Pragmatic-Strategic Reputation-Based Decisions in BDI Agents

Isaac Pinyol
Artificial Intelligence Research Institute (IIIA)
Spanish Scientific Research Council (CSIC)
Campus UAB, Bellaterra, Spain
ipinyol@iiia.csic.es

Jordi Sabater-Mir
Artificial Intelligence Research Institute (IIIA)
Spanish Scientific Research Council (CSIC)
Campus UAB, Bellaterra, Spain
jsabater@iiia.csic.es

ABSTRACT

Computational trust and reputation models have been recognized as one of the key technologies required to design and implement agent systems. These models manage and aggregate the information needed by agents to efficiently perform partner selection in uncertain situations. For simple applications, a game theoretical approach similar to that used in most models can suffice. However, if we want to undertake problems found in socially complex virtual societies, we need more sophisticated trust and reputation systems. In this context, reputation-based decisions that agents make take on special relevance and can be as important as the reputation model itself. In this paper, we propose a possible integration of a cognitive reputation model, Repage, into a cognitive BDI agent. We define a BDI model as a multi-context system whose regular logical reasoning process incorporates reputation information. After introducing the theoretical model we focus on an example to illustrate and analyze the behavior of our BDI agents in several situations.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agents

General Terms

Theory, Design

Keywords

Reputation and trust, Agent reasoning, BDI agents, Multi-context systems

1. INTRODUCTION

Computational trust and reputation models have been recognized as key to design and implementation multi-agent systems [9]. These models manage and aggregate the information needed by agents to efficiently select partners in uncertain situations. In recent years, several models have been developed [13]. For simple applications, a game theoretical approach similar to that used in most models can be sufficient. However, if we want to undertake problems found in socially complex virtual societies, more sophisticated trust and reputation systems based on solid cognitive theories are needed. One such cognitive theory is defined in [4].

This theory [4] proposes that agents evaluate the performances of other agents according to certain criteria. These evaluations (social evaluations from now on) can be only believed by the agents, only communicated by the agents or both believed and communicated. When a social evaluation is believed by a group of agents we refer to it as *image*. On the contrary, when a social evaluation circulates in the society we refer to it as *reputation*.

From this generic overview, the theory then develops a more individualistic vision. From a single agent, a specific image refers to a social evaluation in which the agent believes. At the same time, reputation is the belief of what is said. Therefore, an agent can have a good image of agent A as a seller, and at the same time acknowledge that A has a bad reputation as a seller. Furthermore, at this level, the theory describes a typology of possible decisions that autonomous agents can make involving social evaluations:

- *Epistemic decisions* cover the dynamics of beliefs regarding image and reputation, or in other words, decisions about updating and generating social evaluations.
- *Pragmatic-strategic decisions* are decisions of how to behave with potential partners using social evaluations information, and thus, how agents use these to reason.
- *Memetic decisions* refer to the decisions of how and when to spread social evaluations.

From a computational point of view, not all current state-of-the-art reputation models make a distinction between image and reputation, but all of them compute social evaluations. In fact, the field of reputation models has been mainly focused on epistemic decisions, while little attention has been paid to pragmatic-strategic and memetic decisions. Indeed, agents' decisions about how to use reputation information and how and when to spread them have been designed *ad-hoc* lacking any systematic or formal procedure. As mentioned before, this solution may suffice for simple environments, but in more complex societies pragmatic-strategic and memetic decision can be as important as epistemic decisions, and need more formal approaches as well.

This paper focuses on pragmatic-strategic decisions. Here the Repage cognitive reputation model is chosen as a paradigmatic example, because it is based on the cognitive theory

Cite as: Pragmatic-Strategic Reputation-Based Decisions in BDI Agents, Isaac Pinyol, Jordi Sabater-Mir, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1001–1008

defined in [4], to be integrated in a BDI (*Belief, Desire, Intention*) agent, providing then a formal integration of social evaluations in the agents' reasoning and decisions. To do so, in section 2 we state a preliminary section to briefly introduce Repage model and the formalism we use in the paper. In section 3 we formalize the specification of the cognitive BDI agent whose logical reasoning process incorporates image and reputation information. In section 4 we analyze some relevant reasoning points by presenting an example. Finally, in section 5 we conclude our analysis and propose some future work.

2. PRELIMINARIES

2.1 The Repage Model: Social Evaluations, Image and Reputation

Repage [14] is a computational system designed to be part of the agents architecture and based on a cognitive theory of reputation [4]. It provides social evaluations as image and reputation. A social evaluation is a generic term used to encapsulate the information resulting from the evaluation that an agent (evaluator) might make about another agent's (target's) performance regarding some skill, standard or norm (object of the evaluation). The object of the evaluation relies on which property of the target agents is evaluated. The value of the evaluation indicates how *good* or *bad* the performance resulted to be.

A social evaluation in Repage has three elements: a target agent, a role and a probability distribution over a set of labels. The target agent is the agent being evaluated. The role is the object of the evaluation and the probability distribution the value of the evaluation. The evaluator is the agent making the social evaluation.

The role uniquely identifies a kind of transaction and the classification of the possible outcomes. The current implementation of Repage considers five abstract labels for this classification: *Very bad, Bad, Neutral, Good, Very good*(*VB, B, N, G, VG* from now on), and assigns a probabilistic value to each label, however, we generalize it considering a finite number of labels $w_1, w_2 \ldots$ The *meaning* of each label must be contextualized depending on the role. For instance, we can represent a Repage image predicate as $img_i(j, seller,$[0.4, 0.2, 0.2, 0.1, 0.1]). This indicates that agent i holds an image predicate about agent j in the role of $seller$, and the value of the evaluation is [0.4, 0.2, 0.2, 0.1, 0.1]. This value reflects a probability distribution over the labels *VB, B, N, G, VG*. Then, it means that agent i believes that in the transaction of buying, when agent j acts as a seller, there is a probability of 0.4 to achieve a VB result (in the context of this transaction, this may mean a very low quality of the product), with a probability of 0.2 a B result, etc. For reputation predicates, it is the same as image, but instead, the agent believes that the evaluation is said by all or most of the agents in the group.

