

A Normative and Intentional Agent Model for Organisation Modelling

Joaquim Filipe

*Escola Superior de Tecnologia do Instituto Politécnico de Setúbal
Rua Vale de Chaves, Estefanilha, 2910 Setúbal, Portugal.
JFilipe@est.ips.pt*

Abstract. This paper proposes a new agent model, including normative aspects based on studies from social psychology. The three main model components are the Epistemic, Deontic and Axiologic components. This model structure facilitates the representation of normative organisational knowledge and processes as well as the relations between individual and social interests. Instead of manipulating the notions of belief, intention, desire, commitment, goal and obligation, the only primitives needed are belief and generalized goal.

1 Introduction

Organisations are multi-agent systems, eventually including both human and artificial agents. Organisations are also seen as multilayered Information Systems (IS) themselves, including an informal subsystem, a formal subsystem and a technical system as shown in figure 1, denoting the organizational semiotics perspective on information systems.

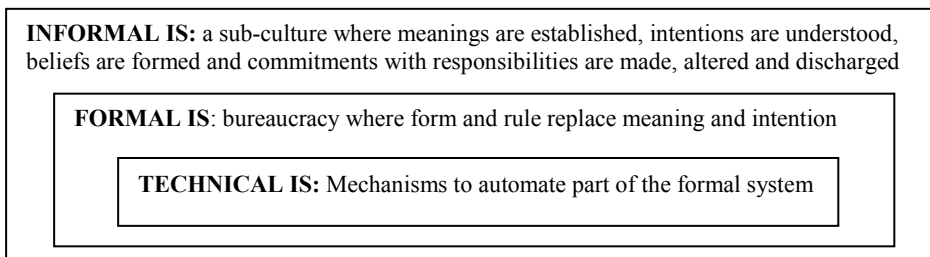


Figure 1: Three main layers of the real information system (Stamper 1996)

We aim at improving the technical subsystem within the constraints defined by the other two. Organisational information systems are inherently distributed, nowadays, thus communication and coordination are major problems in this kind of information systems. Perhaps motivated by the difficult problems there is currently a strong interest on this area, which is an active research field for several disciplines, including Distributed Artificial Intelligence (DAI), Organisational Semiotics and the Language-

Action Perspective, among others. Our approach integrates elements from these three perspectives.

The Epistemic-Deontic-Axiologic (EDA) designation refers to the three main components of the agent structure described in this paper. Here we propose an agent model which, contrary to most DAI proposals, not only accounts for intentionality but is also prepared for social interaction in a multi-agent setting. Existing agent models emphasise an intentional notion of agency – the supposition that agents should be understood primarily in terms of mental concepts such as beliefs, desires and intentions. The BDI model (Rao and Georgeff, 1991) is a paradigm of this kind of agent, although some attempts to include social aspects have been made (Dignum et al., 2000).

We claim that the cognitive notions that constitute the basis of intentional models show only one face of the coin, the other one being social agency. It is required that an adequate agent model, able to function in a multi-agent setting, should emphasise social agency and co-ordination, within a semiotics framework. Our approach focuses on organisational agents who participate in organisational processes involving the creation and exchange of signs, *i.e.* knowledge sharing, in a way that is inherently public, thus depending on the agent's social context, *i.e.* its information field (Stamper, 1996), to support co-ordinated behaviour. An information field is composed by all the agents and objects involved in a certain activity and sharing a common ontology.

A realistic social model must be normative: both human agents and correctly designed artificial agents ought to comply with the different kinds of norms that define the information field where they operate, although exceptions may occur. Private representations of this shared, normative, knowledge are translations of the public knowledge into specific agent mental structures. When designing an artificial agent, the designer must adopt a private knowledge representation paradigm to set up the agent's internal knowledge in such a way that it fits the normative shared ontology.

We postulate that norms are the basic building blocks upon which it is possible to build co-ordination among organised entities and co-ordinated actions are the crux of organised behaviour. We claim that although organised activity is possible either with or without communication, it is not possible without a shared set of norms. Therefore, these socially shared norms define an information field that is a necessary condition for heterogeneous multi-agent co-ordination including both artificial agents and humans.

