THE EDA MODEL: AN ORGANIZATIONAL SEMIOTICS PERSPECTIVE TO NORM-BASED AGENT DESIGN

Joaquim Filipe

Escola Superior de Tecnologia do Instituto Politécnico de Setúbal Rua Vale de Chaves, Estefanilha, 2914 Setúbal, Portugal. JFilipe@est.ips.pt

Kecheng Liu

School of Computing, Staffordshire University Stafford, ST18 0AD, United Kingdom K.Liu@staffs.ac.uk

Keywords: Norms, Organizations, Semiotics, Information Systems, Social Simulation.

Abstract

Semiotics is the formal doctrine of signs. Organizational Semiotics is a particular branch of Semiotics, concerned with understanding organizations as information systems. Information Systems' core concepts such as information and communication are very complex and ill-defined concepts, which should be analyzed in terms of more elementary notions such as semiotic signs. Business processes would then be seen as processes involving the creation, exchange and use of signs. Rejecting the existence of a totally objective reality, we adopt as our philosophical stance a radical relativistic model. In this model both human agents and artificial agents have a constructed knowledge about reality, which requires the agent active participation, and all knowledge is connected to a knowing agent. In this view, organizational activity becomes an information process based on the notion of responsible agent, and we propose a model that accommodates both the social dimension in organizational agents behavior and the relative autonomy that individual agents exhibit in organizations.

The EDA (Epistemic-Deontic-Axiological) model, here proposed, enables the representation of agent informational states and simultaneously defines the conceptual communication framework. Agents use their knowledge (epistemic level) and take into account their obligations and authorizations (deontic level), which they may choose to accept or to violate, to decide what to do next, i.e. to define their goals. In the process they use individual preferences defined in their system of values (axiological level). Organizational concepts and activities, such as power relationships, roles, or contracts, are defined by norms in terms of the basic EDA components. Using an EDA model, it is possible to define an explicit representation of the institutional roles the agent can play, where a role is defined as a set of services plus a set of policies. A service is represented by a procedural abstraction, whereas a policy is represented by a deontic statement, either specifying an obligation-to-do.

We adopt the view that agents try to maximize their individual utility within a context of bounded rationality. Given the well-known theoretical and practical limitations of utility functions, we propose a logical framework, based on prioritized default reasoning, to represent norms and norm-based decision processes.

Organizational agents co-operate by exchanging services and respecting their commitments and the policies that are defined in their roles. However, each agent is autonomous in the sense that, in order to be able to cope with non-programmed situations, it may choose to violate its obligations, depending on the axiological component of its model.

We suggest that this social model can be used both to analyze an organization and to guide the design of mixed organizations where human and artificial agents co-operate. In a multi-agent environment, the perlocutionary effects of agent communicative actions, or speech acts, consist of EDA models update using agent specific pragmatic functions, i.e. each agent updates its own EDA components, by pragmatically interpreting received messages. The details of these pragmatic functions will be reported elsewhere.

1 INTRODUCTION

Organizations are complex, dynamic, non-linear, adaptive, evolving systems. Since organizational behavior results from interactions among a variety of heterogeneous (human and artificial) agents, organizations are poor candidates for analytical models. Organization Theory has produced many models of various types of organizations (Scott, 1992; Simon, 1996; Mintzberg, 1979) but even the most elaborate formal models are too simple to account for the complexity of human organizations.

A possible approach to this type of problems, which sometimes achieves good results, is the usage of system simulation techniques. However, the immediate application of simulation methods to organizational processes faces some complexity unless a pre-simulation analysis is conducted (Barjis and Filipe, 2000). This difficulty is compounded when the simulation is based on human factors and requires the interaction of organizational agents. The application domain must be very well understood before attempting to create a simulation model. This entails a requirements analysis phase, which is actually a preliminary phase of any software engineering project. Many methods have been proposed to tackle this important problem (Ghezzi, 1991) but, in spite of that, a high proportion of computer-based systems fail. Estimates vary between 40% and 50% of projects (Stamper et al., 2000), which seems to indicate that the requirements specifications are often wrong.

In this paper we approach this problem using the Organizational Semiotics stance (Stamper, 1973; Liu, 2000) to provide adequate system requirements and a solid conceptual basis for simulation models. Semiotics (Peirce, 1960), which was traditionally divided into three areas – syntax, semantics and pragmatics – has been extended by Stamper (1973) in order to incorporate three other levels, including a social world level. The type of information signs studied in each level is informally and briefly described in figure 1. A detailed and formal account of these levels may be found in (Stamper, 1996).

					SOCIAL WORLD - Commitments,
Human Information					contracts, law, culture,
Functions				PRAGMATICS - intentions, communication,	
			conversations, negotiations,		
			SEMANTICS - meanings,		
			validity, truth, signification,		
The	IT	SYNTACTICS - formal structure, language, logic,			
Platform Data,		records, deduction, software,			
EMPIRICS - pattern, variety, noise, entropy,					
channel capacity, redundancy, efficiency, codes,					
PHYSICAL WORLD - signals, traces, physical distinctions,					
hardware, component density, speed, economics,					

Figure 1: The semiotics framework (Stamper 1973)

The social world level, especially in what concerns social norms and commitments, is particularly relevant for both information systems analysis and design, because organizational activity is mainly the result of the coordinated behavior of several organizational units/agents.

The recent paradigm shift from centralized data processing architectures to heterogeneous distributed computing architectures, emerging especially since the 1990's, placed social concerns in the agenda of much research activity in computing, particularly in the Distributed Artificial Intelligence field (DAI). In DAI, organizations are modeled as multi-agent systems composed by autonomous agents acting in a coordinated way in order to achieve social goals, in a cooperative manner (Wooldridge and Jennings, 1995; Singh, 1996). Considering artificial agents, we defined 'autonomy' as the characteristic of agents that do not accept external requests unconditionally: they have a choice whether or not to act, and which action to perform. This is one of the major differences to objects, which can be considered the previous evolutionary step in software engineering.

Organizational agents, whether human or artificial, rely essentially on two mechanisms for social coordination: norms and communication.