In section 3.7, we properly formalize these concepts. We refer to [14] for details on the calculus and the internal architecture.

2.2 The BC-logic

In order to integrate image and reputation information in a BDI reasoning process, the BDI agent must be endowed with a belief logic capable to capture the semantics of image and reputation. For this we use BC-logic [11], a probabilistic dynamic belief logic with a set of special modal operators. We are specially interested in the operators B_i and S, the first expressing what is believed by agent i, and the latter, what has been said by all the agents in the group respectively. The dynamic aspect of this logic is introduced by defining a set Π of actions. Then, for $\alpha \in \Pi$ and $\varphi \in BC$, formulas like $[\alpha]\varphi$ indicate that after the execution of α, the formula φ holds.

This logic incorporates specific axioms to reason about the probabilities of formulas by means of the operator Pr and constants $\overline{p}$ such that $p \in [0,1] \bigcap \mathbb{Q}$. It follows that for formulas $\varphi \in BC$, the expression $\overline{p} \leq Pr\varphi$ indicates that the probability of holding φ is higher or equal to p. This logic is based on the Logic of Knowledge and Probability introduced by Fagin and Halpern in [6].

BC-logic allows expressions like $B_i(\overline{p} \leq Pr([\alpha]\varphi))$. This indicates that agent i believes that the probability of holding φ after the execution of action α is at least p. Thereby, the formula $S(\overline{p} \leq Pr([\alpha]\varphi))$ expresses the same but in terms of what all agents have said. To simplify the notation, we will write expressions like $B_i(\overline{p} \leq Pr\varphi)$ as $(B_i\varphi, p)$, and $S(\overline{p} \leq Pr\varphi)$ as $(S\varphi, p)$.

This logic allows us to express image information in terms of beliefs $B_i\varphi$, and reputation information in terms of beliefs about what is said, $B_iS\varphi$ (see section 3.8). By grounding image and reputation into simple elements, we endow the agent with a powerful tool to reason over these concepts.

The complete syntax, semantics and axiomatization of BC-logic can be found at [11]. The belief operator follows the standard K, D, 4 and 5 axioms of modal logic, while operator S has its owns. The most interesting axioms are those that describe the interaction between S and B_i. These are closely related to the concept of *trust* that Demolombe in [5] defined regarding agents as information sources. The relationship of the two operators implies a relation between image and reputation at the belief level [11]. For instance, if for every φ the formula $B_i((S\varphi \to \varphi), p)$ holds, then agent i believes that what all agents say is really true with a probability p.

3. INTEGRATING REPAGE IN A MULTI-CONTEXT BDI AGENT

In this section, we propose a possible integration of Repage in a BDI agent. The underlying idea is to define a BDI agent, specified as a multi-context system, that uses the BC-logic presented in Section 2 to describe the belief base of the agent, and use a mechanism to translate image and reputation Repage predicates into BC-formulas.

3.1 Preliminaries and Related Work

Multi-context systems (MCS) provide a framework that allows several distinct theoretical components to be specified together with the mechanisms that link them together [8]. These systems are composed of a set of contexts (or units), and a set of bridge rules. Each context can be seen as a logic and a set of formulas written in that logic. Bridge rules are the mechanisms with which to infer information from one context to another. Each bridge rule has a set of antecedents (preconditions) and a consequent. The consequent is a formula that becomes true in the specific context when each antecedent holds in its respective context.

The use of MCS offers several advantages when specifying

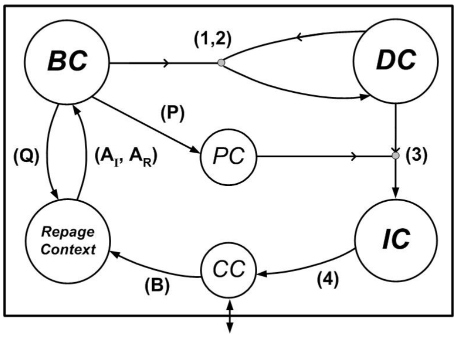

Figure 1: The Repage context embedded in a multi-context BDI agent. Circles represent context and arrows represent bridge rules.

and modeling agent architectures [15]. From a software engineering perspective, MCS supports modular architectures and encapsulation. From a logical modeling perspective, it allows the construction of agents with different and well-defined logics, keeping all formulas of the same logic in their corresponding context. This increases considerably the representation power of logical agents, and at the same time, simplifies their conceptualization.

The BDI architecture defined in [10] uses one context for each attitude; there is the belief context (B), the desire context (D) and the intention context (I). Each of them is equipped with a logic that corresponds to the premises that Rao and Georgeff [12] stated. Bridge rules among contexts determine the relationship between the attitudes and the type of agent: strong realism, realism and weak realism [12]. A communication context (C) is also included.

In [7], this specification is extended by means of a new commitment context, equipped with a deontic logic, creating then a new attitude of obligation. In [2] a multi-context BDI agent is specified and its attitudes are graded. Therefore, beliefs, desires and intentions are multi-valued with grades between 0 and 1. For our BDI model, we take the logic defined for desires and intentions described in [2] and [3].

3.2 The Multi-context BDI Model

The specification of our BDI agent as a multi-context system is formalized with the tuple $Ag = \langle\{BC, DC, IC, PC, CC, RC\}, \triangle_{br}\rangle$. These correspond to Belief, Desire, Intention, Planner, Communication and Repage contexts respectively. The set of bridge rules $\triangle_{br}$ incorporates the rules $1, 2, 3, 4, P, Q$ and B, shown in Figure 3, the bridge rules A_I and A_R shown in Figure 2, and rule B. Figure 1 shows a graphical representation of this multi-context specification. In the next sections we briefly explain each context and bridge rule.