2 The Normative Structure of the EDA model

Norms are typically social phenomena. This is not only because they stem from the existence of some community but also because norms are multi-agent objects (Conte and Castelfranchi, 1995):

- They concern more than one individual (the information field involves a community)
- They express someone's desires and assign tasks to someone else

- Norms may be regarded from different points of view, deriving from their social circulation: in each case a norm plays a different cognitive role, be it a simple belief, a goal, a value, or something else.

Social psychology provides a well-known classification of norms, partitioning them into perceptual, evaluative, cognitive and behavioural norms. These four types of norms are associated with four distinct attitudes, respectively (Stamper, 1996):

- Ontological – to acknowledge the existence of something;
- Axiologic – to be disposed in favour or against something in value terms;
- Epistemic – to adopt a degree of belief or disbelief;
- Deontic – to be disposed to act in some way.

An EDA agent is a knowledge-based system whose knowledge base structure, depicted in figure 2, is based on three components: the Epistemic, the Deontic and the Axiologic.




Figure 2: The EDA Agent Knowledge-Base Structure

The epistemic model component is where the knowledge of the agent is stored, in the form of statements that are accepted by that agent. Two types of knowledge are stored here: declarative knowledge – statements relative to the agent beliefs – and procedural knowledge – statements concerning the know-how of the agent, *e.g.* their plans and procedural abilities.

The importance of norms to action has determined the name we have chosen for the model component where the agent goals are represented. An agent goal may be simply understood as the *desire to perform an action* (which would motivate the designation of *conative*) but it can also be understood, especially in a social context, as the *result of the internalisation of a duty or social obligation* (which would motivate the designation of *deontic*). We have adopted the latter designation not only because we want to emphasise the importance of social obligations but also because personal desires can be seen as a form of ‘generalised’ obligation established by an agent for himself. This provides a *unification* of social and individual drives for action, which simplifies many aspects of the model.

The axiologic model component contains the value system of the agent, namely a partial order that defines the agent preferences with respect to norms¹. The importance of the agent value system is apparent in situations of conflict, when it is necessary to violate a norm. This preference ordering is dynamic, in the sense that it may change whenever the other internal components of the agent model change, reflecting different beliefs or different goals.

¹ It is important to note that these *values* are different from the values used in the evaluation process required to make a rational decision. The latter can be done in the scope of BDI models as illustrated in (Antunes and Coelho, 1999-1&2).

3 Intentions and Social Norms in the EDA Model

The multi-agent system metaphor that we have adopted for modelling organisations implies that organisations are seen as goal-governed collective agents, which are composed of individual agents. This perspective comes in line with the principles of normative agents proposed in (Castelfranchi et al., 1999).

In our model individual agents are autonomous, heterogeneous, rational, social agents. Therefore, they are compelled to make decisions and act in a way that, although not entirely deterministic, is constrained by their rationality. Actually their behaviour would be predictable if we knew all the details of their EDA model components, the environment stimuli, their perception function and also the reasoning machine they use, because the ultimate goal of a rational agent is to maximize its utility.

Typically, Artificial Intelligence (AI) agent models consider intelligent agent decision processes as internal processes that occur in the mind and involve exclusively logical reasoning, external inputs being essentially data that are perceived directly by the agent. This perspective does not acknowledge any social environment whatsoever. In this thesis we start from a totally different perspective, by emphasising the importance of social influences and a shared ontology on the agent decision processes, which then determines agents' activity. We shall henceforth refer to agent as 'it' although the EDA model also applies to human agents. In any case, we are particularly interested in the situations where information systems are formally described, thus making it possible for artificial agents to assist or replace human agents.

An important role of norms in agent decision processes is the so-called cognitive economy: by following norms the agent does not have to follow long reasoning chains to calculate utilities – it just needs to follow the norms.

However, instead of adopting a whole-hearted social sciences perspective, which is often concerned merely with a macro perspective and a statistical view of social activity, we have adopted an intermediate perspective, where social notions are introduced to complement the individualistic traditional AI decision models: a psycho-social perspective, whereby an agent is endowed with the capability of overriding social norms by intentionally deciding so.

Our model enables the relationship between socially shared beliefs with agent individual, private, beliefs; it also enables the analysis of the mutual relationships between moral values at the social level with ethical values at the individual level. However, we have found particularly interesting analogies in the deontic component, specifically in the nature of the entities and processes that are involved in agent goal-directed behaviour, by inspecting and comparing both the social processes and individual processes enacted in the deontic component of the EDA model.

