Cognitive and Affective Motivation in Conceptual Modelling

Antonio L. Furtado   Angelo E. M. Ciarlini

{furtado, angelo}@inf.puc-rio.br

> He made parallel lists of reasons for and against the move, giving each reason a numerical value. The sums being 110 points for remaining in England and 301 for going, he remained [Du, pg. 270].

Abstract. A proposal is presented towards the extension of conceptual models of information systems, in order to allow specification and simulation of the behaviour of agents with an adequate degree of realism. Our method is mainly based on rules to infer the goals of agents from situations holding at given states. In this paper, we argue that the rules should take into account both cognitive and affective characteristics, as can be conveyed, for the various agents, by their individual profiles and current internal states. Such characteristics should also influence the choice of strategies to handle goal interferences in multi-goal/multi-agent environments.

Keywords. Information systems, conceptual modelling, simulation, multi-agents, goals and plans, BDI-models, affective motivation.

Resumo. É apresentada uma proposta para a extensão dos modelos conceituais de sistemas de informação, de modo a permitir a especificação e a simulação do comportamento de agentes com um grau adequado de realismo. Nosso método baseia-se principalmente em regras para inferir os objetivos dos agentes, a partir de situações que valem em determinados estados. Neste artigo, afirmamos que as regras devem levar em conta tanto características cognitivas quanto afetivas, conforme representadas nos perfis individuais e estados internos correntes dos vários agentes. Tais características influenciam também na escolha de estratégias para lidar com interferências entre objetivos em ambientes multi-objetivos/multi-agentes.

Palavras-chaves. Sistemas de informação, modelagem conceitual, simulação, multi-agentes, objetivos e planos, modelos BDI, motivação afetiva.

1 Introduction

Our current research project aims at developing more realistic methods for the *conceptual specification* of information systems. Information systems cover, first of all, domains of practical applications, such as sales, banking, etc. Taking a temporal perspective, we can go beyond static descriptions to follow the *narratives* that arise in the mini-world delimited by the domain, consisting of events caused by the agents' interactions. Thus, in a banking application domain, one can usefully trace stories of clients handling their saving accounts and making investments, and their contacts with the management of the bank. But fiction also supplies domains, such as fairy-tales or detective stories, wherein descriptions and narratives are also amenable to computerized specification and simulation techniques [Ca,Sg]. The ability to handle fictional domains seems particularly relevant to the growing area of entertainment applications [DAB,Si].

To experiment with our methods, we have been developing prototypes, with the help of a software system combining Logic Programming with Constraint Programming [CW]. The prototypes, based on a plan recognition/ plan generation paradigm, are executable, thus allowing to perform *simulations* over the information system specified.

In our previous work [CF], we showed how to elaborate formal specifications at three levels for information systems having a database component:

1. At the **static level**, *facts* are classified according to the Entity-Relationship model. Thus, a fact may refer either to the existence of an entity instance, or to the values of its attributes, or to its relationships with other entity instances. Entity classes may form an *is-a* hierarchy. All kinds of facts are denoted by *predicates*. The set of all facts holding at a given instant of time constitutes a *database state.*

2. The **dynamic level** covers the *events* happening in the mini-world represented in the database. Thus, a real world event is perceived as a *transition* between database states. Our dynamic level schemas specify a fixed repertoire of domain-specific *operations*, as the only way to cause state transitions. Accordingly, we equate the notion of event with the execution of an operation. Operations are formally specified by the facts that should or should not hold as *pre-conditions* and by the facts added or deleted as the *effect* of execution.

3. The **behavioural level** models how agents are expected to act in the context of the system. To each individual agent (or agent class) A, we assign a set of *goal-inference rules*. A goal-inference rule A:S $\rightarrow$ G has, as antecedent, a *situation* S and, as consequent, a *goal* G, both of which are first-order logic expressions having database facts as terms. The meaning of the rule is that, if S is true at a database state, agent A will be motivated to act in order to bring about a state in which G holds. In addition, we indicate the *typical plans* (partially ordered sequences of operations) usually employed by the agents to achieve their goals [FC].

The first two levels encompass an *object-oriented* view of information systems, whereas the third level extends this view to incorporate an *agent orientation* [DLF,DLF,MCY].

