Multi-context specification for Graded BDI-Agents

Ana Casali[1], Lluís Godo[2] and Carles Sierra[2]

[1] PhD Program in Information Technologies
EIA - University of Girona, Spain
UNR - Av Pellegrini 250, 2000 Rosario, Argentina.

[2] Institut d'Investigació en Intel·ligència Artificial (IIIA) - CSIC
08193 Bellaterra, Catalunya, España.

Abstract. In the recent past, an increasing number of multiagent systems have been designed to engineer complex distributed systems. Among the proposed architectures, one of the most widely used is the BDI agent presented by Rao and Georgeff. We consider that in order to apply agents in real domains, it is important for the formal models to incorporate uncertainty representation. With that aim we introduce in this work a general model for graded BDI agents, based on multi-context systems. The multicontext approach provides an adequate platform to specified agents, allowing a suitable representation of the graded mental attitudes, and their interactions. The architecture proposed serves as a blueprint to design different kinds of particular agents.

1 Introduction

In the recent past, an increasing number of multiagent systems (MAS) have been designed and implemented to engineer complex distributed systems. Several theories and architectures have been proposed to give these systems a formal support. A well-known intentional system formal approach is the BDI architecture proposed by Rao and Georgeff [13]. This model is based on the explicit representation of the agent's beliefs (B), its desires (D), and its intentions (I). This architecture has evolved over time and it has been applied in several of the most significant multiagent applications developed up to now.

Modeling different intentional notions by means of several modalities (B, D, I) can be very complex if only one logical framework is used. In order to help in the design of such complex logical systems Giunchiglia et.al. [6] introduced the notion of *multi-context system* (MCS). This framework allows the definition of different formal components and their interrelation. In our case, we propose to 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*. Contexts have been used in diverse applications and particularly, have been used to model agent's architectures [11, 14], as a framework where the different components of the architecture and their interactions, can be neatly represented. Indeed one advantage of the MCS logical approach to agency modeling is that it allows for rather affordable computational implementation. For instance, a portion of the framework described in [11] has been recently implemented using [5].

The agent architectures proposed so far mostly deal with two-valued information. Although the BDI model developed by Rao and Georgeff explicitly acknowledges that an agent's model of the world is incomplete, it makes no use of quantified information about how possible a particular world is to be the actual one. Neither does it allow desires and intentions to be quantified. There are a few works that partially address this issue and emphasize the importance of graded models. Notably, Parsons and Giorgini [12] consider the belief quantification by using Evidence Theory. They set out the importance of quantifying degrees in desires and intentions, but this is not covered by their work.

Our Thesis Project is about the development of a General Multi-context Model for Graded BDI Agents, specifying an architecture able to deal with the environment uncertainty and with graded mental attitudes. In this sense, belief degrees represent to what extent the agent believes a formula is true. Degrees of positive or negative desire 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. We consider the individual aspects of agency, defining clear semantics for the different mental attitudes and suitable logics to represent each one. These are formalized in the respective contexts for beliefs (BC), desires (DC), and intentions (IC), and two functional contexts for planning (PC) and communication (CC). These contexts are briefly schematized in the next Sections. We complete the agent architecture presenting in Section 8, a set of bridge rules. Preliminary results and an example of a Tourist Assistant Agent, illustrating the overall reasoning process of our graded model, can be found in [2]. In this paper, we present a revised version of the IC and we introduce the social context (SC) in our agent model, where the trust in others agents is modeled (outline in Section 6). Finally, in Section 9 we present some conclusions and future lines of work.

2 Graded BDI agent model

The architecture proposed is inspired by the work of Parsons et.al. [11] about multi-context BDI agents. Multi-context systems were introduced by Giunchiglia et.al. [6] to allow different formal components to be defined and interrelated. 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: $\left\langle \{C_i\}_{i\in I}, \Delta_{br} \right\rangle$, 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. They define the logic for the context and its basic behavior is constrained by the axioms. When a theory $T_i \in L_i$ is associated with each unit, the implementation of a particular agent is complete. Δ_{br} can be understood as rules of inference with premises and conclusions in different contexts. The deduction mechanism of these systems is based on two kinds of inference rules, internal rules Δ_i, and bridge rules Δ_{br}. Internal rules allow to draw consequences within a theory, while bridge rules allow to embed results from a theory into another [4]. 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: $A_g = (\{BC, DC, IC, SC, PC, CC\}, \Delta_{br})$.