This was motivated by the close relationship between deontic concepts and agency concepts, and represents a direction of research that studies agency in terms of normative social concepts: obligations, responsibilities, commitments and duties (. These concepts, together with the concepts of power/influence, contribute to facilitate the creation of organisational models, and are compatible with a vision of

organisations as normative information systems as well as with the notion of information field that underlies the organisational semiotics approach.

As will be explained in more detail in the next section, an essential aspect of the EDA model is that the Deontic component is based on the notion of *generalised goal* as a kind of obligation, that encompasses both social goals (social obligations) and individual goals (self obligations). This idea is inspired by Shoham's work (1993). Following a traditional designation in DAI, we designate those individual *generalised goals* that are inserted in the agenda as *achievement goals*, as in (Cohen and Levesque, 1990). Figure 3 describes the parallelism between mental and social constructs that lead to setting a goal in the agenda, and which justifies the adoption of the aforementioned generalised obligation. Here, p represents a proposition (world state). $B_\alpha(p)$ represents p as one of agent α 's beliefs. $O_\alpha^\beta(p)$ represents the obligation that α must see to it that p is true for β . $O_\alpha^\alpha(p)$ represents the interest that α has on seeing to it that p is true for itself – a kind of self-imposed obligation. In this diagram $p \in E_\alpha(W, D)$ means, intuitively, that proposition p is one of the goals on α 's agenda.

Interest is one of the key notions that are represented in the EDA model, based on

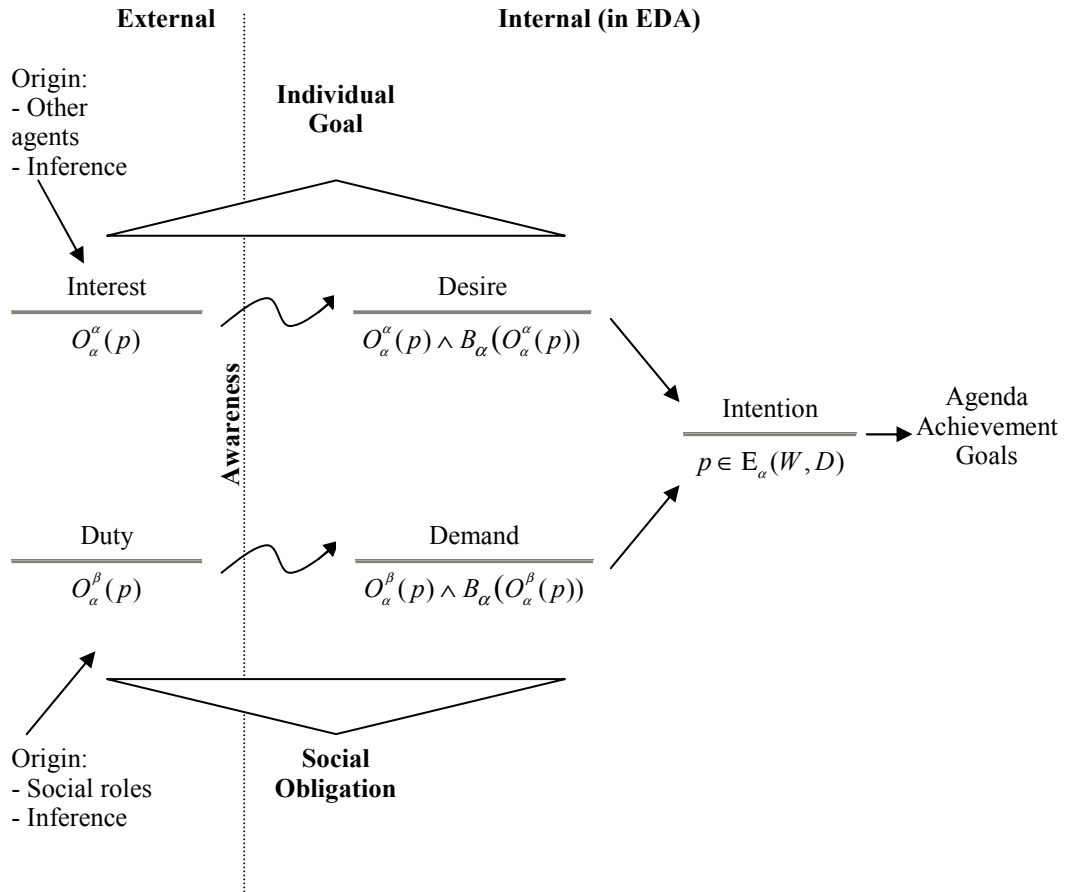


Figure 3: Social and Individual goals parallelism in the EDA model.