Here, we shall concentrate on issues related to the design of goal-inference rules. To begin with, let us criticize some common naive assumptions:

a. Omniscience - Agents cannot be expected to *know* all facts currently holding. An agent may well ignore a fact, and may even have an erroneous notion about it. So, an agent A may fail to behave as predicted by a goal-inference rule A:S → G, of which he is supposed to be aware, simply because he does not know that the motivating situation S holds.

b. Competence for logical reasoning - Similarly, human beings are not equally proficient to apply precise methods — logical inference, probabilities, etc. — to reach conclusions. For example, practical experiments [TK] have demonstrated that people with training in statistics have been found to rate the occurrence of p ∧ q as more probable than facts p or q alone! Likewise, even if A knows that S holds, he may fail to apply an apparently well-understood goal-interference rule leading to G as an implied consequence.

c. Rationality - Far more disturbing is to note how a person well provided of factual knowledge and reasoning skills, after duly concluding that a goal G corresponds to the best course of action available at the moment, can decide against it with no declared justification. An eloquent example is the episode narrated in the epigraph, having as protagonist the English philosopher Herbert Spencer, who chose not to move to New Zealand, despite his conclusion that, according to his own evaluation, this would be more advantageous than staying in England [Du].

The present paper, which is to be read as a research proposal addressing the broad issues, rather than attempting a formal detailed treatment, surveys lines of investigation that may help us to drop these usually unwarranted assumptions, so as to model with an adequate degree of realism the behaviour of agents moving in practical or fictional domains.

2 Internal states and profiles - cognitive elements

As an example of a goal-inference rule, expressed in a semi-formal notation, consider:

employee E: position P is open → employee E wants position P

This formulation suggests that employee E is always able to know whether or not position P is open. To avoid assuming omniscience, an *internal state* can be attributed to each user A, registering the facts that A *believes*, correctly or not, as holding in the current global database state. We now establish that, for a rule A:S → G to affect the behaviour of A, it is not enough that the facts denoting situation S be objectively true; in addition, such facts must be *believed* by A, and, accordingly, be part of A's internal state.

Moreover, it is not enough that, believing S, A concludes that G is desirable. Except in cases where A's behaviour is purely *reactive*, he will still be free to decide, by some presumably objective criterion, whether or not he will actually *commit* to G as a goal and, consequently, adopt or develop a suitable plan Π to achieve it — which characterizes *deliberative* behaviour [SL]. Individual beliefs, rather than global knowledge, and the concept of *intentions*, as the result of purposefully adding commitment to mere desires, are among the basic tenets of **BDI**-models [CL,LL,RG].

Separate research has been applied to investigate what leads to this transition from desires to intentions. An approach that seems quite rational but is unfortunately hard to apply in domains of some complexity is based on the notion of *utility* [RN]. Firstly, it requires that the desirability of a goal G be expressed by a numerical *utility value*. This would seem to be easy whenever a number is naturally attached to G; for instance, G may consist of the possession of an amount m of money. But the same amount m will have a different importance to people of different income levels, and so the utility value u of G, although depending on m, would not be necessarily identical to it. If no quantitative attribute is attached to G, the determination of utility values becomes even harder. One should, at the very least, choose the values so as to ensure the ability to order situations according to their desirability; i.e. if G1 is intuitively more desirable than G2 then their respective utility values u1 and u2 should be determined so as to have u1 > u2.

An additional concern is that reaching a goal G by executing a plan Π is often an uncertain process. Instead of the purported G, the plan may achieve significantly different results G1, G2, ... , Gn, with probabilities p1, p2, ... , pn, respectively. Of course, replacing G by the n possible results of Π requires that different utility values u1, u2, ... , un be assigned to each Gi. The overall utility value of executing Π then becomes a statistical average, to be computed by a *utility function*:

$U(\Pi) = \sum_i p_i \times u_i$, for i = 1, 2, ... , n

Whenever there is more than a plan to reach a goal, the utility functions of all such plans have to be evaluated, and a "rational" agent should choose the plan of maximum utility. An analogous decision problem arises when an agent has to choose between two or more mutually exclusive goals (more about this in section 4), as in Spencer's dilemma. The difficulty of avoiding arbitrariness when determining utility values and the computational effort involved in the maximization calculations are drawbacks that must be recognized, since they can render unpractical the exhaustive comparison of all alternatives.