In recent years, the communication-oriented perspective on organizations has received considerable academic attention (Barbuceanu and Fox, 1997; Weigand, et al., 1998; van Reijswoud and Dietz, 1999). In this perspective, organizations are social systems that are composed of socially acting actors and coordinated by performative (action-oriented) communication (Taylor, 1998; Winograd and Flores, 1986). The normative perspective to organizations has been the subject of research involving various areas such as deontic logic (Normatics, 1991; Dignum et al., 1996; Jones and Sergot, 1996; Santos and Carmo, 1996), social sciences (Conte and Castelfranchi, 1995) or multi-agent systems (Singh, 1991; 1997).

However, in this paper we focus especially on a norm-oriented perspective, adopting the social norm concept found in the semiotics philosophical stance, to propose a new information model (the EDA model) for agents and, based on this model, to simulate the social behavior of organizational multi-agent systems. An important point to note is that the concept of information we adopt is not part of a cognitive model – as considered by some authors (Werner, 1989) – but rather the reverse: the information model incorporates different types of norms, at different semiotic levels, including perceptual norms, evaluative norms, cognitive norms, and deontic norms (Stamper et al., 2000).

2 SEMANTIC ANALYSIS

We are interested to use the organizational semiotics stance as a bridge between organizations and computerbased systems (Filipe et al., 1997, 1999a), which are purely syntactic devices, as implied in the semiotics ladder in figure 1. Since organizational models cover human information functions levels this gap may be philosophically unbridgeable. However, we hypothesize that, using adequate abstractions, it is possible to produce syntactic devices that act *as if* they were able to perform those higher information functions in the semiotics ladder.

The first step in this direction, using a principled method, is to identify the right abstractions to use and their meanings. That is the methodology followed in Organizational Semiotics that approaches the requirements analysis of information systems by focusing on the semantics level first and then building up towards the social world level.

At the semantics level, the analyst's objective is to identify concepts and their relationships, and to represent meanings in a semantic network, where arcs represent ontological dependency and nodes represent (potential) beings. Two node types are particularly important in this network: agents and affordances.

An agent is responsible for controlling the existence of other entities using the following three operations: bring into existence, maintain and terminate (Liu, 2000). In this process, the agent is guided by organizational norms but, since the agent may choose to violate norms, he always retains autonomy. Without choice the concept of agency would be vacuous (Belnap and Perloff, 1991).

An affordance is an invariant that is perceived by an agent. This concept is borrowed from Gibson's work on perception that has lead to his 'theory of affordances' (Gibson, 1979). According to this theory, an agent is not a passive entity: he perceives by detecting invariants in an information-rich environment. The invariants the agent recognizes depend upon the structure of both agent and environment, emerging from the relationships between them. Gibson called these types of invariants 'affordances'. The role of knowledge, both as pre-condition and result of 'active perception', is also emphasized by Piaget.

2.1 Ontology Charts

Semantic analysis consists of a negotiated understanding, between relevant process intervenients, of the meanings of domain specific signs, including agents, concepts, relationships, and other related aspects. The result of semantic analysis may be provided in a graphical format, using what is called 'ontology chart', as shown in an example in figure 2.



Figure 2: Ontology Chart.

For example, the relationship between 'department' and 'school' is one of *ontological dependency* because it defines the existence of the first in terms of the existence of the second: 'school' is the ontological *antecedent* and 'department' is the ontological *dependent*. The dot sign (\bullet) means that the ontological dependency is that of part to whole.

The lines that link 'person' and 'works' (in both concepts of 'works' in the diagram) are labeled, using halfcircles. These labels denote the *roles* that the antecedent plays, in each case, as agent of the action 'works'.

The '@' sign, for example on the link between 'approves' and 'request', means that the consequent ('request') affordance is brought into existence conditionally to the validity of the antecedent ('approves'): the head of the department may or may not 'approve' the 'request' – only after s/he exercises his discretion and approves the request will the 'request' actually start its existence.

For a complete description of the syntax of the *ontology chart*, the reader may consult (Stamper, 1996). This graphical representation shows the existing entities and their ontological dependencies but not the starting and finishing times of each one of them. The dynamics of the organization, established by these start and finish events, may be determined either by norms or by agents taking responsibility and exercising discretion. Norms are a reflex of business rules, social goals, constraints and other structural aspects of the organization.

The essential advantage of ontology charts over, for example, Entity-Relationship (ER) models lies on that *ontological dependencies* shown only on ontology charts are less prone to change than entity relationships. Ontology charts provide a way to establish the semantics of very stable relationships in an organizational environment. Ades (1999) proposed a canonical form for representing these ontological relationships using a semantic normal form (SNF) based on the concept that any consequent may have at most two antecedents.

2.2 Norms, Agents and Responsibility

Meanings become clear when we know precisely *when* activities are performed. The ontology chart shows the ontological dependencies but not the actual existence time of each affordance or agent. The dynamics of the organization is described by start and finish events, which may be caused by: (i) automatically triggered organizational norms, or (ii) agents exercising their autonomy, although (ideally) within boundaries defined by the organizational norms. For example, in a library context, a norm may bring into existence a 'fine' relating a 'borrowed book' to a 'person', automatically if the book is not returned before a certain time. The librarian (agent) may exercise her/his/its agency to terminate the 'fine' affordance, assuming responsibility for that action.

Norms are a reflex of business rules, social goals, constraints and other structural aspects of the organization, being an essential part of an organizational agent's role (Filipe et al., 1999b; 2000) because they define all the regulated aspects that guide the agent's activities, including both obligations (to bring certain affordances into existence) and authorizations (to exercise agency, assuming responsibility for it).

Our approach views a business process as an activity network composed of autonomous agents. Agents represent individuals or collectives, including external stakeholders such as customers, regulators or suppliers, and internal entities such as staff, departments, or systems. In such a social environment, normative knowledge acquires a prominent role. Norms are needed to reduce the communicative overload that would occur if every action would require a process of negotiation before an agreement was reached.