the combination of the deontic operator ‘ought-to-be’ (von Wright, 1951) and the agentive ‘see-to-it-that’ *stit* operator (Belnap, 1991). *Interests* and *Desires* are manifestations of *Individual Goals*. The differences between them are the following:

- *Interests* are individual goals of which the agent is not necessarily aware, typically at a high abstraction level, which would contribute to improve its overall utility. Interests may be originated externally, by other agents’ suggestions, or internally, by inference: deductively (means-end analysis), inductively or abductively. One of the most difficult tasks for an agent is to become aware of its interest areas because there are too many potentially advantageous world states, making the full utility evaluation of each potential interest impossible, given the limited reasoning capacity of any agent.
- *Desires* are interests that the agent is aware of. However, they may not be achievable and may even conflict with other agent goals; the logical translation indicated in the figure, $O_{\alpha}^{\alpha}(p) \wedge B_{\alpha}(O_{\alpha}^{\alpha}(p))$, means that desires are goals that agent α ought to pursue for itself and that it is aware of. However, the agent has not yet decided to commit to it, in a global perspective, *i.e.* considering all other possibilities. In other words, desires become intentions only if they are part of the preferred extension of the normative agent EDA model (Filipe, 2000).

It is important to point out the strong connection between these deontic concepts and the axiologic component. All notions indicated in the figure should be interpreted from the agent perspective, *i.e.* values assigned to *interests* are determined by the agent. Eventually, external agents may consider some goal (*interest*) as having a positive value for the agent and yet the agent himself may decide otherwise. That is why *interests* are considered here to be the set of all goals to which the agent would assign a positive utility, but which it may not be aware of. In that case the responsibility for the *interest* remains on the external agent.

Not all interests become desires but all desires are agent interests. This may seem contradictory with a situation commonly seen in human societies of agents acting in *others’ best interests*, sometimes even against their desires: that’s what parents do for their children. However, this does not mean that the agent desires are not seen as positive by the agent; it only shows that the agent may have a deficient axiologic system (by its information field standards) and in that case the social group may give other agents the right to override that agent. In the case of artificial agents such a discrepancy would typically cause the agent to be banned from the information field (no access to social resources) and eventually repaired or discontinued by human supervisors, due to social pressure (*e.g.* software viruses).

In parallel with *Interests* and *Desires*, there are also social driving forces converging to influence individual achievement goals, but through a different path, based on the general notion of social obligation. Social obligations are the goals that the social group where the agent is situated require the agent to attain. These can also have different flavours in parallel to what we have described for individual goals.

- *Duties* are social goals that are attached to the particular roles that the agent is assigned to, whether the agent is aware that they exist or not. The statement $O_{\alpha}^{\beta}(p)$ means that agent α ought to do p on behalf of another agent β .

Agent β may be another individual agent or a collective agent, such as the society to which α belongs. Besides the obligations that are explicitly indicated in social roles, there are additional implicit obligations. These are inferred from conditional social norms and typically depend on circumstances. Additionally, all specific commitments that the agent may agree to enter also become duties; however, in this case, the agent is necessarily aware of them.