3.3 Belief Context (BC)

This context contains the beliefs of the agent. We use the BC-logic introduced in Section 2. As said before, BC-logic uses the operator B_i for describing the knowledge that is believed by agent i. From now on, we will assume that we are specifying agent i.

The idea is to take advantage of the axioms and inference rules defined for BC-logic to reason about believed knowledge. Furthermore, since the logic allows probabilistic predicates, the agent can reason on uncertain information, and make estimations. Image predicates will be represented as expressions like $(B_i\varphi, g)$, and reputation with expressions like $B_i((S\varphi, g))$. Thus, inference rules and axioms relating B_i and S would be implicitly relating image and reputation information, as we expressed in section 2.2.

3.4 Desire context (DC)

This context deals with the desires of the agent. Like the BDI model described by Rao and Georgeff in [12], they are attitudes that are explicitly represented and that reflect the general objectives of the agent. We consider that desires are graded, and for that, we use the multi-valued logic (DC-logic) based on the Lukasiewicz logic described in [2].

DC-logic includes two fuzzy modal operators[1]: D_i^+ and D_i^-. The intended meaning of $D_i^+\varphi$ is that the formula φ is desired by agent i, and its truth degree, from 0 (minimum) to 1 (maximum), represents the level of satisfaction if φ holds. The intended meaning of $D_i^-\varphi$ is that φ is negatively desired, and the truth degree represents de level of disgust if φ holds. Also, DC-logic includes truth constants $\overline{r}$ where $r \in [0,1] \bigcap \mathbb{Q}$, and the connectives & and $\Rightarrow$ corresponding to the Lukasiewicz conjunction and implication respectively.

We differentiate generic from concrete desires. Generic desires define the general preferences of the agent, and are formulas like $D_i^*\phi$, where $*$ stands from $+$ or $-$ and ϕ does not contain any action. Concrete desires are formulas like $D_i^*[\alpha]\phi$ and define the desire to satisfy *phi* by execution action α. In our model, concrete desires are generated from generic desires and beliefs through bridge rules 1 and 2 (see section 3.8.2).

Because in Lukasiewicz logic the formula $\phi \Rightarrow \varphi$ is 1-true iff the truth value of φ is greater or equal to that of ϕ, and the truth value of $\overline{r}$ is exactly r, formulas like $\overline{r} \Rightarrow D_i^+\varphi$ in the theory of an agent indicate that the level of *satisfaction* of agent i is at least r if φ holds. The same with negative desires and the level of *disgust*. From now on we will write these formulas as $(D_i^+\varphi, r)$ and $(D_i^-\varphi, r)$. The semantics is given in terms of a positive and negative preference distributions over the possible worlds. We refer to [2] for technical details.

3.5 Intention Context (IC)

This context describes the intentions of the agent. Like in the Rao and Georgeff's BDI model [12], intentions are explicitly represented, but in our case generated from beliefs and desires. Also, we consider that intentions are graded, and for this we use the IC-logic defined in [2].

Similar to DC-logic, IC-logic defines the fuzzy modal operator $I_i\varphi$, indicating that agent i has the intention to achieve φ, and its truth degree (from 0 to 1) represents a measure of the trade-off between the benefit and counter-effects of achieving φ. Furthermore, IC-logic is defined in terms of a Lukasiewicz logic in the same way as DC-logic. Also, formulas like $\overline{r} \Rightarrow I_i\varphi$ will be written as $(I_i\varphi, r)$.

3.6 Planner Context (PC) and Communication Context (CC):

The logic in the Planner context is a first-order logic restricted to Horn clauses. In this first approach, this context only holds the special predicate *action*, which defines a primitive action together with its precondition. We look forward

[1]The original logic in [2] does not contain the reference to the agent. We include it to remark the desires of agent i.

$$\mathbf{A}_I: \quad \frac{RC : img_i(j, r, [V_{w_1}, V_{w_2}, \dots])}{\begin{array}{c} BC : (B_i([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1})) \\ BC : (B_i([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2})) \\ \dots \end{array}}$$

$$\mathbf{A}_R: \quad \frac{RC : rep_i(j, r, [V_{w_1}, V_{w_2}, \dots])}{\begin{array}{c} BC : (B_i(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_1}, V_{w_1}))) \\ BC : (B_i(S([\mathcal{R}_r(j)]\mathcal{T}_{r,w_2}, V_{w_2}))) \\ \dots \end{array}}$$

Figure 2: The bridge rules $\mathbf{A}_I$ and $\mathbf{A}_R$ (see Figure 1). They translate Image and Reputation predicates respectively into beliefs expressions in BC.

to introducing plans as a set of actions in the future. Communication context is a functional context as well, and its logic is also a first-order logic restricted to Horn clauses with the special predicates *does* (to perform actions), and rec_{ij} (to notify that agent i has received a communication from agent j).

3.7 Repage context (RC)

The Repage context contains the Repage model. We can assume that Repage predicates are specified in first-order logic restricted to Horn clauses, where the special predicates Img and Rep are defined. We write them as img_i $(j, r,$ $[V_{w_1}, V_{w_2}, \dots])$ and rep_i $(j, r,$ $[V_{w_1}, V_{w_2}, \dots])$, corresponding to the Image and Reputation of agent j playing the role r, from the point of view of i.

When in Repage the role and its labeled weights are defined, the role uniquely identifies which kind of transaction is part of, and each w_k identifies a predicate. To simplify, we can assume that the transaction identified by a role is summarized in a single action. To state this, we presuppose the definition of a mapping $\mathcal{R}_r$ between each role r and its action. In a similar way, we assume a mapping $\mathcal{T}_{r,w_k}$ between each role r and label w_k to a predicate.