The adoption of internal states allows to consider what facts an agent A believes to be true at a given state, dropping thereby assumption (a). On the other hand, individual differences in logical reasoning competence, which underly assumption (b), as well as other relatively stable (i.e. state-independent) personal characteristics of agents, should be captured in *profiles*, to be specified for each agent class with as small granularity as convenient, and even, if necessary, specialized for individual agents. For an initial design of profiles and their corrections and adaptations, as experience may demand, the methods and techniques in the *stereotype* approach to user-modelling [PLL,Ri] seem promising.

Whereas both internal states and profiles might be restricted to diverse *cognitive* elements — basically related to awareness of facts and expertise to apply rules — in order to abolish assumptions (a) and (b), another type of elements must be brought in, if we propose to do without assumption (c) as well. Conceivably, Herbert Spencer decided to stay in England because he "felt" better staying there than moving to a remote country. Now, *feeling* is not a rare determinant in human decision-making, and a recent trend in Artificial Intelligence research is dedicated to what has been called *affective computing* [PVH].

3 Internal states and profiles - affective elements

One must recognize that behaviour is largely influenced, sometimes determined, by *drives* and *emotions*, among other affective elements (e.g. *moods*, not treated here) [Ba,Br,Ve]. There is already some recognition that believable agents, i.e. agents that provide the illusion of life, show emotions even when trying to behave rationally, and, to some extent, act under their influence; this remark is still more crucial in attempts to combine agent technologies with those of the entertainment industry, including cinema, computer games, and virtual reality [WJ,Wo]. *Drives* are basic physical needs, such as hunger and thirst, to which it is legitimate to add social needs, such as the urge to acquire money or prestige. *Emotions* have been classified according to distinct criteria, depending on the purpose of the classification; one popular classification considers six primitive emotions, with the convenient feature that they can be easily mapped into sharply distinguishable facial traits [DBH]: anger, disgust, fear, joy, sadness, surprise. An emotion can, in general, be either taken by itself, e.g. "a person is angry", or with respect to an object, another person or an event [OCF], e.g. "a person is angry at the prospect of losing a property".

Both drives and emotions are amenable to a numerical scale representation, showing their intensity within a lower and a higher limit. A drive or emotion is said to be present in an agent if its intensity measure excedes an appointed threshold. With the passage of time, the intensity of an unsatisfied drive increases. The satisfaction of a drive is accompanied by an increase in "positive" emotions (e.g. joy), whereas leaving it unsatisfied or, on the contrary, reaching an overwhelming regime by going beyond saturation, can stimulate "negative" emotions (anger, sadness). The intensity of an emotion decays after some time. Certain emotions are able to excite or inhibit other emotions, e.g. fear may excite anger and inhibit joy. As would be expected, the assignment of numerical measures is a no less delicate process here, needing to be validated for its adequacy in actual practice. Curious experiments [Ve] to emulate a toddler with purely reactive behaviour have been conducted, dealing with mutually stimulating or inhibiting interactions among various drives and emotions.

The intensities of drives and emotions of an agent A must be recorded as part of the internal state of A. Likewise, *personality traits* of A (e.g. whether or not A is an introvert, or is aggressive, etc.) should be part of A's profile. Thus, both internal states and profiles would contain both a *cognitive component* and an *affective component* which, together, contribute to determine A's behaviour. Indeed, it seems clear that most people decide under the combined influence of rational and affective factors. Both kinds of factors should therefore be taken into consideration as constituents of

pre-conditions and effects of operations and, hence, of situations and goals in goal-inference rules. Additionally, they should both be taken into account in the determination of utility values; in special, the satisfaction of fundamental physical and social drives should be at the root level in goal hierarchies.

Examples of the relevance of affective factors are easy to find. A client will buy a merchandise from a salesman only if he is happy with the salesman's service. Delays in delivery will have the effect of increasing the client's anger against the salesman. The action of watching a movie aims mostly at procuring pleasurable emotions, not rationally determined profit. Some goals, like the purchase of food, owe their desirability to their contribution to satisfy a drive, such as hunger. Computer interfaces offering unrequested advice may cause anger in users with a high degree of expertise on the subject, as their profile should indicate — and preventively lead the system to turn down the advice-giving facility. For Herbert Spencer, the utility value of going to New Zealand may have been reduced by his fear of facing a new physical and human environment. Making or keeping a client happy has been defined as a "softgoal" in the requirements analysis literature [DLF,MCY], due to the imprecision of the notion of "happy"; one may expect that, by numerically measuring emotions and the increasing or decreasing effect that operations can have on their level, may contribute towards the treatment of softgoals as ordinary goals.