- *Demands* are duties that the agent is aware of². This notion is formalised by the following logical statement: $O_{\alpha}^{\beta}(p) \wedge B_{\alpha}(O_{\alpha}^{\beta}(p))$. Social demands motivate the agent to act but they may not be achievable and may even conflict with other agent duties; being autonomous, the agent may also decide that, according to circumstances, it is better not to fulfil a social demand and rather accept the corresponding sanction. Demands become intentions only if they are part of the preferred extension of the normative agent EDA model – see (Filipe, 2000 section 5.7) for details.
- *Intentions*: Whatever their origin (individual or social) intentions constitute a non-conflicting set of goals that are believed to offer the highest possible value for the concerned agent. Intentions are designated by some authors (Singh, 1990) as psychological commitments (to act). However, intentions may eventually (despite the agent sincerity) not actually be placed in the agenda, for several reasons:
 - They may be too abstract to become directly executed, thus requiring further means-end analysis and planning.
 - They may need to wait for their appropriate time of execution.
 - They may be overridden by higher priority intentions.
 - Required resources may not be ready.

The semantics of the prescriptive notions described above may be partially captured using set relationships as depicted in figure 4, below.

When an agent decides to act in order to fulfil an intention, an agenda item is created – we adopt the designation of *achievement goal*. Achievement goals are defined as in (Cohen and Levesque, 1990) as goals that are shared by individuals participating in a team that has a *joint persistent goal*. Following the terminology of (Cohen and Levesque, 1990) agent α has a *weak achievement goal*, relative to its motivation (which in our case corresponds to the origin and perceived utility of that goal), to bring about the joint persistent goal γ if either of the following is true:

- α does not yet believe that γ is true and has γ being eventually true as a goal (*i.e.* α has a normal achievement goal to bring about γ)

² According to the Concise Oxford Dictionary, *demand* is “an insistent and peremptory request, made as of right”. We believe this is the English word with the closest semantics to what we need.

- α believes that γ is true, will never be true or is irrelevant (utility below the motivation threshold), but has a goal that the status of γ be mutually believed by all team members.




Figure 4: Set-theoretic relationships among deontic prescriptive concepts.

However we do not adopt the notion of joint persistent goal for social coordination, as proposed by Cohen and Levesque (1990) because their approach has a number of shortcomings, not only theoretical but also related to the practical feasibility of their model, which are well documented in (Singh, 1996).

4 The EDA Model Internal Architecture

Using the social psychology taxonomy of norms, and based on the assumption that organisational agents' behaviour is determined by the evaluation of deontic norms, given the agent epistemic state, with axiological norms for solving eventual interest conflicts, we propose an intentional agent model, which is decomposed into three main components: the epistemic, the deontic and the axiologic. Additionally there are two external interfaces: an input (perceptual) interface, through which the agent receives and pragmatically interprets messages from the environment and an output (acting) external interface through which the agent acts upon the environment, namely sending messages to other agents³.

A socially shared ontology is partially incorporated in an agent cognitive model whenever it is needed, *i.e.* when the agent needs to perform a particular role. In this case, beliefs are incorporated in the Epistemic component, obligations and responsibilities are incorporated in the Deontic component and values (using a partial order relation of importance) are incorporated in the Axiologic component – all indexed to the particular role that the agent is to play.

³ In this thesis we restrict our attention to the semiotic, symbolic, types of agent activity, ignoring substantive physical activities.

Figure 5 depicts the EDA model and its component relationships.

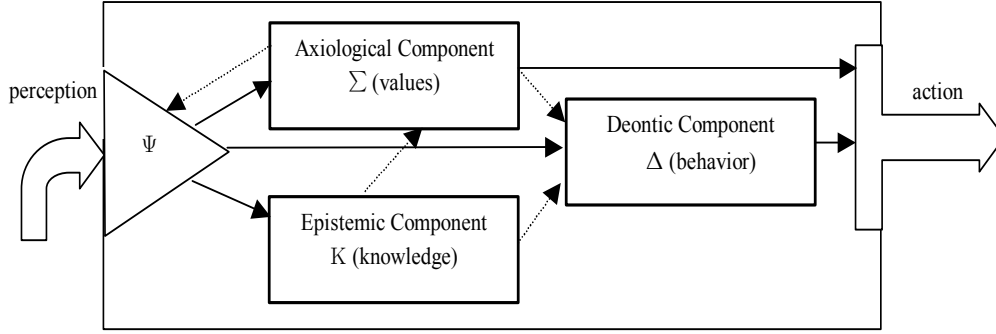


Figure 5: The EDA model component relationships.

Ψ is a pragmatic function that filters perceptions, according to the agent ontology, using perceptual and axiologic norms, and updates one or more model components.