We illustrate this with an example: In a typical market, the transaction of buying certain product involves two agents, one playing the role of buyer and the other playing the role of seller. From the point of view of the buyer, if she wants to evaluate other agents that play the role of seller, she knows that the associated action is *buy*. So, $\mathcal{R}_{seller}$ maps to *buy*. In the same way, the agent must know the meaning of each label w_k of Repage. Then, we can define that $\mathcal{T}_{seller,w_1}$ is $veryBadProduct$, $\mathcal{T}_{seller,w_2}$ is $okProduct$, etc.

In this mapping, the Repage predicate $img_i(j, seller, [0.2, 0.3, \dots])$ indicates that agent i believes that there is a probability of 0.2 that after executing the action $\mathcal{R}_{seller}$ (buy) with agent j as a seller, she will obtain a $\mathcal{T}_{seller,w_1}$ ($veryBadProduct$); with 0.3 that she will obtain $\mathcal{T}_{seller,w_2}$ ($OKproduct$), etc. With reputation predicates it is similar, but the concept is quite different. In this case it indicates that agent i believes that the corresponding evaluation is said by the agents in the group.

3.8 Bridge Rules

3.8.1 Bridge Rules A_I and A_R

Bridge rules A_I and A_R (see Figure 2) are in charge of generating the corresponding beliefs from images and reputations respectively. Notice that given a Repage social evaluation, these bridge rules generate one belief for each weight w_k. Both bridge rules use the belief operator (B_i) over certain formula, but meanwhile rule A_I states a knowledge that agent i believes as true, A_R states a knowledge that agent i believes to be said. They follow the definition of image and reputation we have given in section 2 and in the Repage context in section 3.7.

$$\mathbf{1}: \quad \frac{\begin{array}{c} DC : (D_i^+\varphi, d_\varphi) \\ BC : (B_i([\alpha]\varphi), p_\psi), BC : B_i(\psi \rightarrow \varphi) \end{array}}{DC : (D_i^+[\alpha]\varphi, g(d_\varphi, p_\psi))}$$

$$\mathbf{2}: \quad \frac{\begin{array}{c} DC : (D_i^-\varphi, d_\varphi) \\ BC : (B([\alpha]\psi), p_\psi), BC : B_i(\psi \rightarrow \varphi) \end{array}}{DC : (D_i^-[\alpha]\varphi, g(d_\varphi, p_\psi))}$$

$$\mathbf{3}: \quad \frac{\begin{array}{c} DC : (D_i^+[\alpha]\varphi, \delta), PC : action(\alpha, P), PC : P \\ DC : (D_i^-[\alpha]\psi_1, \delta_{\psi_1}), \dots, (D_i^-[\alpha]\psi_n, \delta_{\psi_n}) \\ \delta - \sum_{k=1}^n \delta_{\psi_k} \geq 0 \end{array}}{IC : (I_i[\alpha]\varphi, f(\delta, \sum_{k=1}^n \delta_{\psi_k}))}$$

$$\mathbf{4}: \quad \frac{IC : (I_i[\alpha]\varphi, \epsilon_{max})}{CC : does(\alpha)}$$

$$P, Q, B: \quad \frac{BC : B_i\varphi}{PC : \lceil B_i\varphi \rceil}, \quad \frac{BC : B_i\varphi}{RC : \lceil B_i\varphi \rceil}, \quad \frac{CC : rec_{ij}\varphi}{RC : \lceil rec_{ij}\varphi \rceil}$$

Figure 3: The bridge rules 1,2, 3, 4, P and Q respectively (see Figure 1)

3.8.2 Bridge Rules 1, 2, 3, 4

Bridge rules 1 and 2 (see Figure 3) transform generic desires to more concrete and realistic desires. To do this, these bridge rules merge generic desires from DC (with absolute values of satisfaction or disgust) with the information contained in BC, which includes the probability to achieve the desire by executing certain action. The result is a desire whose gradation has changed, becoming more realistic. This is calculated by the function g. If we define it as the product of both values, we obtain an expected level of satisfaction/disgust. Notice that we require that the belief information implies the achievement of the desired predicate.

Bridge rule 3 generates intentions. It takes into account both the expected level of satisfaction and the cost of the action. At the same time, executing an action to achieve certain formula can generate undesirable counter-effects. Thus, bridge rule 3 also takes into account the possible negative desires that can be reached by executing this action. In this bridge rule, for each positive realistic desire (D^+), we must include all negative desires (D^-) that can result from the same action. In this way we have the value of the positive desire (δ^+) and the sum of all negative desires (δ^-) that can be achieved by executing the same action. The strength of the intention that is created is defined by a function f. Different f functions would model different behaviors. In our examples we use the following definition: $f(\delta^+, \delta^-) = max(0, \delta^+ - \delta^-)$.

Finally, bridge rule 4 instantiates a unique intention (the one with maximum degree) and generates the corresponding action in the communication context.

3.8.3 Bridge Rules P, Q and B

Bridge rules P and Q allow the planner and Repage context respectively to be aware of the beliefs of the agent. The planner context uses this information to build plans, actions and their preconditions. Repage uses the information to configure the mappings $\mathcal{R}$ and $\mathcal{T}$, and to solve cognitive dissonances (see the example in Section 4, case 5).

Rule B reflects the reaction of the communication context once it receives communicated images, communicated reputation, third party images from other agents and fulfillment predicates. The content of these communications is directly introduced in Repage, which will update its information.

4. PUTTING THE MODEL TO WORK

In this section we analyze the reasoning processes performed by an executable version of the model presenting an example.

The base scenario we use involves a BDI agent (i) that, as a manager of a small restaurant, needs to periodically order wine to refill the stock. In this scenario, several providers are available. The information our agent wants to capture about them includes reliable information, for instance the price she will have to pay, but also uncertain information such as the delivery time of the orders and the quality of the wine. While reliable information is introduced as beliefs of probability 1, uncertain information will result in beliefs of lower probability values.