This multicontext specification of our model of agent allow us to set a clear and independent semantics for each mental attitude. We do not intent to define a global semantics for our multicontext graded BDI agent. This approach let us to represent separately the different graded mental attitudes, considering that they must be treat with a suitable logic in each case. The overall behavior of the system will be result of the logic representation of each intentional notion in the different contexts and the bridge rules. These rules, constitude an operational part of the system behavior and permit us to modify the different theories.

In order to represent and reason about graded notions of beliefs, desires and intentions, we decide to use a modal many-valued approach. In particular, we shall follow the approach developed by Hájek et al. [8, 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. In this proposal, we choose the infinite-valued Łukasiewicz logic but another selection of many-valued logics may be done for each unit.

3 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. We consider the probability theory as the uncertainty model. Other models might be used as well by just modifying the corresponding axioms.

To reason about the credibility of crisp propositions, we define the BC language for belief representation, following Godo et al.'s [7], based on Łukasiewicz logic. In order to define the basic crisp language, we extend a propositional language L to represent actions, taking advantage of Dynamic logic [10]. These

actions, the environment transformations they cause, and their associated cost must be part of any situated agent's beliefs 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. More concretely, given a set Π_0 of symbols representing elementary actions, it can be defined the set Π of plans (composite actions). Then L_D is defined in the usual way [2], if $\psi \in L$ then $\psi \in L_D$ and if $\alpha \in \Pi$ and $\varphi \in L_D$ then $[\alpha]\varphi \in L_D$. The interpretation of $[\alpha]\, A$ is *"after the execution of α, A 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 crisp 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 believable". Formulae of BC are of two types:
- *Crisp (non B-modal):* they are the (crisp) formulae of L_D.
- *B-Modal:* they are built from elementary modal formulae $B\varphi$, and truth constants, using the connectives of Łukasiewicz many-valued logic:

• If $\varphi \in L_D$ then $B\varphi \in BC$
• If $r \in Q \cap [0,1]$ then $\overline{r} \in BC$
• If $\Phi, \Psi \in BC$ then $\Phi \to_L \Psi \in BC$ and $\Phi \& \Psi \in BC$ (where & and $\to_L$ correspond to the conjunction and implication of Łukasiewicz logic)

Other Łukasiewicz logic connectives can be defined from &, $\to_L$ and $\overline{0}$, for example $\neg_L \Phi$ is defined as $\Phi \to_L \overline{0}$. Since in Łukasiewicz logic a formula $\Phi \to_L \Psi$ is 1-true iff the truth value of Ψ is greater or equal to that of Φ, modal formulae of the type $\overline{r} \to_L B\varphi$ express that the probability of φ is at least r and will be denoted as $(B\varphi, r)$.

3.1 Belief Semantics and Axiomatization

The semantics for the language BC is defined, as usual in modal logics, using a Kripke structure. We have added to such structure a ρ function in order to represent the world transitions caused by actions, and a probability measure μ over worlds. Thus, we define the Kripke structure $M_B = \langle W, e, \mu, \rho \rangle$ where:

- W is a non-empty set of possible worlds.
- $e : V \times W \to \{0,1\}$ provides for each world $w \in W$ a Boolean evaluation of each propositional variable $p \in V$, that is, $e(p,w) \in \{0,1\}$.
- $\mu : 2^W \to [0,1]$ is a finitely additive probability measure such that for each crisp φ, the set $\{w \mid e(\varphi, w) = 1\}$ is measurable [8].
- $\rho : \Pi_0 \to 2^{W \times W}$ assigns to each elementary action a set of pairs of worlds denoting world transitions.