Philosophically there is little doubt that responsibility must rest ultimately with a human agent. Only when artificial agents develop a conscience will they become socially and legally liable. However, it is common to have artificial agents performing normative tasks, *as if* they were responsible agents. For example, an ATM machine is delegated with the responsibility to deliver cash against the presentation of a proper plastic card and a correct personal identification code. However, if the machine makes a mistake it becomes clear that the responsibility ultimately lies with the bank, or its representative, who must accept the consequence of the mistake, according to certain social norms.

3 THE INFORMATION FIELD PARADIGM

Organizations are distributed information systems. They are composed by many agents that act together, in a responsible way, to achieve organizational goals. The multi-agent systems paradigm is thus an adequate computational metaphor to model an organization, as long as normative aspects are taken into account.

3.1 Organizational Structure

The Information Field paradigm views information systems in a social perspective (adequate for organizational modeling). The organizational model is a three-tiered information system – Informal-Formal-Technical – which Stamper (1996) coined as the 'organizational onion' (figure 3).

INFORMAL IS: a sub-culture where meanings are established, intentions are understood, beliefs are formed and commitments with responsibilities are made, altered and discharged
FORMAL IS: bureaucracy where form and rule replace meaning and intention
TECHNICAL IS: Mechanisms to automate part of the formal information system

Figure 3: Three main layers of the real information system (inspired in Stamper 1996).

This model shows that parts of the organization work without computer support, either informally or formally. Organizations have a structure: people get things done in an organized way in so far as their behavior exhibits regularities. However, organized behavior may or may not involve explicitly declared formal norms.

Organized behavior without written norms and based upon informal coordination is actually the main characteristic of an organizational structure that is designated by Mintzberg (1979) as "simple structure";

Stamper (1996) claims that this type of behavior always exists in every organization: it constitutes the 'informal information system'. However, when the size of the organization increases or the operations become too complex (e.g. requiring coordination between several specialized people), decision makers may face cognitive overloads, thus at this stage organizations typically become what Mintzberg defines as 'bureaucratic machines'. In Stamper (1996) the bureaucratic level is considered an extra layer of the general information system, which is denoted the 'formal information system'.

One of the advantages of bureaucracy is that routine information tasks are performed mechanically, by people who do not need to understand the precise meaning or purpose of the signs they are handling. This deskilling is also a serious disadvantage in the sense that seldom contributes to self-satisfaction of the organization members and it represents a waste of cognitive potential for the organization. Therefore, the existence of a Formal System has been usually a pre-condition for the automation of business processes and the existence of the organizational Technical Information System.

Information Technology has found its application field mainly in the automation of all kinds of repetitive semiotic tasks for which it was possible to find an algorithm. Computer-based systems embody norms typically in the form of software and provide a way to execute tasks more efficiently than human agents but still within a human normative framework. Agent technology provides a valuable, new, paradigm for automating tasks that were not possible to automate in the past because of distributed computing and/or communication-intensive requirements over a network of heterogeneous information systems. This technology may be used both in intra-organizational environments and inter-organizational environments, giving agents a level of autonomy that emulates some characteristics of human organizational agents and provides a degree of responsibility.

3.2 Agents in Multiple Information Fields

Different norm systems may co-exist within the same organization or across organizations. Typically an information field is shared by a number of agents that use the norms embedded in the information field to interact socially. However, if we consider the case of human organizations, it is easy to see that an agent may belong simultaneously to different information fields, playing different roles in each one, thus requiring some



Figure 4: Information Fields as Interacting Systems of Norms.

form of accommodating different, possibly conflicting, systems of norms (figure 4).

In order for an agent to act coherently in such an environment, it must first identify the information field that supports the role it is playing in each situation. This identification leads to the selection of an adequate system of norms that the agent will use to behave in a socially acceptable way.

Coordination of collaborative multi-agent systems is achieved both by normative behavior and by communicative behavior. When collaborative agents are playing roles in the same, shared, information field their normative behavior is fully compatible and mutually known, therefore there is no need for extensive communication. On the other hand, if the roles agents are playing belong to different information fields then mutual expectations (or default behaviors) may be quite different from actual behaviors. In this case, increased communication is required in order to obtain the desired coordination, with a corresponding increase in costs.

4 THE EDA MODEL

Social psychology provides a well-known classification of norms, partitioning them into perceptual, evaluative, cognitive and behavioral norms. These four types of norms are associated with four distinct attitudes, respectively (Stamper et al., 2000):

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

Our agent model is based on these attitudes and the associated norms, which we characterize in more detail below:

- Perceptual norms, guided by evaluative norms, determine what signs the agent chooses to perceive. Then, when a sign is perceived, a pragmatic function will update the agent EDA model components accordingly.
- Cognitive norms define entity structures, semantic values and cause-effect relationships, including both beliefs about the present state and expectations for the future. Conditional beliefs are typically represented by rules, which being normative allow for the existence of exceptions.
- Behavioral norms define what an agent is expected to do. These norms prescribe ideal behaviors as abstract plans to bring about ideal states of affairs, thus determining what an agent ought to do. Deontic logic is a modal logic that studies the formal properties of normative behaviors and states.
- Evaluative norms are required for an agent to choose its actions based on both epistemic and deontic attitudes. If we consider a rational agent, then the choice should be such that the agent will maximize some utility function, implicitly defined as the integral of the agent's axiological attitudes.

Using this taxonomy of norms, and based on the assumption that an organizational agent behavior is determined by the evaluation of deontic norms, given the agent epistemic state, we propose an intentional agent model, which is decomposed into three components: the epistemic, the deontic and the axiological.

Together, these components incorporate all the agent informational contents, according to the semiotics ladder depicted in figure 1, where it is shown that information is a complex concept, and requires different viewpoints to be completely analyzed.



Figure 5: The EDA agent model.

 \Box is a pragmatic function that filters perceptions, according to the agent perceptual and axiological norms, and updates one or more model components.

 \Box is an axiological function, that is used in two circumstances: to decide which signs to perceive and to decide which actions to execute.

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

 \Box is a set of available plans, either explicit or implicit, that the agent may choose to execute.

The main advantage of this model is that, since it is supported by a sound philosophical background, it avoids the *ad hoc* trap that undermines some current system architectures.