Σ is an axiologic function that is used mainly in two circumstances: to help decide which signs to perceive, and to help decide which goals to put in the agenda and execute.

K is a knowledge-based component, where the agent stores its beliefs both explicitly and implicitly, in the form of potential deductions based on logical reasoning.

Δ is a set of plans, either explicit or implicit, the agent is interested in and may choose to execute.

The detailed description of each component, including its internal structure, is provided in (Filipe, 2000). In this paper we focus on the system behaviour. The next sections show how in EDA we specify ideal patterns of behaviour and also how we represent and deal with non-ideal behaviours.

5 Organisational Modelling with Multi-Agent Systems using the EDA Model

The EDA model may apply to both human and artificial agents, and is concerned with the social nature of organisational agents:

- Firstly, because it accounts for a particular mental structure (Epistemic-Deontic-Axiologic) that is better, for our purposes, than other agent mental structures proposed in the literature to model agent interaction. Specifically, we intend to use it for modelling information fields where social norms influence individual agents and are used by an agent to guide inter-subjective communication and achieve multi-agent co-ordination.

- Secondly, because the model is based on normative notions that are not only intended to guide the agent behaviour in a socially accepted way, but also to identify what sanctions to expect from norm violations, in order to let the agent take decisions about its goals and actions, especially when co-ordination is involved. The EDA model is based on the claim that multi-agent notions such as social commitment, joint intentions, teamwork, negotiation and social roles, would be merely metaphorical if their normative character were not accounted for.

Given its social-focused nature, the EDA agent notion may be used to model and implement social activities, involving multi-agent co-ordination. However, although the agent paradigm described in this thesis is suited to model team work and joint problem solving, the major novelty with respect to other current agent models is the normative flavour. EDA agents are able to co-ordinate on the basis of shared norms and social commitments. Shared norms are used both for planning and for reasoning about what is expected from other agents but, internally, EDA agents keep an intentional representation of their own goals, beliefs and values.

Co-ordination is based on commitments to execute requested services. Commitments are represented not as joint intentions based on mutual beliefs, as is the case of the Cohen-Levesque model, upon which the BDI paradigm is based, but as first-class social concepts, at an inter-subjective level. The organisational memory embedded in the representation of socially accepted best-practices or patterns of behaviour and the handling of sub-ideal situations is, in our opinion, one of the main contributions that a multi-agent system can bring about.

6 Representing Ideal and non-Ideal Patterns of Behaviour

The EDA model is a norm-based model. Norms are ultimately an external source of action control. This assumption is reflected specially in the Deontic component of the EDA model.

Standard Deontic Logic (SDL) represents and reasons about ideal situations only. However, although agent behaviour is guided by deontic guidelines, in reality an agent who always behaves in an ideal way is seldom seen. The need to overcome the limited expressiveness of SDL, and to provide a way to represent and manipulate sub-ideal states has been acknowledged and important work has been done in that direction, *e.g.* by Dignum *et al.* (1994).

Contrary to SDL, the Deontic component of the EDA model is designed to handle sub-ideal situations. Even in non-ideal worlds, where conflicting interests and obligations co-exist, we wish to be able to reason about agent interests and desirable world states in such a way that the agent still is able to function coherently.

Behaviours may be represented as partial plans at different abstract levels. A goal is a very high abstract plan, whereas a sequence of elementary actions defines a plan at the instance level. The Deontic component is similar, in this sense, to what Werner (1989) called the agent intentional state.

However, in our model, agent decisions depend both on the available plans and a preference relationship defined in the axiologic component. This value assignment,

which is essential for determining agent intentions, *i.e.* its preferred actions, can change dynamically, either due to external events (perception) or to internal events (inference), thus dynamically modifying the agent's intentions.

Although our representation of ideal behaviours is based in deontic logic, we acknowledge the existence of problems with deontic logic, partially caused by the fact that the modal 'ought' operator actually collapses two operators with different meanings, namely 'ought-to-do' and 'ought-to-be'. Our solution, inspired in (Hilpinen, 1971) and (Belnap, 1991), is to use a combination of action logic and deontic logic for representing agentive 'ought-to-do' statements, leaving the standard deontic operator for propositional, declarative, statements. Agentive statements are represented as $[\alpha \text{ stit} : Q]$ where α stands for an agent and Q stands for any kind of sentence (declarative or agentive). An 'ought-to-do' is represented using the conventional 'ought-to-be' modal operator combined with an agentive statement, yielding statements of the form $O[\alpha \text{ stit} : Q]$ or, for short: $O_\alpha(Q)$.