Extension of e to L_D formulae: e is extended to L using classical connectives and to formulae with action modalities –as $[\alpha]\, A$, setting: $e([\alpha]\, A, w) = min\, \{e(A, w_i) \mid (w, w_i) \in \rho(\alpha)\}$.
Extension of e to B-modal formulae: e is extended by means of Łukasiewicz logic truth-functions and the probabilistic interpretation of belief as follows:

- $e(B\varphi, w) = \mu(\{w' \in W \mid e(\varphi, w') = 1\})$, for each crisp φ
- $e(\overline{r}, w) = r$, for all $r \in Q \cap [0,1]$
- $e(\Phi \& \Psi, w) = \max(e(\Phi) + e(\Psi) - 1, 0)$
- $e(\Phi \to_L \Psi, w) = \min(1 - e(\Phi) + e(\Psi), 1)$

Finally, the truth degree of a formula Φ in a Kripke structure $M = \langle W, e, \mu, \rho \rangle$ is defined as $\|\Phi\|^M = \inf_{w \in W} e(\Phi, w)$. Notice, that a theory T in this context may include formulae like $B([\alpha]\varphi, r)$. According to the belief semantics, this formula means that the probability that the actual world is w' where $[\alpha]\varphi$ is true, is at least r. This probability captures in some sense, the probability of failure of plan α. As mentioned in Section 2, to set up an adequate axiomatization for our belief context logic we need to combine axioms for the crisp formulae, axioms of Łukasiewicz logic for modal formulae, and additional axioms for B-modal formulae according to the probabilistic semantics of the B operator. Hence, axioms and rules for the Belief context logic BC are as follows:

1. Axioms of propositional Dynamic logic for L_D formulae (see e.g. [10]).
2. Axioms of Łukasiewicz logic for modal formulae (see e.g. [8]).
3. Probabilistic axioms
   $B(\varphi \to \psi) \to_L (B\varphi \to B\psi)$
   $B\varphi \equiv \neg_L B(\varphi \wedge \neg\psi) \to_L B(\varphi \wedge \psi)$
   $\neg_L B\varphi \equiv B\neg\varphi$
4. Deduction rules for BC are: modus ponens, necessitation for $[\alpha]$ for each $\alpha \in \Pi$ (from φ derive $[\alpha]\varphi$), and necessitation for B (from φ derive $B\varphi$).

Actually, one can show that the above axiomatics is sound and complete with respect to the intended semantics described in the previous subsection (cf. [8]).

4 Desire Context

In this context, we represent the agent's desires. Desires represent the *ideal* agent's preferences regardless of the agent's current perception of the environment and regardless of the cost involved in actually achieving them. We deem important to distinguish what is positively desired from what is not rejected. According to the works on bipolarity representation of preferences by Benferhat et.al. [1], positive and negative information may be modeled in the framework of possibilistic logic. Inspired by this work, 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. As in BC logic, we will use a modal many-valued logic to formalise graded desires with Łukasiewicz logic as the base many-valued logic. In this context the agent's preferences will be expressed by a theory T containing quantitative expressions about positive and negative preferences, like $(D^+\varphi, \alpha)$ or $(D^-\psi, \beta)$, as well as qualitative expressions like $D^+\psi \to_L D^+\varphi$ (resp. $D^-\psi \to_L D^-\varphi$), expressing that φ is at least as preferred (resp. rejected) as ψ.

A complete formalization of semantics, in terms of generalized possibilistic Kripke structures, and a correct axiomatics can be seen in [2].

5 Intention Context

In this context, we represent the agent's intentions. They, as well as desires, 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 w and the cost of transforming it, into a world w' 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. The positive and negative desires are used as pro-active and restrictive tools respectively, in order to set intentions. Note that intentions depend on the agent's knowledge about the world, which may allow –or not– the agent to set a plan to change the world into a desired one. 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. Thus, if the agent's IC theory T contains the formula $I_\alpha\varphi \rightarrow_L I_\beta\varphi$ then the agent will try φ executing the plan β before than executing plan α. On the other hand, if T has the formula $I\psi \rightarrow_L I\varphi$ then the agent will try φ before ψ

We define IC Language in the similar way as we did with BC. Let L denote the basic propositional language and Π the set of actions corresponding to the dynamic propositional language L_D. 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. 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 DC. $I\varphi$ will represent that the agent intends φ by means of the "best plan" known. We use Łukasiewicz multivalued logic to represent the degree of the intentions. As in the other contexts, if the degree of $I_\alpha\varphi$ is δ, it may be considered that the truth degree of the expression "the goal φ is intended by means of action α" is δ. If the degree of $I\varphi$ is γ, it may be considered that the truth degree of the expression "the goal φ is intended" is γ. The computation of the degree of $I_\alpha\varphi$ for each feasible plan α is left to a suitable bridge rule (see (1) in next Section 8).