4.1 The Epistemic Component

In fact, since we adopt a radical subjectivist philosophical perspective, saying that all knowledge is connected to a knowing agent, instead of reflecting an objective reality, we are not interested in representing objective facts but rather in representing agents' beliefs. In this sense the most correct designation for this component would be *doxastic* instead of *epistemic*. However, following a more popular terminology, in the following we will use the term epistemic although with the reserves pointed out above.

The semiotics methods proposed in (Stamper et al., 1988) regarding requirements analysis and specification, state that the analysis process should start with a semantic analysis phase. The results of this phase can be displayed graphically as an ontology chart. However, since we are interested in trying to partially automate some of the organizational processes, we need a formal model.

A formal model of an organization must enable the representation for agents, affordances, and their ontological relationships. Furthermore, cognitive norms need to be included in the epistemic component of the agent informational model in order to provide an intensional form of knowledge representation.

Let $A = \{ \Box_1, \Box_2, ..., \Box_k \}$ be the set of agents and let $\Phi = \{ \Box_1, \Box_2, ..., \Box_m \}$ be the set of affordances, represented in the ontology chart; Let $P = \{ \Box_1, \Box_2, ..., \Box_n \}$ be the set of relationships between them. An affordance may depend ontologically on one or two antecedents, which can be agent(s) or affordance(s). Formally, using the BNF notation: $\Box_i = \Box_i(\Box_k | \Box_i, \Box_k | \Box_k | \Box_i, \Box_k)$.

Two main approaches to formal definition of knowledge (belief) have been proposed:

- The sentential approach: every agent knows every proposition that is stored in its knowledge base (Konolidge, 1986).
- The possible-worlds approach: an agent knows every proposition that is true in all worlds that are 'considered' possible (Hintikka, 1962), i.e. that are compatible with what it knows.

Each approach has advantages and disadvantages: the possible-worlds approach provides an elegant semantics but is based on the assumption that agents are perfect reasoners, which is not realistic, because agents would then know all logical consequences of their knowledge – the logical omniscience problem. The sentential approach does not have this problem because it is a syntactic approach: it does not assign semantic content to knowledge. Alternative approaches exist, e.g. (Singh and Asher, 1993), that seek to avoid both these problems but are technically more complex than the approaches mentioned. Therefore, for the sake of simplicity we will adopt the sentential approach in the following.

Since agents are situated in time, we need a time structure for our model. Let $T = \{\tau_1, \tau_2, ..., \tau_t\}$ be the set of time instants ordered by a relation p. Some authors take the temporal precedence relation to be linear in the past and assume time is branching into the future, which makes p a partial relation between moments along the same history (time tree branch). Temporal expressions are constructed using some variant of CTL* (Emerson, 1990) and the semantics is usually formalized using a possible-worlds approach. As indicated before, we decided to avoid the possible-worlds approach, thus we have adopted a linear-time relation and defined a language \Im_L that

allows two types of time expression \Im : either *relative* to an event \Im^* or *absolute* \Im^* . Time intervals $\Delta \Im$ are defined using a pair of two time expressions, to refer to all the time instants between the interval lower and upper limits. Using the BNF notation, a time expression \Im is defined as:

$$\mathfrak{I} = \mathfrak{I}^{\kappa} | \mathfrak{I}^{\Lambda} \text{ where: } \begin{cases} \mathfrak{I}^{\Lambda} = \tau_{i} \\ \mathfrak{I}^{\kappa} = \tau_{event} + \tau_{i} \end{cases} (\tau_{event} \text{ is the instant when the event occurs}) \end{cases}$$

and a time interval $\Delta \Im$ is defined as:

 $\Delta \mathfrak{I} = \begin{bmatrix} \mathfrak{I}_i & \mathfrak{I}_j \end{bmatrix}$, where $\mathfrak{I}_i \mathbf{p} \mathfrak{I}_j$

The Epistemic Component contains belief statements, cognitive norms and an incomplete¹ inference machine. Belief statements are represented by: B(A, Π , $\Delta\Im$), where A is an agent; Π is a proposition, expressed in terms of affordances and/or agents, instantiating a valid ontological pattern as determined in the ontology chart; and $\Delta\Im$ is a time interval that indicates the period of existence of the belief. For example, if A='I', Π ='at home' and $\Delta\Im$ ='today between 10:30 and 16:00', then the system would hold the following belief: B(I, at home, [today+10:30 today+16:00]), assuming that the time constant 'today' referred to 0:00h today.

Cognitive norms represent conditional beliefs: $\Sigma = B(A, \Pi, \Delta \Im)$ where Σ is a statement defined in any of the agent model components. For example, assuming that there is a deontic norm $O^{\theta}_{\alpha}(P, \Delta \Im, \sigma)$ meaning that agent α ought to bring about P in $\Delta \Im$, subject to accepting a sanction σ from agent θ if P does not hold in $\Delta \Im$, then $O^{\theta}_{\alpha}(P, \Delta \Im, \sigma) = B(a, O^{*}_{\theta}(P, \Delta \Im, *), \Delta \Im)$, where '*' represent unknown values, meaning that if the indicated deontic condition holds then it is possible to deduce that agent α believes that agent θ wants P to hold in $\Delta \Im$.

These norms are actually equivalent to defeasible rules in default reasoning (Reiter, 1980). The application of conflicting rules leads to multiple coherent sets of beliefs (theory extensions), which can be maintained by a TMS (*truth maintenance system*). The TMS identified the conflicting extensions but the selection of one of these extensions requires the usage of a preference relation – defined in the axiological component.

4.2 The Deontic Component

The deontic component of our agent information model is where all the possible behaviors of the agent are defined. Some authors might classify this component as *conative*² but that would elude the attitude-based classification derived from the classification of norms where we would like to base our model. Our concept partially (but not completely) departs from the traditional use of the word *deontic*, where the *deontic* qualification suggests ideal behavior. For example, standard deontic logic (SDL) represents and reasons about ideal situations. However, although an agent behavior is guided by deontic guidelines, in reality it is seldom seen an agent who behaves always in an ideal way. 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 some work has been done in that direction, e.g. by Dignum et al. (1994) and also by Carmo and Jones (1998). Here, we propose a method for representing goals as abstract agent behaviors, in the Deontic component of the EDA model, based

¹ Incomplete, in the sense that it does not imply logical omniscience.

