a propositional language *L* by introducing two (fuzzy) modal operators D^+ and D^- . $D^+\varphi$ reads as " φ is positively desired" and its truth degree represents the agent's level of satisfaction would φ become true. $D^-\varphi$ reads as " φ is negatively desired" and its truth degree represents the agent's measure of disgust on φ becoming true.

In this context the tourist's desires will be expressed by a theory containing quantitative expressions about positive and negative preferences, These formulae express in different degrees what the tourist desires, e.g. $(D^+(\textit{mountain}), 0.8)$ or $(D^+(\textit{rafting}), 0.6)$, or what it rejects, e.g. $(D^-(\textit{northregion}), 0.9)$. These desires are the proactive elements of the recommender T-Agent and they start a chain of intra and inter-context deductions in order to determine which is the best touristic plan to recommend to the user.

3.3 Social Context

The aim of considering a Social Context (SC) in the T-agent architecture is to model the social aspects of agency. To do so, a key issue is the modeling of the agent's trust on other agents. In an agent community different kinds of trust are needed and should be modeled [4]. In [3] we used the notion of trust to asses the quality of the information received from other agents. Here, we consider the trust in the touristic package suppliers that interact with the T-Agent in order to evaluate the risk of touristic plans. Assuming we have a multiagent system scenario with a finite set of agents: $\{\textit{agent}_i\}$, $i \in I_{AG}$, the language for this context is a basic language *L* extended by a family of modal operators T_{ij} , where $i, j \in I_{AG}$. We consider the trust of an \textit{agent}_i toward an \textit{agent}_j about φ , $T_{ij}\varphi$, may be graded taking values in $[0,1]$, to express different levels of trust. Like in the other contexts, we use a many-valued approach for trust modelling. When the agent holding the trust is clear from the context we remove its subindex, that is, $T_{ij}\varphi$ becomes $T_j\varphi$. As for the modal formulae, we follow the intuition that the trust of $\varphi \wedge \psi$ may be taken as the minimum of the trusts

in φ and in ψ , hence we interpret the trust operator T_{ij} as a necessity measure on non-modal formulae, adding the corresponding axiomatics. The theory for SC in the T-Agent has formulae like $(T_j[\alpha]\varphi, t)$ expressing that the trust of the T-Agent toward an *agent_j* about a plan α directed to a goal φ , has degree greater than t . For this application, we consider that the trust depends only on the kind of touristic plan that the operator offers. Hence, we have proposed a plan classification based on a tourism ontology. For instance, we consider the region of the country as a classification element, since there are tour-operators that are good for plans in a particular region, but not in others. We believe that the trust in a provider *agent_j* is fundamental for the T-Agent to evaluate the risk in endorsing a plan α offered by *agent_j*. Then, as it was mentioned in the IC description, we introduce the trust degree as another variable that must be weighted in the computation of the intention degree. In previous works as in [8], it was considered that the plan quality could be computed as a weighted sum of a *standard rating* (combination of the benefit obtained by the plan execution and its cost) and a *cooperative rating* (evaluated from the trust in the agents involved). For the T-Agent, we propose a weighted combination of the different variables that is formalized in a bridge rule (see (1) in subsection 3.6).

3.4 Intention Context

This unit is used to represent the agent's intentions. Together with the desires, they represent the agent's preferences. However, we consider that intentions cannot depend just on the benefit of reaching a goal φ , but also on the world's state and the cost of transforming it into one where the formula φ is true. By allowing degrees in intentions we represent a measure of the cost/benefit relation involved in the agent's actions towards the goal. Moreover, when the execution of a plan involves the delegation of some actions to other agents, there is some risk that must be contemplated. We present two kinds of graded intentions, intention of a formula φ considering the execution of a particularly plan α , noted $I_\alpha\varphi$, and the final intention to φ , noted $I\varphi$, which take into account the best path to reach φ . Then, for each $\alpha \in \Pi$ we introduce a modal operator I_α , and a modal operator I , in the same way as we did in the other contexts. The intention to make φ true must be the consequence of finding a *feasible* plan α , that permits to achieve a state of the world where φ holds.