5.1 Semantics and axiomatization for IC

The semantics defined in this context shows that the value of the intentions depends on the formula intended to bring about and on the benefit the agent gets with it. It also depends on the agent's knowledge on possible plans that may change the world into one where the goal is true, and their associated cost. The models for IC are Kripke structures $M_I = \langle W, e, u\rangle$ where W and e are defined in the usual way, and $u : \Pi \rightarrow [0,1]^{W\times W}$ assigns to each action α a utility distribution $u_\alpha : W \times W \rightarrow [0,1]$ such that: $u_\alpha(w', w) = 0$, if $(w', w) \notin \rho(\alpha)$ and if $\alpha' = \langle\beta; \alpha; \gamma\rangle$ then $u_{\alpha'} \leq u_\alpha$. $u_\alpha(w', w)$ can be understood as the utility of reaching state w from state w', by means of action α. Then, the upper and lower bounds for the range of utilities of reaching φ from w' by means of α are:

- $U^+(w', I_\alpha\varphi) = \sup\{u_\alpha(w', w) : e(w, \varphi) = 1\}$
- $U^-(w', I_\alpha\varphi) = \inf\{u_\alpha(w', w) : e(w, \varphi) = 1\}$

Then, the truth evaluation of the $I_\alpha\varphi$ at w' can be in principle any value in the interval $[U^-(w', I_\alpha\varphi), U^+(w', I_\alpha\varphi)]$. There are in such a case extreme alternatives: a pessimistic attitude would correspond to take $e(w', I_\alpha\varphi) = U^-(w', I_\alpha\varphi)$, an optimistic attitude would be to take $e(w', I_\alpha\varphi) = U^+(w', I_\alpha\varphi)$. Notice that $U^+(w', .)$ and $U^-(w', .)$ are a possibility and a guaranteed possibility measures on formulas, respectively, so they can be axiomatized in Lukasiewicz logic [8].

6 Social Context

We introduce the Social Context (SC) to deal with the social aspects of agency. This context has the purpose of filter the agent's information interchange. The incoming information must be analyzed and filtered depending on the trust that the agent has in its source. To equip an agent with the social aspects, it is important to model and support the agent's trust. In an agent community different kind of trust are needed and should be modeled. In a first stage, the purpose of the SC in our model of agent, is to filter all the information coming from other agents. We inspired our work in the Belief, Inform and Trust (BIT) logic presented by Liu [9]. One of the central ideas formalized in BIT logic is that if $agent_i$ is informed by $agent_j$ about φ, the $agent_i$'s believes about φ depends on the trust the $agent_i$ has in $agent_j$ respect to φ. In the framework of this logic all the formulae, are crisp.

Considering we have a multiagent system scenario with a finite set of agents: $\{agent_i\}$, $i \in A$, the language for this context is a basic language L extended by a family of modal operators T_{ij}, where $i, j \in A$. We think that the trust of an $agent_i$ towards an $agent_j$ about a formula φ, $T_{ij}\varphi$, may be graded taking values in [0,1], to express different levels of trust. A belief-based degree of trust has been discussed in [3]. As in the other contexts, we used a many-valued treatment for the trust of an agent towards others. Then, if the degree of $T_{ij}\varphi$ is τ, we shall consider that the truth degree of the sentence "$agent_i$ trusts in $agent_j$ about φ" is τ. We again chose the Łukasiewicz logic as the many-valued logic.