² Conative refers to action and control, denoting an attempt to perform an action. The meaning of conation is "the conscious drive to perform apparently volitional acts with or without knowledge of the origin of that drive".

on the deontic operator 'ought-to-be' and leave preference determination and goal selection to be managed by a separate component - the axiological component (described in the next section).

4.2.1 Behavior Representation

Behaviors 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 others (Werner, 1989) have called the agent intentional state. However, in our model, agent decisions depend both on the available plans and a preference relationship defined in the axiological 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.

The plan ontology $\Box = \{\Box_1, \Box_2, ..., \Box_n\}$ may be modified by three different sources, who can add, remove or change deontic elements:

- The agent designer, through programming.
- Other agents, e.g. through communication or norms.
- The agent itself, through learning (adaptation).

A plan is typically given to the agent at a very abstract level, by specifying the goal that it ought to achieve, and then the agent should be able to decompose it into simple, executable, actions. This decomposition can be achieved by a means-ends process. By representing plans declaratively, as behavioral 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 directed reasoning in knowledge-based systems.

The Deontic Component contains the following kinds of elements:

- Atomic actions.
- Goal statements (agent abstract plans).
- Behavioral norms (socially accepted behaviors).

Atomic actions correspond to services that the agent is *able* to do, are represented procedurally and cannot be decomposed. These services may be offered to other agents or be reserved only for internal use. They are typically given to the agent embedded in organizational roles and may be used as building blocks for constructing complex plans.

Both goals and behavioral norms are represented declaratively, thus they may be manipulated by an inference machine, according to adequate syntactic rules, as suggested above. We believe that for the sake of representational simplicity and reasoning power, it would be useful to represent them in a unified way, independently of their source.

Our representation of ideal behaviors is inspired in deontic logic. However, 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 and Perloff, 1990), 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 $[\Box \text{ stit: } S]$ where \Box stands for an agent and S 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: $O[\Box \text{ stit: } S]$.

4.2.2 Unified Deontic Goals

A goal G represents an abstract plan. This corresponds, according to agency logic (Belnap and Perloff, 1990), to an agentive statement that an agent \square sees to it that the state of affairs P obtains at time \square_2 by an agent \square choice at time \square_1 . We claim that it is possible to use the same syntax and semantics both for representing agent goals and behavioral norms. This would make it possible to build up detailed plans from general goals and behavioral norms using means-ends analysis – the same inference mechanism commonly used in knowledge-based systems for backward-chaining.

We propose the following expressive modal representation for goals:

 $G_i = O_{\alpha}^{\theta}(P, \tau, \sigma) = O([\square \text{ stit } P] \text{ in-time-window } \square \text{ subject-to-sanction } \square \text{ by } \square)$

where O is the standard deontic operator 'ought-to-be'; $[\Box \text{ stit } P]$ is an agency statement, saying that agent \Box sees-to-it-that proposition P becomes true. This means that \Box will perform a plan to bring about P in time window \Box , where \Box is a time expression, specifying the time window during which proposition P is intended to be satisfied. \Box may be specified in absolute time or relative to some event; \Box indicates the sanction cost of violation.

This means that all agent goals can be represented as obligations, whether they are self-imposed obligations, moral obligations or commitments established in the course of their social activity.

The violation V_i of goal G_i is detected when the proposition $\sim P$ is true, in time window \Box . A sanction is a conditional behavioral norm (contrary-to-duty) that is brought into existence when a violation is detected, with a typically non-positive *estimated* value of \Box for \Box , and which may be performed by agent \Box , to whom \Box has committed to achieve *P*.

Assumptions regarding the sanction cost:

The expected cost has a probabilistic nature: $\sigma = \sum_{j} p_{j}c_{j}$, where c_{j} is each of the costs to be assumed by

agent \Box due to violation V_i , and p_i are their respective probabilities of occurrence. These probabilities are normative values defined in the epistemic component.

- The cost of seeing to it that *P* must be less than the sanction cost, otherwise the goal would never be intended.

Goals can be written in the scope of conditional statements, becoming behavioral norms, similarly to if-then rules. Depending on their sources, goals may have specific characteristics, the semantics of which are informally given below, for the instantiated case:

- Designer goals are implanted in the agent by the system designer. In this case \Box is the agent designer and both \Box and \Box are constants, defined by the agent designer. When a violation of this goal is detected, the agent reports an exception. This exception may be resolved programmatically if the designer provided a control for it, or interactively, if the exception has not been predicted by the designer or if he did not provide any control.
- Other agent goals include social normative goals (embedded in social roles), and communicated goals (transmitted in messages)
 - In the first case the sanction-enforcing agent □ is a collective agent, e.g. the organization, who
 must be represented by some (eventually) unspecified individual agent, for all practical
 purposes.
 - In the second case the sanction-enforcing agent \Box is another organizational agent, e.g. the boss or a client, who transferred goal G_i to \Box . The transference may occur either through a 'strategic action' based on a claim to power or through a 'communicative action' based on a validity claim (Habermas, 1984):
 - A claim to power is based on an implicit threat to agent □, supported by a power structure, e.g. the organization hierarchy. In this case, the goal transference has no other implications in the agent deontic component. It may however have axiological consequences.
 - A validity claim promotes an agreement between the two agents: this typically involves an exchange, i.e. agent □ also acquires a goal where □ is the sanction-enforcing agent. To succeed, this exchange must have a positive balance for both agents, i.e. the relative value of sanctions and benefits must be positive for both agents.
- Learned goals, generated by the agent itself, are goals where the sanction-enforcing agent □ coincide with □. There is no sense in ascribing negative values to sanction costs in this case. What the agent loses when it violates such a goal is an opportunity cost, i.e. a gain that does not materialize.