A theory for IC in the T-Agent represents those desires the user can intend by different feasible plans. Using this set of graded intentions, the T-Agent derives the final intention and the best recommended touristic plan. This theory is initially empty and will receive from a suitable bridge rule formulae like $(I_\alpha\varphi, i)$ for all the desires φ and for all the feasible plans α that the Planner context PC finds (see subsection 3.5). We consider that the degree of the intention is a function of different variables: the degree d of the desire that intends to satisfy $(D^+\varphi, d)$, the degree of belief r of having the desire after the execution of the plan $(B[\alpha]\varphi, r)$, the normalized cost of the plan c , and the reputation of the tourist supplier o of the plan $(T_o[\alpha]\varphi, t)$. The intention degree is computed as

some weighted average $i = f(d, r, c, t \mid w_d, w_r, w_c, w_t)$ by a bridge rule (see (1) in subsection 3.6) that gathers the different degrees d, r, c, t from the appropriate units and their corresponding weights w_d, w_r, w_c, w_t are set to match the tourist's requirements and constraints. For instance, a tourist with little money will increase the weight of the minimum cost criterion and a distrustful user will give more importance to the trust factor. Different functions will define distinct behaviors of the T-Agent. Moreover, for a particular function f , by choosing diverse set of weights the T-Agent can reach different degrees of intentions for a goal φ by a plan α . This allows the T-Agent to take more flexible decisions modeling the user's needs.

3.5 Planner and Communication Contexts

The nature of these contexts is functional and they are essential components of our model. In this work we only draft their functionalities in relation with the mental contexts presented. The Planner Context (PC) has to look for feasible plans in a repository of the touristic packages offered by the different supplier agents. All the touristic plans offered are introduced in the PC via the Communication Context. Within this context, we propose to use a first order language restricted to Horn clauses, where a theory of planning includes at least the following special predicates:

- $plan(\alpha, P, A, c)$ where $\alpha \in \Pi$ is the touristic package, P is the set of preconditions; A are the postconditions and $c \in [0, 1]$ is the normalised cost.
- $fplan(\varphi, \alpha, P, A, r, c)$ representing the feasible plan α towards the goal φ , where r is the belief degree of actually achieving φ by performing plan α .
- $bestplan(\varphi, \alpha, P, A, r, c)$ similar to the previous one, but only one instance with the best feasible plan is generated.

Each plan in order to satisfy a goal φ must be feasible, that is, the current state of the world must satisfy the preconditions, the plan must make true the positive desire the plan is built for, and cannot have any negative desire as post-condition. These feasible plans are computed within this unit using an appropriate planner that takes into account beliefs and desires injected by bridge rules from the BC and DC units respectively.

The Communication unit (CC) makes it possible to encapsulate the agent's internal structure by having a unique and well-defined interface with the environment. The theory inside this context will take care of the sending and receiving of messages to and from other agents in the multiagent society where our graded BDI agent lives.

3.6 Bridge Rules

For our T-Agent, we define a collection of basic bridge rules to set the interrelations between contexts. In this Section we comment the most relevant rules and we give an overview of how the T-Agent works.

As already mentioned in the previous section, there are bridge rules from BC and DC to PC that, from the positive and negative desires, the beliefs of

the agent regarding what the user can or cannot achieve through a particular plan, generate predicate instances in the PC unit that are used by the planner program to build the feasible plans.

Regarding intentions, there is a bridge rule that infers the degree of $I_\alpha\varphi$ for each feasible plan α that allows to achieve the goal φ . The intention degree is thought as a trade-off among the benefit of reaching a goal, the cost of the plan and the trust in its provider. The following bridge rule computes this value from the degree of $D^+\varphi$, the degree of belief $B[\alpha]\varphi$, the cost of the plan α and the trust t in the tourist supplier σ :

$$\frac{DC : (D^+\varphi, d), PC : fplan(\varphi, \alpha, P, A, r, c), SC : (T_\sigma[\alpha]\varphi, t)}{IC : (I_\alpha\varphi, f(d, r, c, t))} \quad (1)$$

Different functions f allow to model different agent behaviors. For instance, if we consider an *equilibrated agent* the function might be defined as a weighted average, where the different weights w_i are set according to the user's interests: $f(d, r, c, t) = (w_d d + w_r r + w_c (1 - c) + w_t t) / (w_d + w_r + w_c + w_t)$