The models for SC are defined in a similar way as we did in the other contexts using a Kripke structure. 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 formulas. Then, the corresponding axiomatics is set in a similar way than in IC for the pessimistic attitude. In a multiagent system scenario, if the $agent_i$ is informed by $agent_j$ that φ is true, then the statement $N_{ij}\varphi$ will be a crisp formula in the Communication context (CC). One of the central axioms for trust in the BIT logic [9] is: $(B_i I_{ij}\varphi \wedge T_{ij}\varphi) \rightarrow B_i\varphi$, where $i, j \in A$. We present a multi-context version of this axiom. As belief, inform and trust formulae are represented in different contexts, we use a bridge rule (see (5) in section 8) to formalize it, and we extend this rule to a many-valued framework.

7 Planner and Communication Contexts

The nature of these contexts is functional and they are needed components of our model. In this work we only draft their funtionalities in relation with the mental contexts presented. There is much work to do respect Planner and communication contexts, but is out of the scope of this work. The Planner Context (PC) has to build plans which allow the agent to move from its current world to another, where a given formula is satisfied. This change will indeed have an associated cost according to the actions involved. Within this context, we propose to use a first order language restricted to Horn clauses (PL), where a theory of planning includes the following special predicates:

- *action*(α, P, A, c_α) where $\alpha \in \Pi_0$ is an elementary action, $P \subset PL$ is the set of preconditions; $A \subset PL$ are the postconditions and $c_\alpha \in [0,1]$ is the normalised cost of the action.

- *plan*(φ, α, P, A, c_α) where $\alpha \in \Pi$ is a composite action representing the plan to achieve φ, P are the pre-conditions of α, A are the post-conditions $\varphi \in A$, c_α is the normalized cost of α.

- *bestplan*(φ, α, P, A, c_α) similar to the previous one, but only one instance with the best plan is generated.

Each plan 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 deduced by a bridge rule (see (1) in the next Section 8).

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. This unit has a propositional language with the modality N_{ij}, where $N_{ij}\varphi$ represents "$agent_i$ is informed by $agent_j$ about φ". The theory inside this context will take care of the sending and receiving of messages to and from other agents in the Multi Agent society where our graded BDI agents live.

8 Bridge Rules

For our BDI agent model, we define a collection of basic bridge rules to set the interrelations between contexts. In this Section we comment the most relevant rules of an $agent_i$[1].

From the positive and negative desires –represented by $\nabla(D^+\varphi)$ and $(D^-\psi, \mathit{threshold})$ respectively, the beliefs of the agent, and the possible transformations using actions, the Planner can build feasible plans. Furthermore, a filter is used to select the plans with a belief degree of achieving the goal after its execution, greater than some b-threshold. The following bridge rule does this:

$$\frac{\begin{array}{c} D : \nabla(D^+\varphi), D : (D^-\psi, \mathit{threshold}), P : \mathit{action}(\alpha, P, A, c), \\ B : (B([\alpha]\varphi), \mathit{bthreshold}), B : B(A \rightarrow \neg\psi) \end{array}}{P : \mathit{plan}(\varphi, \alpha, P, A, c)} \quad (1)$$

[1] We only explicit the i subscript in bridge rule (5) where another agent takes place

The intention degree trades off the benefit and the cost of reaching a goal. There is a bridge rule that infers the degree of $I_\alpha\varphi$ for each feasible plan α that allows to achieve the goal. This value is deduced from the degree of $D^+\varphi$ and the cost of the plan α. This degree is calculated by function f as follows:

$$\frac{D : (D^+\varphi, d), P : plan(\varphi, \alpha, P, A, c)}{I : (I_\alpha\varphi, f(d, c))} \tag{2}$$

Different functions model different individual behaviors. For example, the function might be defined as $f(d, c) = (d+(1-c))/2$, if we consider an *equilibrated agent*. Then, by means of an inference in the IC, $(I\varphi, i)$ will be computed, where i is the maximum degree of all the $I_\alpha\varphi$, where α is a feasible plan for φ. The *best plan* that allows to get the maximum intention degree to φ, will pass to the Planner.