The representation of behaviours using this kind of agentive statement has several attractive properties, including the fact that it is a declarative representation (with all the flexibility that it provides) and that there exists the possibility of nesting plans as nested *stits*.

A plan is typically given to the agent at a very abstract level, by specifying the goal that it ought to achieve. The agent should then be able to decompose it into simple, executable, actions. This decomposition can be achieved by a means-ends process. By representing plans declaratively, as behavioural norms, the process becomes similar to backward chaining reasoning from abstract goals to more specific ones, until executable tasks are identified, the same way as in goal oriented reasoning in knowledge-based systems. This similarity enables the adoption of methods and tools from that area, being especially useful the inference engine concept.

However, the Deontic component does not control the Agenda directly, *i.e.* it is not responsible for setting the agent goals directly, because the prospective agent goals – similar to *desires*, in the BDI model – must be analysed by the Axiologic component first, which computes their value accordingly to an internal preference relation, taking into account possible obligation violations.

7 Generalised Goals

Goal-governed behaviour can be represented, according to agency logic (Belnap, 1991), by an agentive statement where an agent α sees to it that the state of affairs Q obtains at time τ_2 by the same agent α performing a choice at time τ_1 : $[\alpha \text{ stit} : Q]$. We claim that it is possible to use the same syntax and semantics both for representing goals derived from both agent interests and behavioural social norms. This would make it possible to build up detailed plans from agent interests and behavioural norms using means-ends analysis – the same inference mechanism commonly used in knowledge-based systems for goal oriented backward-chaining.

We propose an unification of different concepts under the notion of generalised goal. This unification stems from the following axioms:

- All actions have a responsible agent, who performs the action.
- The performer may not be the ultimate responsible for the action and, in the case of artificial agents, it never is: the last agent in the responsibility chain must be a person.
- An action is performed for satisfying someone's interest.
- If the client for whom the action is performed is not satisfied he can ask for a sanction to be imposed by the information field that defines the context of the action.

We propose the following formal representation for generalised goals:

$$G_i = O_{\alpha/\rho}^{\beta/\chi}(Q, \tau, \sigma) \supset O([\alpha \text{ stit} : Q] \text{ in-time-window } \tau \text{ subject-to-sanction } \sigma)$$

This can be read as: α (under the responsibility of ρ) has the goal of ensuring Q , in reply to a request of β (under the control of χ); α is the performing agent, β is the client that requested the service, χ is a controller agent – which Castelfranchi (1993) calls a witness, Singh (1996) calls a context, and we associate to the information field. ρ is the responsibility chain for α ⁴; O is the standard deontic operator ‘ought-to-be’; $[\alpha \text{ stit} : Q]$ is an agency statement, saying that agent α sees-to-it-that proposition Q becomes true. This means that α will perform a plan to bring about Q in time window τ , where τ is a time expression specifying the time window during which proposition Q is intended to be satisfied; τ may be specified in absolute time or relative to some event; σ indicates the sanction cost of violation.

If agent α is different from agent β then we have a social commitment; otherwise we have an intention. In the latter case there is no controller agent χ and the sanction σ is typically null (represented by \emptyset) with expected value zero.

8 Conclusions

The EDA model is a norm-based, theroretically sound, agent model that takes into account not only the intentional aspects of agency but also the social norms that prescribe and proscribe certain agent patterns of behaviour. The main components of this model (Epistemic, Deontic and Axiological) have a direct relationship with the types of norms that are proposed in the social psychology theory supporting the model. In this paper, however, we focused our attention essentially in the Deontic component, where the normative social aspects are more important, namely where ideal and sub-ideal behaviours are represented.