This situation can be formalized in multiple ways. We can define four possible pairwise disjoint predicates for the quality of the wine: *poorWine, averageWine, goodWine, excellentWine* (pW, aW, gW and eW from now on) and five pairwise disjoint predicates for the delivery time: $days(0,1)$, $days(1,3)$, $days(3,5)$, $days(5,10)$, $days(10,\infty)$ indicating respectively a delivery time of less than 1 day, between 1 and less than 3 days etc. Also we define the predicates *paid(X), paidLess(X), paidMore(X)* to indicate that the agent has paid X, less than X and more than X respectively. The predicate *budget(X)* indicates that the money she has in the budget is X. This knowledge and the implication among predicates must be introduced also as beliefs.

The interaction model defining the purchase of wine indicates that providers act as *wineSeller*s, but agent i wants to evaluate them in the two independent dimensions: the quality of the wine and the delivery time. Thus, Repage uses the roles $wineSeller(quality)$ and $wineSeller(dTime)$. The mapping $\mathcal{R}$ (see section 3.7) of these two roles points to the same action *buyWine* (*buy* from now on), which then summarizes the entire interaction model. The mapping $\mathcal{T}$ of the role $wineSeller(quality)$ relates w_1 to $poorWine$, w_2 to $averageWine$ etc, and the mapping $\mathcal{T}$ of the role $wineSeller(dTime)$ relates w_1 with $days(0,1)$, w_2 with $days(1,3)$, etc.

4.1 The Initial Knowledge

In this world, our agent knows the existence of four providers represented by *alice, bob, charlie* and *debra* respectively. Our agent is aware of their prices, and so this knowledge is introduced as beliefs:

$$(B_i[buy(alice)]hasWine \wedge paid(1000), 1)$$
$$(B_i[buy(bob)]hasWine \wedge paid(900), 1)$$
$$(B_i[buy(charlie)]hasWine \wedge paid(400), 1)$$
$$(B_i[buy(debra)]hasWine \wedge paid(1300), 1)$$

Bridge rule P introduces the information above into the planner context in order to generate the corresponding plans (simple actions in this case). It follows then, that in PC we might find

$$action(buy(alice), hasMoreMoney(1000))$$

indicating that the action of buying wine from *alice* is preconditioned on the budget having more than 1000.

4.2 Study cases

4.2.1 Exploring the space: case 1

Our agent is new to the business and only *trusts* her own direct experiences. Since she is just starting the business, she is mostly concerned about the quality of the wine rather than the delivery time. She has a budget of 1100 (*budget*(1100)) for the purchase. Regarding her desires, she would be satisfied with paying up to 1350 for an excellent wine. With the same strength she would be satisfied paying up to 800 for a good wine. In any case, she needs the wine. What she does not want is a poor or average wine. Lower on her priority list is obtaining the wine quickly, but still a long delivery time is not desired. These preferences can be formalized as desires in the DC as follows:

$$(D_i^+ hasWine \wedge paidLess(1350) \wedge eW, .9)$$
$$(D_i^+ hasWine \wedge paidLess(800) \wedge gW, .9)$$
$$(D_i^+ hasWine, .7),\ (D_i^- pW, 1), (D_i^- aW, .8)$$
$$(D_i^- days(10,\infty), .5),\ (D_i^- days(5,10), .4)$$

Since she does not have any information about the providers, Repage predicates contain the maximum possible uncertainty. For instance, for *charlie* the corresponding image predicates are:

$$img_i(charlie, wineSeller(quality), [.25, .25, .25, .25])$$
$$img_i(charlie, wineSeller(time), [.2, .2, .2, .2, .2])$$

Under these conditions the reasoning process leads to an almost random choice between three agents (*charlie*,*bob* and *alice*) to achieve the desire *hasWine*. In the following lines we briefly explain the most relevant steps.

Bridge rule A_I generates beliefs in the BC from images. As said before, the epistemic decision is not done at this rule but inside Repage, which computes image and reputation. In the case of *charlie* this rule is activated regarding the role $wineSeller(quality)$ as:

$$\frac{RC: img_i(charlie, wineSeller(quality), [.25, .25, .25, .25])}{\begin{array}{c}BC: (B_i[buy(charlie)]pW, .25),\ (B_i[buy(alice)]aW, .25),\\ (B_i[buy(charlie)]gW, .25),\ (B_i[buy(alice)]eW, .25)\end{array}}$$

All possible outcomes after buying from *charlie* have the same probability. In BC, because of the assumption that the quality and delivery time dimensions are stochastically independent, probabilistic inference rules of *BC*-logic are applied. For example, from $(B_i[buy(alice)]\ eW, .25)$ and $(B_i[buy(alice)]\ days(0,1), .2)$ can be deduced

$$(B_i[buy(charlie)]eW \wedge days(0,1), .05)$$

where .05 is the product of .25 and .2. In particular, and for the interest of our example, the following belief is also generated:

$$(B_i[buy(charlie)]hasWine \wedge paid(400) \wedge eW, .25)$$

Bridge rule 1 and 2 are executed for each generic positive and negative desire respectively. For instance, rule 1 is fired for the first desire as follows:

$$\frac{\begin{array}{c}DC: (D_i^+(hasWine \wedge paidLess(1350) \wedge eW), .9)\\ BC: (B_i[buy(charlie)](paid(400) \wedge eW, .25))\\ BC: B_i(paid(400) \rightarrow paidLess(1350))\end{array}}{(D_i^+[buy(charlie)](hasWine \wedge paidLess(1350) \wedge eW), g(.9, .25))}$$

If we consider that $g(x,y) = x \cdot y$, the resulting grade of the positive concrete desire is .225. It indicates that performing the action of buying from *charlie* to obtain an excellent wine and paying less than 1350 has an expected level of satisfaction of .225. Of course, for the same desire bridge rule 1 can be executed several times because different actions can lead to the same desire. Negative desires fire bridge rule 2 generating concrete negative desires. They indicate the expected level of disgust if the action is executed.