We shall now revise the example goal-inference rule indicated at the beginning of section 2. Always in a semi-formal notation, the situation antecedent will be rewritten, combining terms which express cognitive and affective elements taken from the agent's profile and internal state. The objective fact that position P must be open for the antecedent to be true is replaced by the agent's belief that this fact holds; a personality trait of the agent (ambition), a social drive (to support his family) and another person's feeling with respect to his performance are added as affective elements; and notice the mixture of cognitive/affective aspects, in that the agent has a belief about an affective element of the other agent's internal state.

employee E: E believes that position P is open and E is ambitious and E needs a larger income to support his family and E believes that the section manager appreciates his work with intensity $\geq n1 \rightarrow$ employee E wants position P with utility value n2

4 Multi-goal/multi-agent environments

4.1 Willensky's classification of goal interferences

At a given state, some of the goals resulting from the application of goal-inference rules may form one or more sets of mutually interfering goals. Robert Willensky [Wi] notes that interferences can be separately characterized, on the one hand, as *negative* or *positive*, and, on the other hand, as *internal* (involving goals of the same agent) or *external* (goals of different agents). On the basis of these two dimensions, he proposes the following classification of goal interferences :

   a. *goal conflict* : negative, internal;
   b. *goal competition*: negative, external;
   c. *goal overlap* : positive, internal;

d. *goal concord*: positive, external.

In both types of negative interferences, the goals involved are somehow incompatible. The incompatibility may be caused by resource limitations; e.g. a customer wants to buy two merchandises, but the amount of money he has available would only be enough for one purchase. Another possible cause of incompatibility is that fulfilling the goals would lead to mutually exclusive states; e.g. two members of an organizing committee want to appoint different dates or sites for a conference. Also, it is common to happen that, to pursue a goal, an agent may endanger a preservation goal; e.g. an employee wants to watch a game, when he should be present at a meeting — thereby risking to lose his job.

There are several strategies to handle conflicting goals. It may be possible to *resolve the conflict* by obtaining more resources, if their insufficiency is the problem, or by changing the circumstances in case of mutually exclusive states (for example, by obtaining that the date of a meeting be postponed). Another strategy is to *abandon* one or more goals, partially or totally. To decide which goals should be retained, one criterion is to search the maximum total utility value. But, instead of partial or total abandonment, one may prefer *substitution* across similar goals. For instance, "watching a game" and "hearing a concert" are different specializations of "entertainment" in a goal is-a hierarchy; so, it would not be unnatural to trade one for the other.

In some occasions the conflict is not between the goals themselves, but arises from the plans initially adopted for their achievement. If this is the case, the adoption of an alternative plan may remove the problem. For instance, visiting clients in separate towns in the same day may be impossible to a salesman if he travels by train, but may become feasible by airplane.

Similarly, for competition, sometimes it may be possible to *avoid the competition*; if two employees dispute the same position, the creation of a second position might accomodate both of them. *Abandonment* is another strategy, either spontaneous or induced. However, *going ahead with the competition* can also be considered, with two options: an agent may try either to *outdo*, i.e. do better than his competitor, or to *undo*, i.e. act to frustrate the competitor's plan; e.g. an employee can take training courses to improve his efficiency, or, alternatively, he can try to show that the other employee is disqualified to fill the position.

Overlapping goals are often more profitably reached by *developing a single plan* to reach them; e.g. a person may want to buy two merchandises and then realizes that they are available in stores not far away from each other, so that both stores can be visited by car, with no detour, when coming home from office. Overlapping may also work as a *reinforcement*: the agent may find that performing a certain action may aim at more than one goal, as when a person observes that a vegetarian diet serves both the goal of improving health and of reducing expenses with food. Or an action may *subsume* a recurring goal, as when an employee decides to rent an apartment closer to the working place where he must be present everyday.

In cases of goal concord, agents may *join efforts* in order to reach their goals more efficiently and without wasting resources. This strategy may consist of *pooling*

resources, *dividing tasks*, and is implemented, in general, by *coordinating or merging the plans of the agents.* Treaties of mutual assistance between nations with a common enemy serve as an example. It is curious to note that what initially appears as competition can sometimes be treated as concord, as happens with competing industrial firms forming a cartel to jointly raise the price of a product.