We consider agents to be goal-governed systems. Using the *generalised goal* concept, all agent goals can be represented as obligations, encompassing both agent self-imposed obligations and social obligations – derived from moral obligations or commitments established in the course of their social activity. The proposed modal operator – G_i – overcomes the excessive strength of *stit*, and implicitly permits the

⁴ The responsibility chain is a list constituted by one or more agents that are responsible for the commitment to service performance and may be brought to scene if there is a breakdown during the service performance or if the performing agent is unable to handle the service satisfactorily.

representation of an important organisational activity: delegation, which Belnap's *stit* does not allow. Furthermore, the inclusion of a '*committed-to*' agent also extends *stit* and permits the bringing in of the notion of social commitment, based on the notion of obligation towards other agents and subject to sanctions.

Instead of manipulating the notions of belief, intention, desire, commitment, goal and obligation, the only primitives that we need are *belief* and *generalised goal*. All other notions are constructed from these ones. The interested reader can find more details about this in (Filipe, 2000).

References

1. Antunes, L. and H. Coelho, 1999. Rationalising decisions using multiple values, *Proceedings of the European Conference on Cognitive Science*, Siena, Italy.
2. Antunes, L. and H. Coelho, 1999. Decisions based upon multiple values: the BVG agent architecture. *Proceedings of the Portuguese AI Conference EPIA99*, Springer-Verlag LNAI (Lecture Notes in Artificial Intelligence) n° 1695, pp. 297-311, Évora, Portugal.
3. Belnap, N., 1991. Backwards and Forwards in the Modal Logic of Agency. *Philosophy and Phenomenological Research*, vol. 51.
4. Castelfranchi, C., 1993. Commitments: from Individual Intentions to Groups and Organisations. In *Working Notes of AAAI'93 Workshop on AI and Theory of Groups and Organisations: Conceptual and Empirical Research*, pp.35-41.
5. C. Castelfranchi, F. Dignum, C. Jonker and J. Treur, 1999. Deliberate Normative Agents: Principles and Architectures. In N. Jennings and Y. Lesperance (eds.) *Proceedings of ATAL-99*, Orlando, USA.
6. Cohen, P. and H. Levesque, 1990. Intention is Choice with Commitment. *Artificial Intelligence*, **42**:213-261.
7. Conte, R., and C. Castelfranchi, 1995. *Cognitive and Social Action*, UCL Press, London.
8. Dignum, F., J-J Ch. Meyer and R. Wieringa, 1994. A Dynamic Logic for Reasoning About Sub-Ideal States. In J. Breuker (Ed.), *ECAI Workshop on Artificial Normative Reasoning*, pp.79-92, Amsterdam, The Netherlands.
9. F. Dignum, D. Morley, E. Sonenberg, and L. Cavendon. Towards socially sophisticated BDI agents. In *Proceedings of the Fourth International Conference on MultiAgent Systems, ICMAS-2000*, Boston, USA.
10. Filipe, J., 2000. Normative Organisational Modelling using Intelligent Multi-Agent Systems. Ph.D. thesis, University of Staffordshire, UK.
11. Hilpinen, R., 1971. *Deontic Logic: Introductory and Systematic Readings*. Reidel Publishing Co., Dordrecht, The Netherlands.
12. Rao, A. and M. Georgeff, 1991. Modeling Rational Agents within a BDI architecture. *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning* (Nebel, Rich, and Swartout, Eds.) pp.439-449. Morgan Kaufman, San Mateo, USA.

13. Shoham, Y., 1993. Agent-Oriented Programming. *Artificial Intelligence*, **60**, pp.51-92, Elsevier Science Publishers.
14. Stamper, R., 1996. Signs, Information, Norms and Systems. In Holmqvist et al. (Eds.), *Signs of Work, Semiosis and Information Processing in Organizations*, Walter de Gruyter, Berlin, New York.
15. Singh, M., 1990. *Towards a Theory of Situated Know-How*. 9th European Conference on Artificial Intelligence.
16. Singh, M., 1996. Multiagent Systems as Spheres of Commitment. *Proceedings of the International Conference on Multiagent Systems (ICMAS) - Workshop on Norms, Obligations and Conventions*. Kyoto, Japan.
17. Werner, E., 1989. Cooperating Agents: A Unified Theory of Communication and Social Structure. In Gasser and Huhns (Eds.), *Distributed Artificial Intelligence*, pp.3-36, Morgan Kaufman, San Mateo, USA.
18. von Wright, G., 1951. Deontic Logic. *Mind*, **60**, pp.1-15.