These negative desires are used in bridge rule 3 to take into account possible counter-effects of satisfying certain desire. Rule 3 is executed only one time for each positive concrete desire. For example, considering the desire above (we omit the predicate $hasWine$):

$$\frac{\begin{array}{c} DC : (D_i^+[buy(charlie)](paidLess(1350) \wedge eW), .225) \\ DC : (D_i^-[buy(charlie)][days(10, 1000)], .08) \\ DC : (D_i^-[buy(charlie)][days(5, 10)], .08) \\ DC : (D_i^-[buy(charlie)][aW], .2) \\ DC : (D_i^-[buy(charlie)][pW], .25) \\ PC : action(buy(charlie), budgetMore(400)) \\ PC : budget(1100) \rightarrow budgetMore(400) \end{array}}{IC : (I_i[buy(charlie)](paidLess(1350) \wedge eW), f(.225, .61))}$$

In this case, notice that the expected level of satisfaction of achieving the desire by buying from *charlie* is .225 but its counter-effects bring an expected level of disgust of .61. Taking $f(\delta^+, \delta^-) = max(0,\ \delta^+ - \ \delta^-)$, this intention has a grade of 0. Why would we perform an action if we expected from it to obtain more disgust that benefit?.

If the intention had the maximum degree, bridge rule 4 would generate the corresponding action. In this case, the intentions with a grade higher that 0 are:

$$(I_i[buy(charlie)], hasWine, .14), (I_i[buy(bob)], hasWine, .14)$$
$$(I_i[buy(alice)], hasWine, .14)$$

As expected, since Repage does not have any information and our agent needs to buy wine, a random choice can be made among these possibilities. Buying from *debra* is not considered because in rule 3 the precondition of having a budget greater than 1300 does not hold (see the action definition in the planner context). Assuming that she picks $(I_i[buy(charlie)], hasWine, .14)$, bridge rule 4 is fired executing the action $buy(charlie)$.

The result of this transaction fulfills the agent's desires in terms of delivery time and quality. This information is inserted into Repage by means of the bridge rule B. Repage evaluates the outcomes and updates the values of image and reputation. In the next reasoning process, this information will be introduced as beliefs by bridge rule A_I and A_R, as we have shown at the beginning of this case.

Continuing with our example, we suppose that *charlie* delivers the wine quite fast, in less than one day, but the quality of the wine is not very good. This makes Repage update image predicates as

$$img_i(charlie, wineSeller(quality), [.4, .4, .1, .1])$$
$$img_i(charlie, wineSeller(time), [.45, .25, .1, .1, .1])$$

We recall here that $w_1, w_2, \ldots$ in the role $wineSeller(quality)$ correspond to pW, aW,... meanwhile in the role $wineSeller(time)$ they correspond to $days(0, 1)$, $days(1, 3)$, ... respectively.

4.2.2 Receiving Reputation Information: case 2

After a while, our agent needs to buy more wine. She has exactly the same desires as before and the same budget, so she is mainly interested in the quality of the wine rather than delivery time. But this time, her image information about *charlie* has changed. Furthermore, we assume that she has received several reputation communications, about both *charlie* and *alice*. This information makes Repage generate the following reputation predicates:

$$rep_i(charlie, wineSeller(quality), [.5, .3, .1, .1])$$
$$rep_i(alice, wineSeller(quality), [.1, .2, .2, .4])$$

The reputation information regarding *charlie* coincides more or less with the image our agent has about him. This is not the case with *alice*. Through bridge rule A_R these predicates generate beliefs into BC. For *charlie*:

$$\frac{RC : rep_i(charlie, wineSeller(quality), [.5, .3, .1, .1])}{\begin{array}{c} BC : B_i((S[buy(charlie)]pW, .5)),\ B_i((S[buy(charlie)]aW, .3)), \\ B_i((S[buy(charlie)]gW, .1)),\ B_i((S[buy(charlie)]eW, .1)) \end{array}}$$

Note that these beliefs refer to what others say, not what our agent really believes. Since our agent only *trusts* herself, she does not take into account these predicates. In terms of the BC-logic it indicates that there is no relationship between operator S and operator B_i so far. This situation is also common: we can accept that a given person has a bad reputation, that most people *say* this, even when we believe the opposite [4].

Under these conditions, the reasoning process is similar to the previous case. This time though, *charlie* is no longer a possible choice, since the last experience with him was bad regarding the quality of the wine. Bridge rule 3 generates the intention to buy from *charlie* with a very low grade, in fact zero, since it is likely a poor or average wine would be delivered. In this case, the generated intentions are

$$(I_i[buy(bob)], hasWine, .14), (I_i[buy(alice)], hasWine, .14)$$

Our agent chooses *alice*. This time we suppose the result is in tune with the expectations of our agent; she obtains a good wine, even though the delivery time is not very fast. Repage updates image predicates regarding *alice* as follows:

$$img_i(alice, wineSeller(quality), [0, 0, .15, .85])$$
$$img_i(alice, wineSeller(time), [0, 0, 0, .1, .9])$$

4.2.3 Keeping the same desires: case 3

Maintaining the exact same desires as case 1 and 2, the next time that our agent wants to buy wine, she has the following intentions whose grade is higher that 0:

$$(I_i[buyWine(bob)]hasWine, .14),\ (I_i[buyWine(alice)]hasWine, .35)$$
$$(I_i[buyWine(alice)]hasWine \wedge paidLess(1350) \wedge eW, .365)$$

Since *alice* provided wine that was mostly excellent, and this is the main concern of our agent, she chooses again to buy from *alice*, but to satisfy the desire $hasWine \wedge paidLess(1350) \wedge eW$. The option to buy from bob appears due to the reminding uncertainty around his performance. We suppose that the resulting transaction confirms the same results as the previous case: an excellent wine but a long delivery time.