4.2 Cognitive and affective considerations

A major cognitive requirement in multi-agent environments has to do with the need to establish communication in order to adapt interfering goals and the corresponding plans. We saw that each agent perceives the external world in terms of beliefs, which are part of his internal state. Communication [Ma] between agents then means the ability of one agent to act on the other agent's internal state, changing his beliefs, typically by an exchange of information. *Speech acts* [CP] thus provide an additional repertoire of operations — such as *inform* and *request* — noting that the latter is essential whenever an agent A1 wants another agent A2 to perform an operation which A2, but not A1, is authorized to execute. Operations corresponding to speech acts, besides being included in plans, intermingled with the domain-specific operations, can serve as a basis for *agent communication languages* [FFM].

But speech actions go beyond their cognitive effect. They are associated with emotions, which, in turn, may be manifested by facial expressions [PBS].

More generally, affective considerations certainly influence the choice of strategies for handling the various cases of interfering goals. Temperament traits, which we propose to model as part of the agents' profiles, may establish a preference for either goal abandonment or for aggressive outdo, or even undo, competitive acting. A prototype reported in [El] has been developed to help training salesmen by simulating their interaction with clients with four different personalities: dominant, political, steady, and wary; the same actions of a salesman were expected to elicit different reactions in each case.

A study of emotions that stresses interpersonal relationships [OO], and was used in the training prototype mentioned above, attempts to formally characterize what is meant by a number of words and phrases expressing emotions closely related to behaviour, grouped as follows:

Well-being:	joy, distress;
Fortunes-of-others:	happy-for, gloating, resentment, sorry-for;
Prospect-based:	hope, satisfaction, relief, fear, fears-confirmed, disappointment;
Attribution:	pride, admiration, shame, reproach;
Attraction:	love, hate;
Well-being&attribution:	anger, remorse.

Gloating, for example, as analyzed in the corresponding expression in the situation calculus formalism, means to be pleased about an event undesirable for another agent. Reproach is disapproving of the action of another agent, assuming that the action is considered blameworthy. Love and hate (or like and dislike) are not decomposed into simpler terms, being considered primitive and hence unexplainable.

Such kinds of emotions may well play a role in the choice of strategies. In a pair of employees competing for the same position, one employee may find that an undo strategy is justified if he feels reproach for past actions of the other employee. On the contrary, he may spontaneously abandon his claim to the position, especially if he has a benevolent personality, in view of his admiration for the competitor. Individual agents may reconsider their goals to better suit the needs of a group to which they belong; in [OO], for instance, an agent can exhibit pride or shame for, respectively, a praiseworthy or blameworthy act attributed to a "cognitive unit" of which he is a member.

Going further, if the agents involved are not individual persons or groups, but rather industrial firms or some other kind of *organization*, it becomes far more difficult to characterize their activity in cognitive and affective terms. For human agents, computer scientists seek the orientation of Cognitive Psychology [KW,RH]. For organizations, fortunately, some clues are provided by Management Science, in particular from studies on Theories of Organization. Showing that the various proposed theories can be classified according to the metaphor through which they visualize what the concept of "organization" signifies, Gareth Morgan [Mo] argues convincingly that all classes of theories have important contributions to offer; for instance, whereas mechanistic theories stress a rational concern with efficiency and profit, other theories detect practices inherent in the company's traditional "culture", or the pressure of hidden agendas emerging from political struggles for power, etc.

4.3 Choosing strategies for goal interferences - an outline of the problem

The discussion above suggests that the goals generated by goal-inference rules may need to be revised in view of the various kinds of interferences, and that the behavioural level of specification of an information system should include some indication on how agents will proceed towards this revision process.

Roughly speaking, the following steps should be executed, whenever a system reaches a new state:

1. All applicable goal-inference rules are activated, and each goal G_A of each agent A is included in a set Γ.

2. Cases of interference among goals of the same agent or of different agents are identified and classified according to Willenky's two-dimension criterion [Wi]. For each case i of interference the interfering goals are collected in a tagged *interference-set* κ/C_i, where the tag κ indicates which of the four types of interference is involved.

3. Each agent A chooses a strategy (possibly from the repertoire proposed by Willensky) to handle each of his goals G_A present in each interference-set.