We also need bridge rules to establish the agent's interactions with the environment, meaning that if the agent intends φ at degree i_{max}, the maximum degree of all the intentions. Then the agent will focus on the plan α_b –*bestplan*– that allows the agent to reach the most intended goal φ:

$$\frac{I : (I\varphi, i_{max}), P : bestplan(\varphi, \alpha_b, P, A, c_\alpha)}{C : C(does(\alpha_b))} \tag{3}$$

Through the communication unit the agent perceives all the changes in the environment and particularly receives the information from other agents. This information is introduced in the belief context, by the following bridge rule :

$$\frac{C_i : N_{ij}\varphi, S_i : (T_{ij}\varphi, \tau)}{B_i : (B\varphi, \tau)} \tag{4}$$

Figure 1 shows the graded BDI agent proposed with the different contexts and some of the bridge rules relating them.

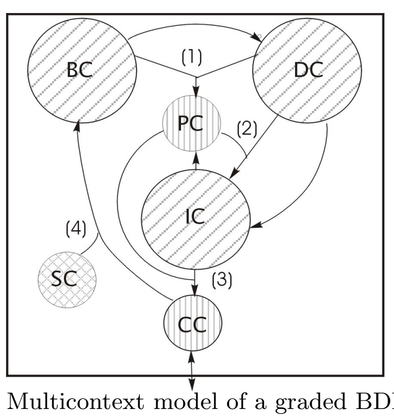

Fig. 1. Multicontext model of a graded BDI agent

9 Conclusions and Future Work

We have presented a BDI agent model that allows to explicitly represent the uncertainty of beliefs, graded desires and intentions. This architecture is specified using multi-context systems and is general enough to be able to model different types of agents. The agent's behavior is then determined by the different uncertainty measures used, the specific theories established for each unit, and the bridge rules. As for future work, we are considering two directions. On the one hand we want to continue with the extension of our multi-context agent model to a multiagent scenario, including other kind of relations with other agents. On the other hand, from an computational point of view, our idea is to implement each unit as a prolog thread, equipped with its own meta-interpreter, in charge of manage inter-thread (inter-context) communication. This implementation will support both, the generic definition of graded BDI agent architectures and the specific instances for particular types of agents. The implementation will also allow us to experiment and validate the formal model presented.

References

1. Benferhat S., Dubois D., Kaci S. and Prade, H. Bipolar representation and fusion of preferences in the possilistic Logic framework. In *KR-2002*, 421-448, 2002.
2. Casali A., Godo Ll. and Sierra C. Graded Models for BDI Agents. In Proceedings of *CLIMA V*, Leite J. and Torroni P. eds., 18-33. Lisboa, Portugal, 2004. Also in Proceedings of *EUMAS'04*, Barcelona, Spain, 2004.
3. Castelfranchi C. and Falcone R. Social Trust: A Cognitive Approach, in *Trust and Deception in Virtual Societies*, Kluwer Academic Publishers, 55-90, 2001.
4. Ghidini C. and Giunchiglia F. Local Model Semantics, or Contextual Reasoning = Locality + Compatibility *Artificial Intelligence*,127(2):221-259, 2001.
5. Giovannucci A. Towards Multi-Context based Agents Implementation. IIIA-CSIC Research Report, 2004.
6. Giunchiglia F. and Serafini L. Multilanguage Hierarchical Logics (or: How we can do without modal logics) *Journal of Artificial Intelligence*, vol.65, 29-70, 1994.
7. Godo, L., Esteva, F. and Hajek, P. Reasoning about probabilities using fuzzy logic. *Neural Network World*, 10:811–824, 2000.
8. Hájek, P. *Metamathematics of Fuzzy Logic*, vol.4 of Trends in Logic. Kluwer, 1998.
9. Liau C. J., Belief, Information Acquisition, and Trust in Multiagent Systems - a modal formulation. *Artificial Intelligence* 149, 31-60, 2003.
10. Meyer J. J. Dynamic Logic for Reasoning about Actions and Agents. *Workshop on Logic-Based Artificial Intelligence*, Washington, DC, 1999.
11. 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.
12. Parsons, S. And Giorgini P. On using degrees of belief in BDI agents. In Proceedings of *IPMU-1998*, Paris, 1998.
13. 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.
14. Sabater,J., Sierra, C., Parsons,S. and Jennings N. R. Engineering executable agents using multi-context systems. *Journal of Logic and Computation*12(3): 413-442, 2002.