4.2.3 EDA Models Update

Within deontic action contexts it is possible to define pragmatic functions to update agents EDA models as a result of social interaction. A pragmatic function interprets signs that are exchanged amongst organizational agents and adds or revises declaratives stored in the Epistemic module, agentives stored in the Deontic module and preferences stored in the Axiological module.

In the important case of message-based communication, our model uses a pragmatic function based on Speech Act theory, to analyze the messages illocutionary points and content. The different kinds of speech acts are then converted, including imperatives, declaratives, commissives and expressives, into Beliefs, Goals and Values, updating the EDA model structures.

4.3 The Axiological Component

Preferred goals become intentions (Dignum et al., 1996). The axiological component provides the development of agent intentions based on agent goals and utility evaluation. Evaluation is a resource-consuming process. Ideally each component would work in parallel with the others. In that case, the agent would be able to re-evaluate goals continuously and update intentions accordingly, otherwise a meta-component must be included in the agent model to decide when to evaluate and when to act.

Since plans may, at some abstraction level, be in conflict so that the execution of one precludes the execution of another, it is possible that the set of goals that are defined in the ideal world are not coherent. This makes it necessary to take decisions, which theoretically should aim at minimizing a cost function, in a control sense. However, besides the fact that the unconstrained minimization of a general non-linear function is a very difficult

analytical problem, the mathematical definition of such a function has eluded researchers until now, thus we will adopt an heuristic perspective, commonly used in the Artificial Intelligence field, to approximately identify the relative costs of the various options available in each situation and to select the best one, as far as the agent axiological component.

The utility-based approach is often found in Economics, but many deontic logicians find it difficult to accept because moral values should be independent of individual utility. However, although we don't share this view, we will not discuss it here because, firstly, our objective is to model real-world organizations that base their behavior on economic values, rather than morally perfect organizations and, secondly, we envision the possibility of incorporating ethical and moral norms in the axiological component.

The axiological component of the EDA model provides preference relations both for the deontic component and for the epistemic component. In both cases, norms are represented as default rules. The problem is how to establish a preference amongst norms that would enable to solve dubious or conflicting situations.

A standard solution is to define a partial order between every pair of norms. For example, (Brewka, 1994) provides an extension to Reiter's default logic – Prioritized Default Logic (PDL) – a meta-level approach to generate preferred extensions of default logic:

Reiter defines a default theory as a pair $\Delta = (W, D)$, where *D* is a set of default rules and *W* is a set of first-order logic well formed formulas. A prioritized default theory is a triple $\Delta = (W, D, p)$, where p is a strict partial order over *D*, such that rule r₁ has priority over r₂ iff (r1, r2) $\in p$, or r1 p r2.

Given a set of formulae *E*, a default rule $a \rightarrow b \in D$ is active in *E* iff: $a \in E, \neg b \notin E, b \notin E$

Based on the notion of active default rule, Brewka presents the following definition:

E is an extension of Δ , generated by a total order = , containing p iff $E = \bigcup E_i$ where $E_0 = Th(W)$ and for $i \ge 0$:

$$E_{i+1} = \begin{cases} E_i & \text{If there is no active default rule in } E_i \\ Th(E_i \bigcup\{c\}) & \text{where c is the consequent of the minimal } = \text{ active default rule in } E_i \end{cases}$$

Traditional meta-level methods for priority assignment like PDL have some drawbacks: it is not always possible to determine beforehand the priority relation between rules. Furthermore, that relation may depend on additional knowledge not explicitly available. Some researchers have observed that reasoning about priorities is often an integral part of the problem solving activity in certain domains such as law. This has lead to new approaches that specify priorities at the representation level instead of using a meta-level. One of the most important advantages is that it becomes possible to specify general priority assignment rules instead of forcing the specification of a list of new priorities every time a new domain rule is introduced. For an overview of these methods, see (Cachopo, 1997).

5 RELATED WORK

Although inspired mainly in the semiotics stance, and the norms-attitudes relationships at different psychosociological levels, related to organizational modeling, the EDA agent model is related to several other models previously proposed, mainly in the DAI literature.

One of these is the BDI model (Belief, Desire, Intention) proposed by Rao and Georgeff (1990). This model is based on a theory of intentions, developed by Bratman (1987). The BDI architecture states that beliefs, desires and intentions are part of the agent mental states. Based upon an interpreter performing a perceive-decide-act loop, BDI uses 'beliefs' basically as a symbolic way of indicating the state information the agent has. Desires are just an anthropomorphic way to represent agent goals. Desires may be incoherent and impractical; therefore they have to be 'filtered' in order to select a plan, which then becomes an intentional goal. Typically the agent adopts an intention by picking up a plan from a plan library. The BDI architecture defines an agent internal architecture without much concern for its integration in a multi-agent system. The BDI perspective is more concerned with capturing the properties of human intentions, and their functions in human reasoning and decision making, whereas the EDA model is a norm-based representation of beliefs, goals and values, based on a semiotics view of information and oriented towards understanding and modeling social cooperation. BDI agents can easily abstract from any social environment because they are not specifically made for multi-agent systems modeling.

Another related model is the ICE architecture (Werner, 1996). ICE is a short for $I^2C^2E^2$, or information, intention, communication, cooperation, evaluation and empowerment. The ICE architecture is based on agents mental states represented by a tuple $R_A = (I_A, S_A, V_A)$, where I_A is the agent *information state*, S_A is the agent *strategic or intentional state*, and V_A is the *evaluation state* of the agent. The notion of information adopted by Werner is quite different from the semiotics notion of information, which seems to incorporate Werner's. Furthermore, Werner does not commit to a symbolic knowledge representation, leaving open all representational

issues, in all agent states. However, ICE shares with EDA an orientation towards modeling social cooperation.