4. A new set of goals Γ' is obtained by changing certain of the interfering goals, according to the chosen strategies. Instead of or in addition to such changes,

the strategies may prescribe requirements to be met when generating plans to reach the (possibly altered) goals.

Step 2 is already nontrivial. Detecting a case of interference is a difficult problem; for a conflict, for instance, one must prove that the joint satisfaction of certain goals would lead to a (logically) inconsistent state or — in a special case of inconsistency, usually involving numeric calculation — would consume more than the available resources. And step 3 is a particularly hard decision problem, since, in general, more than one strategy is applicable in each case, and suitable choice criteria should be provided (more on that in the sequel). And the complexity of the overall process increases considerably when we consider further difficulties such as:

a. A goal G_A may figure in more than one interference-set and the strategies separately chosen to handle it may lead to different and irreconcilable changes and/or planning requirements.

b. Modified goals may cause further interferences, which should be in turn detected and duly handled.

c. The sequence, arbitrary or not, adopted by agents to examine their interfering goals may be significant, in that a different sequence may result in a different revised set of goals Γ'.

While difficulty (a) means that the process may lead to inconsistencies, (b) suggests that it may fail to terminate and (c) that it may not be deterministic. And notice that we must avoid reinstalling an omniscience assumption: agents do not necessarily know about every initial goal of other agents and about the successive changes introduced along the revision process. Communications, as enabled by the speech acts mentioned in the previous section, are a vital component of the process, ideally conducing to a *negotiation* phase, in which the chances are greater for a mutually satisfactory converging solution.

To assist in the execution of step 3, it may be useful to introduce an additional kind of rules in the behavioral schema: the *strategy-selection rules*, with the following structure:

<agent>: <interference-set>, <selection-formula> $\rightarrow$ <strategy-directive>

where, <selection-formula> is a conditional expression whose terms may refer to cognitive and affective information conveyed by the agent's profile and by the agent's internal state, and also to the utility of the agent's goals. The selection-formula may be designed so as to achieve an adequate degree of specialization for each rule, ranging from very general rules, often determined by personality traits indicated in the profile, to specialized opportunistic rules, finely tuned to the circumstances of the situation as seen from the perspective of the agent's current internal state.

Affective personality traits, such as aggressiveness, may be so intense that the agent may, in all circumstances, adopt some preferred strategy, such as, for example, outdo or undo to deal with any case of competition. On the contrary, a non-obsessive agent may be more flexible in his reactions, allowing himself an ample choice among

the possible strategies, based on a detailed assessment of the situation, sometimes taking into account the present feelings of the agent with respect to the other agents [OO] involved in the interference, as well as his beliefs about their situation and goals. And one may even find the conventional agent with an authentically "rational" personality (which Herbert Spencer proved <u>not</u> to have), for whom the utility values predominate, leading to the strategy affording the largest gain.

The consequent <strategy-directive>, to be adopted by the agent if the interference typified by the interference-set occurs and the selection-formula is true, must be based on one of the strategies appropriate for the type of interference involved; as indicated before, it may involve changes to one or more goals of the agent that are present in the interference-set and/or requirements for the plans to be generated for such goals.

For instance, if the interference is "competition", one of the possible strategies is "undo"; if a rule determines the selection of undo, the strategy-directive would be, in words, "keep the goal and add a secondary goal to hinder the competitors' objectives". As another example, if the interference is "concord", the selected strategy may be "join efforts", with a strategy-directive such as "when developing a plan to reach the goal, try to merge it with the other agents' plans". Further possible strategy-directives are, among others, "drop goal", "replace goal", etc.

An example strategy-selection rule, selecting an undo strategy for competition, is sketched below:

employee E: competition/{E wants position P, some other employee X wants position P}, (E has an aggressive personality, E believes X is technically qualified for the position, E reproaches X for past misdeeds with intensity $\geq n3$) $\rightarrow$ E keeps his goal and adds the goal of bringing X's past misdeeds to the attention of the section manager

After a final set of goals Γ' is eventually obtained, plans must still be adopted for their achievement, either taken from a library of typical plans or newly developed. It is during this later phase that strategy-directives such as "merge plans", which might have been recommended by a strategy-selection rule in case of concord, would be put to work. Additionally, this is the moment to analyze any negative interferences caused by the plans adopted and to find how they should be modified — here, again, the possibility of negotiation among agents should be considered. In fact, although software tools may help in the processes of revising goals and plans, a fully automatic implementation seems far beyond the state of the art.