The information supplied by the above bridge rule to the IC unit allows this unit to derive, for each goal φ , a formula $(I\varphi, i)$ where i is the maximum degree of all the $(I_\alpha\varphi, i_\alpha)$ formulas, where α is a feasible plan for φ . The plan α_b that allows to get the maximum intention degree i to φ will be set by the PC unit as the *best plan*. Finally, we also need rules to establish the agent's interaction with the user, meaning that if the T-Agent intends φ at degree i_{max} , then the T-Agent will recommend the plan α_b -*bestplan*- that will allow the tourist to reach the most intended goal φ :

$$\frac{IC : (I\varphi, i_{max}), PC : bestplan(\varphi, \alpha_b, P, A, c)}{CC : C(recommends(\alpha_b))} \quad (2)$$

3.7 Implementation

We are now implementing a prototype of this T-Agent in a multi thread version of prolog. Following previous work on implementation of BDI agents [6], we are implementing each mental unit (BC, DC, IC and SC) as a prolog thread, equipped with its own meta-interpreter. The meta-interpreter purpose is to manage inter-thread communication, i.e. all processes regarding bridge rule firing and assertion of bridge rule conclusions into the corresponding contexts. For efficiency reasons, the PC is implemented in the same thread than the BC as they have fluid information interchange when looking for feasible plans. The Communication unit is planned to be implemented in Java as a graphical user interface. This unit will be also in charge of the interchange of messages with the touristic supplier agents.

4 Conclusions and Future Work

We have presented a Travel Assistant Agent specification using our graded BDI agent model. This model allows us to define architectures that explicitly

represent the uncertainty of beliefs, graded desires and intentions. Using this framework we defined the T-Agent, a recommender agent for touristic plans in Argentina. The user's profile is incorporated in the T-Agent by introducing his preferences (positive and negative) and the importance he gives to the different variables that weigh in the selection of the plan. This profile together with the touristic information, constitute the knowledge base for the T-Agent's reasoning. With the specification of this concrete agent we aim at showing that our general model is useful and flexible to define particular recommender agents. As for future work, we are working to complete the implementation of the T-Agent architecture presented. This will also allow us to implement a number of particular agents of the T-Agent's family. These specific instances will be obtained by modifying different elements of the model, as the uncertainty model used in the mental contexts or the function that determines the intention degree, among others. This implementation will allow us to experiment and validate the formal model proposed.

Acknowledgments

The authors acknowledge partial support of the Spanish projects AECI PCI-Iberoamérica A/3541/05, TIN2004-07933-C0301 and TIC2003-08763-C02-00.

References

1. Benferhat S., Dubois D., Kaci S. and Prade, H. Bipolar representation and fusion of preferences in the possibilistic Logic framework. In *KR-2002*, 421-448, 2002.
2. Casali A., Godo Ll. and Sierra C. Graded BDI Models For Agent Architectures. J. Leite and P. Torroni (Eds.) *CLIMA V, LNAI 3487*, 126-143, 2005.
3. Casali A., Godo Ll. and Sierra C. Multi-Context Specification for Graded BDI Agents. Proceedings of *CONTEXT-05*, Research Report LIP 6, Paris, 2005.
4. Castelfranchi C. and Falcone R. Social Trust: A Cognitive Approach, in *Trust and Deception in Virtual Societies*, Kluwer Academic Publishers, 55-90, 2001.
5. Ghidini C. and Giunchiglia F. Local Model Semantics, or Contextual Reasoning = Locality + Compatibility *Artificial Intelligence*, 127(2):221-259, 2001.
6. Giovannucci A. Towards Multi-Context based Agents Implementation. IIIA-CSIC Research Report, 2004.
7. Godo, L., Esteva, F. and Hajek, P. Reasoning about probabilities using fuzzy logic. *Neural Network World*, 10:811-824, 2000.
8. Griffiths N. and Luck M. Cooperative Plan Selection Through Trust. In Proceedings of MAAMAW-99, 162-174, Spain, 1999.
9. Montaner M., López B. and de la Rosa L. A Taxonomy of Recommender Agents on the Internet. In *Artificial Intelligence Review* 19: 285-330, Kluwer, 2003.
10. Parsons, S., Sierra, C. and Jennings N.R. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3): 261-292, 1998.
11. Parsons, S. And Giorgini P. On using degrees of belief in BDI agents. In Proceedings of *IPMU-1998*, Paris, 1998.
12. Rao, A. And Georgeff M. Modeling Rational Agents within a BDI-Architecture. In *KR-92*, 473-484 (ed R. Fikes and E. Sandewall), Morgan Kaufmann, 1991.