Singh (1991; 1996; 1997) also provides a social perspective to multi-agent systems. He adopts a notion of commitment that bears some similarity with our goals, in the sense that it relates a proposition to several agents, defining the concept of 'sphere of commitment' (Singh, 1997) – a framework that emphasizes the interplay between commitments and social structure. A social commitment is defined as a four-place relation c = C(x, y, G, p) involving a proposition p and three agents: x (the debtor), y (the creditor) and G (the context group). The proposition p is the discharge condition of commitment c. Furthermore, Singh also refers to the concept of social norm, defining social norms as meta-commitments, i.e. commitments about propositions that refer to other commitments. However, in spite of the conceptual proximity between Singh's commitments and the EDA deontic goals, there is no reference in Singh's work to any kind of utility function or priorities associated with commitments. Singh does not approach the problem of resolving conflicting commitments or norm violations, although he recognizes that an agent's commitments *typically* constrain him to act in accordance with them.

Jennings (1994) proposes a social coordination mechanism based on commitments and conventions, supported by the notions of joint beliefs and joint intentions.

The work of Santos and Carmo (1996) is also related to ours, sharing a normative perspective of organizational behavior, that has motivated the incorporation of deontic notions in both kinds of organizational models. Their representation is also based on a combination of action logic and deontic logic, as in the Deontic component of the EDA model. They approach the important problem of indirect action and responsibility in organizations, and propose to represent goal delegation using a generalized action operator, G_jQ to be read 'agent i ensures that Q' either directly using his ability or indirectly using is power/influence eventually supported by the underlying organizational structure. These modal operators can build logical constructs like G_iG_jQ , to denote that agent i ensures Q through agent j. This operator is similar to our unified goal operator, where, in the previous example, θ would correspond to i and \Box would correspond to j.

6 CONCLUSIONS AND FUTURE WORK

The EDA model described here is based on the organizational semiotics stance, where normative knowledge and norm-based coordination is emphasized. The main model components (Epistemic, Deontic and Axiological) reflect a social psychology classification of norms, therefore provide a principled norm-based structure for the agent internal architecture that is also oriented towards a norm-based social interaction in organizations.

The EDA architecture integrates also a number of important ideas gathered mainly from the DAI field and from deontic logic. Some of the most important ones were described in the previous section. We recognize the need for a semantics to underpin the proposed model but, at the present, we have focused mainly on conceptual issues.

Particularly important for organizational modeling is the notion of 'commitment'. Many notions of commitment have been defined, both in DAI and deontic logic, all sharing some common aspects, as shown in the previous section. However, we agree with Staffan Hägg (1998), who defines a commitment as a "contract between the involved agents to reach and preserve a specified goal during a specified time". Although we didn't formally define our notion of commitment, we do see commitments in terms of goals, emerging as a pragmatic result of social interaction. We believe that multi-agent commitments can be modeled as related sets of deontic-action statements, distributed across the intervening agents, based on the notion of unified goals as proposed in the deontic component of our model.

An axiological component seems to be a necessary part of any intelligent agent, both to establish preferred sets of agent beliefs and to prioritize conflicting goals. Since we adopt a unified normative perspective both towards epistemic issues and deontic issues, both being based on the notion of norm as a default or defeasible rule, the axiological component is conceptualized as a meta-level Prioritized Default Logic. This methodological decision permits to inherit the results of active research on this relatively recent subject.

In a multi-agent environment the mutual update of agents' EDA models is essential as a result of perceptual events, such as message exchange. However, the specification of the EDA update using a pragmatic function is still the subject of current research, and will be reported in the near future. A related line of research that is being pursued at the moment involves the software simulation of EDA models, which raises some software engineering questions, related to the implementation of heterogeneous multi-agent systems implementation, where interaction aspects become a key issue, requiring a pragmatic interpretation of the exchanged messages.

REFERENCES

Ades, Y., 1999. Semantic Normal Form Compliance. 2nd International Workshop on Organizational Semiotics, Almelo, Holland.

Barbuceanu, M. and M. Fox, 1997. Integrating Communicative Action, Conversations and Decision Theory to Coordinate Agents. Proceedings of Autonomous Agents'97, Marina del Rey, USA.