And different attitudes can orient the processes. One may have in mind the creation of a *decision-support environment*, with methods and tools to help agents to accommodate negative interferences and take the maximum advantage from the positive interferences. In such environment, each agent should have means to gather all the information needed to understand the current situation and to anticipate the consequences of his own activity, as favoured or disturbed by the other agents' moves.

But, alternatively, one may want to run a free *simulation environment*, where the purpose is to find what possible futures would arise from different lines of actions that the agents might be expected to pursue, without intervention or advice from the system. In this latter environment, strategy-selection rules would still make sense, but would simply reflect the spontaneous preferences of each agent, perhaps as observed

from his past conduct, and may handle an interference in a less than optimal way. Furthermore, an agent could adopt any plan of his liking (interfering or not with other plans). Even plans with obstacles that might be detected beforehand would be admitted; thus, the execution of ultimately failing plans would proceed, and they would produce their consequences until a failure condition causes their interruption.

5 Concluding remarks

As an addition to the use of goal-inference rules for the design and simulation of the behaviour of agents, which we have formally defined in a temporal logic in previous work [CVF], and tested with the help of a prototype [CF], we introduced — as a more far-reaching research proposal:

- internal states and profiles of agents, so as to provide a more realistic cognitive and affective basis to formulate and apply the goal-inference rules.

- strategy-selection rules, also to be based on the cognitive and affective elements kept in internal states and profiles, in order to deal adequately with goal interferences in multi-goal/multi-agent environments.

Our proposal requires a seemingly arbitrary representation of certain attributes in numerical terms; we refer, in special, to utility values and measures for affective elements such as drives and emotions. The arbitrariness should be reduced by whatever previous analysis is possible, and then corrected and tuned through experimental validation with representative sample cases. In what regards the development of prototypes, we recall that Constraint Programming features [KKR] offer an invaluable help to handle numerically expressed pre-conditions and effects in plan generation.

The additions suggested still require considerable research effort and, as usually happens with any attempt to expand a theory, may be costlier and more time-consuming than simpler schemas presently used in practice for conceptual specification and simulation. Nevertheless, we argue that they may prove indispensable, in many real-life contexts, to effectively model and thereby predict the behaviour of agents.

References

[Ba]   J. Bates - The role of emotion in believable agents. *Communications of the ACM*, 7, 37 (1994) 122-125.

[Br]   C. Breazeal - A motivational system for regulating human-robot interaction. Proc. of the Fifteenth National Conference on Artificial Intelligence, Madison (1998) 54-61.

[Ca]   J. Carroll - The deep structure of literary representations. *Evolution and Human Behavior* 20 (1999) 159-173.

[CF]   A. E. M. Ciarlini and A. L. Furtado - Understanding and Simulating Narratives in the Context of Information Systems. Proc. of the 21th International Conference on Conceptual Modeling (2002).

[CL]   P. R. Cohen, H. J. Levesque - Intention is Choice with Commitment. *Artificial Intelligence* 42, 2-3 (1990) 213-261.

[CP]   P. R. Cohen and C. R. Perrault - Elements of a plan-based theory of speech acts. In *Readings in Natural Language Processing*. B. J. Grosz, K. S. Jones and B. L. Webber (eds.). Morgan Kaufmann (1986) 423-440.

[CVF]   A. E. M. Ciarlini, P. A. S. Veloso and A. L. Furtado - A Formal Framework for Modelling at the Behavioural Level. Proc. of the European-Japanese Conference on Information Modeling and Knowledge Bases (2000) 67-81.

[CW]   M. Carlsson and J. Widen - *Sicstus Prolog Users Manual, Release 3.0.* Swedish Institute of Computer Science (1995).

[DAB] G. Davenport, S. Agamanolis, B. Barry, B. Bradley and K. Brooks - Synergistic storyscapes and constructionist cinematic sharing. *IBM System Journal*, 39, 3 & 4 (2000) 456-469.

[DLF]   A. Dardenne, A. van Lamsweerde and S. Fickas: Goal-directed requirements acquisition. *Science of Computer Programming* 20, 1-2 (1993) 3-50.

[DBH]   G. Donato, M. S. Bartlett, J. C. Hager, P. Ekman and T. J. Sejnowski - Classifying facial actions. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 21, 10 (1999) 974-989.