4.2.4 Changing Desires: case 4

This time our agent accepts the suddenly request to host a big birthday banquet that will take place in less than three days. Her cellar is not prepared for this event, so, she needs to order more wine. In this situation, her desires are different, since delivery time is now a key issue while the quality of the wine drops in importance:

$$(D_i^+ hasWine \wedge paidLess(1350) \wedge days(0, 1), .9)$$
$$(D_i^+ hasWine \wedge paidLess(800) \wedge days(1, 3), .7)$$
$$(D_i^- pW, .2), (D_i^- aW, .2)$$
$$(D_i^- days(10, \infty), .8),\ (D_i^- days(5, 10), .7)$$

Thanks to her previous interactions with the providers our agent already has some information about their performance. In this case, the only intention with a degree higher that 0 is

$$(I_i[buyWine(charlie)]hasWine \wedge paidLess(1350) \wedge days(0,1), .095)$$

She picks *charlie*, and the results are as the first time she bought from him in case 1: a short delivery time but a low quality.

4.2.5 Using Reputation Information: case 5

Several weeks after the successful banquet, our agent recuperates her initial desires and needs to order wine again. During this time she has heard about both *bob* and *debra*'s reputations which indicates both offer excellent wines and that furthermore *debra* is capable to deliver the order in a day. This is not the case with *bob*:

$$rep_i(bob, wineSeller(quality), [0, 0, .05, .95])$$
$$rep_i(bob, wineSeller(time), [.1, .2, .3, .3, .1])$$
$$rep_i(debra, wineSeller(quality), [0, 0, 0, 1])$$
$$rep_i(debra, wineSeller(time), [1, 0, 0, 0, 0])$$

This information is introduced through rule A_R as beliefs on what is said by the community. Unfortunately for our agent, *alice* notifies that she will not be available this time because she will be on holidays. Because of that, and because the reputation information she received in case 2 was in concordance with what she really believed, our agent starts *trusting* what others say, so she decides to take into account reputation information. Thus, the following predicate is introduced in BC for each φ:

$$B_i((S\varphi \rightarrow \varphi, .8))$$

From now on we will refer to this condition as condition S. This means that from now on, our agent believes that what is said ($S\varphi$) is really true with a probability of .8. This has big implications on the belief base of the agent. For instance, regarding *bob* in the role of $wineSeller(quality)$, rule A_R generates, among others, the following belief: $(B_iS[buy(bob)]$ eW, .95), meaning that people say there is a probability of .95 the wine will be excellent when buying from *bob*. Since our agent believes the condition S, it can be deduced that $(B_i[buy(bob)]\ eW, .76)$, where .76 is the product of .8 and .95 (see [11] for a formal proof on this result). This belief reflects what our agent really believes, and therefore, this information will be taken into account in the reasoning process. In this case, the only non-zero graded intention generated is

$$(I_i[buy(bob)]hasWine \wedge paidLess(1350) \wedge eW, .568)$$

We recall here that when Repage detects a reputation predicate about agent j regarding role r, and detects the predicate S (through bridge rule Q which introduce each belief into the Repage context), Repage eliminates any images predicate that were generated by default.

We suppose in this situation that the results are not as the agent expects, obtaining an average wine. Thus, Repage image predicates are updated as:

$$img_i(bob, wineSeller(quality), [.3, .4, .2, .1])$$
$$img_i(bob, wineSeller(time), [.1, .2, .3, .3, .1])$$

4.2.6 Image and Reputation Interference: case 6

Note that in the previous situation, the image about *bob* in the role $wineSeller(quality)$ contradicts *bob*'s reputation in the same role. This has already happened in case 2 with *alice*, but condition S was not believed. Now however, a problem arises: from the previous image and rule A_I, the belief $(B_i[buy(bob)]pW, .3)$ is generated. From the reputation predicate regarding *bob* and condition S, the belief $(B_i[buy(bob)]eW, .76)$ is also generated.

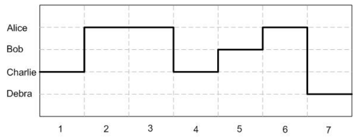

Figure 4: The choices of the agent throughout the situations explained in this section

Obtaining an excellent wine and a poor wine are disjoint predicates, so the generated beliefs results in an inconsistency knowledge base. This situation can be seen as a cognitive dissonance, i.e., two pieces of reliable information that clash. BC-logic does not allow inconsistency, and therefore, in the most general case the agent has two options:

(1)Solve the dissonance: If the BC-logic cannot accept inconsistencies, it could never happens. A source of inspiration for solving it lies in the area of belief revision. A belief revision (BR) system for certain logic is a set of postulates describing how knowledge is introduced, updated and withdrawn in the logic. When possible inconsistencies appear, the BR system has the power to specify *what* to change in the belief base to preserve consistency, for instance, by defining selection functions to pick formulas to eliminate.

Taking this idea, to solve the inconsistency we could eliminate image or reputation information, withdraw predicate S (or update it), or aggregate the information. What is clear is that Repage could deal with the problem since it is the generator of image and reputation. We recall here that, thanks to bridge rule Q Repage is aware of the agent's beliefs, and in particular, of the condition S. The original definition of Repage incorporates the *analizer*, a module prepared to solve cognitive dissonances regarding image and reputation that should be extended and redefined to incorporate the agents' beliefs. However, the internals of Repage are outside the scope of this paper, though we plan to work on it in the near future.

(2)Live with it: Another possibility is to redefine the BC-logic to permit these kind of inconsistencies. If the agent does not have enough information about how to resolve it, we can let the general reasoning process deal with the inconsistencies. We plan to analyze this option as well as part of the future.