- Barjis J. and J. Filipe, 2000. Simulation of Organizational Processes Combining Semantic Analysis and Petri Nets, submitted to the *Organizational Simulation Workshop of the Summer Computer Simulation Conference*, Vancouver, Canada.
- Belnap, N. and M. Perloff, 1990. Seeing to it that: a Canonical Form for Agentives. In *Knowledge Representation and Defeasible Reasoning*, Kyburg, Loui and Carlson (Eds.), Kluwer Academic Publishers, Dordrecht.
- Bratman, M., 1987. Intentions, Plans and Practical Reasoning. Harvard University Press, USA.
- Brewka, G., 1994. Reasoning about Priorities in Default Logic. In Proceedings of AAAI-94, AAAI Press, Seattle, USA.
- Cachopo, J., 1997. A Utilização de Prioridades no Raciocínio Não-Monótono. M.Sc. Thesis, IST, Lisbon.
- Carmo, J. and A. Jones, 1998. Deontic Logic and Contrary to Duties. (Unpublished) manuscript kindly provided by authors. Conte, R., and C. Castelfranchi, 1995. *Cognitive and Social Action*, UCL Press, London.
- Conte, R., and C. Castenraneni, 1995. Cognitive and Social Action, OCL Fress, London.
- Dignum, F., J.-J.Ch.Meyer, R. Wieringa and R. Kuipers, 1996. A Modal Approach to Intentions, Commitments and Obligations: Intention plus Commitment Yields Obligation. In Brown and Carmo (Eds.), *Deontic Logic, Agency and Normative Systems*, Springer Verlag, Berlin.
- Dignum, F., J-J Ch. Meyer and R.Wieringa, 1994. A Dynamic Logic for Reasoning About Su-Ideal States. In J.Breuker (Ed.), *ECAI Workshop on Artificial Normative Reasoning*, pp.79-92, Amsterdam, The Netherlands.
- Emerson, E., 1990. Temporal and Modal Logic. In J van Leeuwen (Ed.), Handbook of Theoretical Computer Science, pp. 996-1072. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands.
- Filipe, J., K. Liu and B. Sharp, 1997. A Semiotic Approach to Knowledge Sharing using Intelligent Agents. In *IWCS'97* (*First International Workshop on Computational Semiotics*), Paris, France.
- Filipe, J., B. Sharp and K. Liu, 1999a. An Intelligent Agent-Based Methodology For Legacy Information Systems Integration. In *Proceedings of UKAIS'99*, York, United Kingdom.
- Filipe, J., K. Liu and B. Sharp, 1999b. A Semiotic Approach to Organizational Role Modelling using Intelligent Agents, In 2nd International Workshop on Organizational Semiotics, Almelo, Holland,.
- Filipe, J., K. Liu and B. Sharp, 2000. Organization Simulation Based on Normative Knowledge and Role Modelling. In: *AAAI Spring-Symposium*, Stanford, USA.
- Ghezzi, C., M. Jazayeri and D. Madrioli, 1991. *Fundamentals of Software Engineering*. Prentice-Hall, Englewood-Cliffs, NJ. Gibson, J. 1979. *The Ecological Approach to Visual Perception*. Mifflin Company, Boston.
- Gibson, J. 1979. The Debogeta Approach to Visual Perception. Within Company, Boston.
- Habermas, J. 1984. The Theory of Communicative Action: Reason and Rationalization of Society. Polity Press. Cambridge.
- Hägg, S., 1998. Commitment in Agent Cooperation, Applied to Agent-Based Simulation. In Proceedings of the 5th International Conference on Intelligent Autonomous Systems (IAS-5), Sapporo, Japan.
- Hilpinen, R., 1971. Deontic Logic: Introductory and Systematic Readings. Reidel Publishing Co., Dordrecht, The Netherlands.
- Hintikka, J., 1962. Knowledge and Belief. Cornell University Press. Ithaca, USA.
- Jennings, N., 1994. Cooperation in Industrial MultiAgent Systems , World Scientific Publishing, Singapore.
- Jones, A. and M. Sergot, 1996. A Formal Characterization of Institutionalized Power. Journal of the IGPL.
- Konolidge, K., 1986. A Deduction Model of Belief. Morgan-Kaufman, Inc.
- Liu, K., 2000. Semiotics Information Systems Engineering. Cambridge University Press. Cambridge.
- Mintzberg, H., 1979. The Structuring of Organizations. Prentice-Hall, Englewood Cliffs, N.J., USA.
- Peirce, C. 1960. Collected papers of Ch. S. Peirce (1931-1935), C. Hartshorne and P. Weiss (Eds). Cambridge, Mass.
- Rao, A. and M. Georgeff, 1991. Modeling Rational Agents within a BDI architecture. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning* (Nebel, Rich and Swartout, Eds.),439-449. Morgan Kaufman, San Mateo, USA.
- Reiter, R., 1980. A Logic for Default Reasoning. In Artificial Intelligence 13, pp. 81-132, North-Holland, Amsterdam.
- Santos F. and J. Carmo, 1996. Indirect Action, Influence and Responsibility. In Brown and Carmo (Eds.), *Deontic Logic, Agency and Normative Systems*, Springer Verlag, Berlin.
- Scott, W., 1992. Organizations: Rational, Natural and Open Systems, third edition, Englewood Cliffs, N.J.: Prentice-Hall.
- Simon, H. 1996. The Sciences of the Artificial (3rd Edition). MIT Press, USA.
- Singh, M. and N. Asher, 1993. A Logic of Intentions and Beliefs. Journal of Philosophical Logic, (22),513:544.
- Singh, M., 1991. Intentions, Commitments and Rationality, in Proceedings of the Annual Conference of the Cognitive Science Society.
- Singh, M., 1996. Multiagent Systems as Spheres of Commitment. In Proceedings of the Int'l Conference on Multiagent Systems (ICMAS) Workshop on Norms, Obligations and Conventions. Kyoto, Japan.
- Singh, M., 1997. Commitments in the Architecture of a Limited, Rational Agent. In Lawrence Cavedon (Ed.), *Intelligent Agent Systems: Theoretical and Practical Issues*, Springer Verlag LNAI 1209, Berlin.
- Stamper, R., 1973. Information in Business and Administrative Systems. John Wiley & Sons.
- Stamper, R., 1996. Signs, Information, Norms and Systems. In Signs of Work, Semiosis and Information Processing in Organizations, Holmqvist et al. (Eds.), Walter de Gruyter, Berlin, New York.
- Stamper, R., K. Althaus and J. Backhouse, 1988. MEASUR: Method for Eliciting, Analyzing and Specifying User Requirements. In *Computerised Assistance during the Information Systems Life Cycle*, Olle, Verrijn-Stuart and Bhabuts (Eds). Elsevier-Science Publishers, Holland.
- Stamper, R., K. Liu, M. Hafkamp and Y. Ades, 2000. Understanding the Roles of Signs and Norms in Organizations a Semiotic Approach to Information Systems Design, Behavior and Information Technolog, 19(1), 15-27.

- Taylor, J., 1998. The Limits of Rationality in Communication Modelling a semiotic reinterpretation of the Concept of Speech Act. In Proceedings of the 3rd Workshop on Language Action Perspective on Communication Modelling, Jönköping.
- van Reijswoud, V. and J. Dietz, 1999. Business Process Redesign with DEMO. In Proceedings of the 4th Workshop on Language Action Perspective on Communication Modelling, Copenhagen.
- Weigand, H., W. Heuvel and F. Dignum, 1998. Modelling Electronic Commerce Transactions a Layered Approach. In *Proceedings of the Language Action Perspective on Communication Modelling Workshop*, Jönköping, Sweden.
- Werner, E., 1989. Cooperating Agents: A Unified Theory of Communication and Social Structure. In *Distributed Artificial Intelligence* (Gasser and Huhns, Eds.), pp.3-36, Morgan Kaufman, San Mateo, USA.
- Werner, E., 1996. Logical Foundations of Distributed Artificial Intelligence. In O'Hare and Jennings (Eds) Foundations of Distributed Artificial Intelligence, Wiley-Interscience, New York.

Winograd, T. and F. Flores 1986. Understanding Computers and Cognition. Addison Wesley Publishing Company.

Wooldridge, M. and N. Jennings, 1995. Agent Theories, Architectures and Languages: A Survey. In Wooldridge and Jennings (Eds.) *Intelligent Agents*, Springer-Verlag, LNAI 890.