[DLF]   A. Dardenne, A. v. Lamsweerde and S. Fickas - Goal-directed requirements acquisition. *Science of Computer Programming* 20 (1993) 3-50.

[El]   C. Elliott: Using the Affective Reasoner to Support Social Simulations. Proc. of the 13th International Joint Conference on Artificial Intelligence (1993) 194-201.

[FC]   A. L. Furtado and A. E. M. Ciarlini: Constructing libraries of typical plans. Proc. of the 13th International Conference on Advanced Information Systems Engineering (CAiSE), Interlaken (2001) 124-139.

[FFM]   T. Finin, R. Fritzson, D. McKay and R. McEntire - KQML as an agent communication language. Proc. of the Third International Conference on Information and Knowledge Management (1994).

[KKR] P. C. Kanelakis, G. M. Kuper and P. Z. Revesz - Constraint query languages. Journal of Computer and System Sciences 51 (1995) 26-52.

[KW]   S. Kaiser and T. Wehrle - Emotion research and AI: some theoretical and technical issues. *Geneva Studies in Emotion and Communication* 8, 2 (1994).1-16.

[LL]   H. J. Levesque and G. Lakemeyer - *The logic of knowledge bases*. MIT Press (2000).

[Ma]   G. Mantovani - Social context in HCL: a new framework for mental models, cooperation and communication. *Cognitive Science* 20, 2 (1996) 237-269.

[MCY]   J. Mylopoulos, L. Chung and E. Yu - From object-oriented to goal-oriented requirements analysis. Communications of the ACM 42, 1 (1999) 31-37.

[Mo]   G. Morgan - *Images of Organization: the Executive Edition*. Berrett-Koehler Publishers, 1998.

[OCF]   A. Ortony, G. L. Clore and M. A. Foss - The referential structure of the affective lexicon. *Cognitive Science* 11, 3 (1987) 341-364.

[OO]   P. O'Rorke and A. Ortony - Explaining Emotions. *Cognitive Science* 18, 2 (1994) 283-323.

[PBS]   C. Pelachaud, N. I. Badler and M. Steedman - Generating facial expressions for speech. *Cognitive Science* 20 (1996) 1-46.

[PLL]   P. Persson, J. Laaksolahti and P. Lönnqvist - Proc. of Stereotyping characters: a way of triggering anthropomorphism? Socially Intelligent Agents - The Human in the Loop - AAAI Fall Symposium (2000).

[PVH]   R. W. Picard, E. Vyzas and J. Healey - Toward machine emotional intelligence: analysis of affective physiological state. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 23, 10 (2001) 1175-1191.

[RG]   A. S. Rao and M. P. Georgeff - Modeling rational agents within a BDI-architecture. Proc. of the International Conference on Principles of Knowledge Representation and Reasoning (1991) 473-484.

[RH]   D. Rousseau, B. Hayes-Roth - A social-psychological model for synthetic actors. Knowledge Systems Laboratory of the Department of Computer Science, Report KSL 97-07. Stanford University (1997).

[Ri]   E. Rich - User modeling via stereotypes. *Cognitive Science* 3 (1979) 329-354.

[RN]   S. Russell and P. Norvig - *Artificial Intelligence - a Modern Approach*. Prentice-Hall (1995).

[Sg]   N. M. Sgouros - Dynamic generation, management and resolution of interactive plots. *Artificial Intelligence* 107 (1999) 29-62.

[Si]   Special issue on digital entertainment - Scientific American, 283, 5, November (2000).

[SL]   A. Sloman and B. Logan - Architectures and tools for human-like agents. *Communications of the ACM* 42, 3 (1999) 71-77.

[TK]   A. Tversky and D. Kahneman - Extensional versus intuitive reasoning: the conjunction fallacy in probability judgement. Psychological Review, 90 (1983) 293-315.

[Ve]   J. D. Velásquez - Modeling emotions and other motivations in synthetic agents. Proc. of the Fourteenth National Conference on Artificial Intelligence, Providence (1997) 10-15.

[Wi]   R. Willensky - *Planning and Understanding - a Computational Approach to Human Reasoning*. Addison-Wesley (1983).

[WJ]   M. Wooldridge and N. R. Jennings - Intelligent agents: theory and practice. *Knowledge Engineering Review* 10, 2 (1995).

[Wo]   M. Wooldridge - *Reasoning about Rational Agents*. MIT Press (2000).