Turning again to the example above, assuming that Repage has considered that image and reputation information are almost equally important, and that the information is contradictory, the resulting image is near the maximum uncertainty regarding *bob*. In this situation, the agent picks *alice*.

4.2.7 Increasing the budget: case 7

To conclude, we want to show the effect of a simple environment change. In this case, our agent decides to increase the wine budget to 2000. With exactly the same desires and the same reputation and image information as before, the reasoning process generates the maximum intention to buy from *debra*. This provider was always filtered out at bridge rule 3 because the precondition of buying from *debra* (to have more that 1300) was never fulfilled. Thus, the intention to buy from *debra* is only slightly higher than buying from *alice*.

4.3 Implementation Details

The scenario and each one of the situations have been implemented in Prolog[2]. An implementation of logical systems usually entails the simplification or limitation of some aspects of the logic. In our case, we assume that each logical formula is expressed as a horn clause and that modal operators are first-order predicates. Also, we do not accept logically omniscient agents that use a forward-reasoning engine, even when some implementations of multi-context systems use this approach [15]. Instead, we take advantage of the backward-reasoning engine of Prolog.

Note that the multi-context system specification of our BDI agent models an agent whose purpose is to execute a single action. This action is generated through rule 4 by choosing the intention of maximum grade. For this choice the agent must generate all possible intentions, which are created through rule 3 from desires, and so on. This schema follows a backward-reasoning algorithm that can be implemented in Prolog.

Thus, considering predefined knowledge as Prolog predicates, and inference rules and bridge rules as Prolog rules, the agent's reasoning can be started by asking Prolog to satisfy the predicate $does(A)$. While this is an oversimplification of what should be understood as multi-context systems, for simple examples the results are coherent and useful. We plan to study implementation issues in the future, an the effects of the simplifications in the desirable properties of the system.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a possible integration of a cognitive reputation model, Repage, in a BDI agent architecture. The agent has been specified using multi-context systems, where each attitude has been represented as a context. We used BC-logic, a probabilistic dynamic logic to represent the beliefs of the agent, allowing probabilistic reasoning. In particular we show how Repage social evaluations, image and reputation, are translated into probabilistic formulas written in BC-logic. The full reasoning process is done by also allowing graded desires and graded intentions, and stating appropriate bridge rules to relate them.

From the example it should be clear that on one hand epistemic decisions play a crucial role in the pragmatic-strategic decisions of the agent, and that a formal model for its integration improves the conceptualization of the reasoning process. On the other hand, the consequences of pragmatic-strategic decisions may effect the epistemic decisions.

As said before, in the future we are interested in studying the resolution of cognitive dissonances, situations in which the agent cannot decide which action to perform due to contradictory information. This research direction is somehow related to argumentation issues. Parsons *et al.* in [10] use a multi-context BDI agent to build an argumentation framework that we could adapt in our model.

Another important part of this research line involves the empirical study of certain properties regarding image and reputation through simulations and the implementation of the model using a logic-based multiagent platform, like JASON[1]. One point that we are specially interested is in the study on how graded trust conditions affect the overall performance of societies, and therefore, how the relation between image and reputation is relevant in determining the dynamics of the society.

[2]The source code can be download at `http://www.iiia.csic.es/~ipinyol/sourceAAMAS09.zip`.

6. ACKNOWLEDGMENTS

This work was supported by the EC under the FP6 programme [eRep project CIT5-028575]; the Spanish Education and Science Ministry [AEI project TIN2006-15662-C02-01, AT project CONSOLIDER CSD2007-0022, INGENIO 2010] and the Generalitat de Catalunya [2005-SGR-00093]. Also, we would like to give special thanks to Mao-Mei Liu for the useful comments on the manuscript.

7. REFERENCES

[1] R. H. Bordini, J. F. Hübner, and M. Wooldridge. *Programming Multi-Agent Systems in AgentSpeak Using Jason*. John Wiley and Sons, Ltd, 2007.

[2] A. Casali, L. Godo, and C. Sierra. Graded models for bdi agents. In J. Leite and P. Torroni, editors, *CLIMA V, Lisboa, Portugal*, pages 18—33, 2004.

[3] A. Casali, L. Godo, and C. Sierra. A logical framework to represent and reason about graded preferences and intentions. In *Proc. of KR'08, Sydney, Australia*, 2008.

[4] R. Conte and M. Paolucci. *Reputation in artificial societies: Social beliefs for social order*. Kluwer Academic Publishers, 2002.

[5] R. Demolombe. To trust information sources: a proposal for a modal logical framework. *Trust and deception in virtual societies*, pages 111–124, 2001.

[6] R. Fagin and J. Halpern. Reasoning about knowledge and probability. *J. ACM*, 41(2):340–367, 1994.

[7] D. Gaertner, P. Noriega, and C. Sierra. Extending the bdi architecture with commitments. In *Proceedings of the CCIA'06*, pages 247—257, 2006.

[8] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logic (or: How we can do without modal logics). *Journal of AI*, 65:29—70, 1994.

[9] M. Luck, P. McBurney, O. Shehory, and S. Willmott. *Agent Technology: Computing as Interaction (A Roadmap for Agent Based Computing)*. AgentLink, 2005.

[10] S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[11] I. Pinyol, J. Sabater-Mir, and P. Dellunde. Probabilistic dynamic belief logic for image and reputation. In *Proc. of the CCIA'08, Empuries, Spain.*, 2008.

[12] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proc. of KR'91*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA, 1991.

[13] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artif. Intel. Rev.*, 24(1):33–60, 2005.

[14] J. Sabater-Mir, M. Paolucci, and R. Conte. Repage: Reputation and image among limited autonomous partners. *JASSS*, 9(2), 2006.

[15] J. Sabater-Mir, C. Sierra, S. Parsons, and N. R. Jennings. Engineering executable agents using multi-context systems. *J. Logic and Comp.*, 12(3):413–442, 2002.