NORMAJINI: A NORM-BASED MULTI-AGENT SYSTEM ARCHITECTURE USING JINI

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

Keywords: Architectures, Prototype Systems, Norms, Institutions, Open environments.

Abstract

This paper describes a normative multi-agent system architecture and a software system which implements that architecture. Our purpose is to provide a framework for designing agents with a social competence. We assume that this social competence in organizations is related to multi-agent co-ordination, which in turn involves information sharing, either dynamically (through communication) or statically (through institutions or culture), based on norms of various types (e.g. perceptual, cognitive, behavioral, evaluative). Since 'information' is an ill-defined word we prefer to adopt the semiotics framework, which uses the 'sign' as the elementary concept. Information as a composition of signs is then analyzed at different levels, including syntax, semantics, pragmatics and the social level.

Based on different properties of signs, found at different semiotic levels, we propose here a new agent model, designated by EDA (an acronym for its three component modules: Epistemic-Deontic-Axiological), to represent agent informational states and simultaneously define its conceptual communication framework. The EDA model components incorporate 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 or an authorization-to-do. EDA agents are designed to promote social, co-operative, behavior. Agents use their role description both prescriptively, exhibiting a goal-directed behavior, and proscriptive, acting and co-operating within the boundaries of role policies. However, agents are autonomous in the sense that they are given enough decision power to allow them to choose to occasionally violate their obligations, depending on evaluations provided by the axiological component of the EDA model.

The multi-agent social architecture defined in this paper is flexible enough to accommodate changes in social structure, including changes in role specification, role instantiation and even the dynamics of institutional relationships, including role removal or creation. Amongst current technological infrastructures, one of the best to provide a suitable environment for the implementation of multi-agent systems, with high adaptability in dynamic environments, is JINI® – a distributed computing infrastructure developed by Sun™. Besides communicative capabilities, social agents require some degree of intelligence, at least in order to interpret pragmatically their perceptions, such as the messages they receive from other agents. This aspect has been implemented with the help of a rule-based development tool – JESS (Java Expert System Shell) – developed by Sandia Labs. To demonstrate the adequacy of JINI and JESS to implement normative multi-agent systems, based on the EDA model, we present here a simple but illustrative case study and explain some of JINI advantages as a communication infrastructure for multi-agent systems programming. We also analyze the JESS suitability to represent and interpret norms, giving some examples.

1 INTRODUCTION

Organizational Semiotics is a particular branch of Semiotics, the formal doctrine of signs (Pierce, 1960), concerned with analyzing and modeling 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. Since organizational activity is an information process based on the notion of responsible co-operative agent, we propose a model that accommodates both the social dimension in organizational agents behavior and the relative autonomy that individual agents exhibit in real organizations.

The proposed model is an intentional model, based on three main components, trying to capture relevant agent mental attitudes. In section 7 we provide a comparative critical analysis of other intentional models that have been proposed in the Distributed Artificial Intelligence (DAI) literature. Although a formal description can be provided, in this paper we take a design approach, showing how the EDA architecture can be used to model a case study and then how it is implemented on top of JINI and JESS. Since multi-agent logic is not sufficiently developed to achieve the logical consistency and completeness of organizational models, computational simulation provides a way to assess the adequacy of organizational models. Furthermore, computer simulation permits to gain insights into organizational phenomena and to explore processes and configurations that are difficult or impossible to investigate with other methods.

However, previously to design a simulation model it is necessary to understand system requirements, using an adequate method for systems analysis.

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.

2 ORGANIZATIONAL SEMIOTICS

In this paper we approach this problem using the Organizational Semiotics stance (Stamper, 1973; Liu, 1993) to provide adequate system requirements and a solid conceptual basis for simulation models. Semiotics, 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. A detailed and formal account of these levels may be found in (Stamper, 1996).

This approach is different from mainstream computer science because instead of adopting an objectivist stance – where it is assumed the existence of a single observable reality, external to the agent, which requirements modeling methods try to capture with the help of entity-relationship models and data-flow diagrams – it adopts a social subjectivist stance. This means that for all practical purposes nothing exists without a perceiving agent nor without an agent engaging in action (Stamper et. al., 2000). This philosophical stance ties every item of knowledge to an agent who is, in a sense, responsible for it. Thus, adopting this stance,

> (…) *the information model avoids the concept of truths that have been delivered from on high, independently of a knowing agent. Truth is something that the individual agents and group agents have to decide upon and for the consequence of which they have to take responsibility (Stamper, 1996 p.371).*

The distinction between entity, attribute and relationship is therefore replaced by the concepts of agent, affordance and norm (a socially defined pattern). These concepts are related in a way that indicates ontological dependency[1].

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 1.

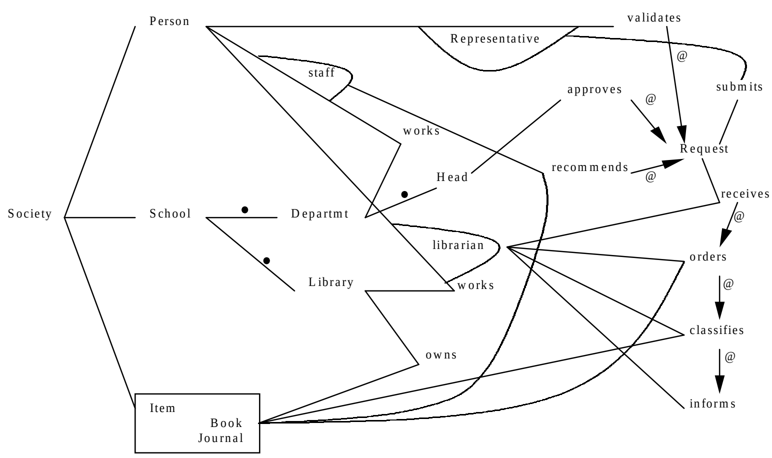

Figure 1: Ontology Chart.

[1] The word ontology is used here in a different sense than the one commonly found in AI. By ontology dependency we mean an existential relationship, i.e. that a concept's existence depends on others' existence.

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 (●) 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 half-circles. 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 Paradigm Shift

This approach is different from mainstream computer science because instead of adopting an objectivist stance – where it is assumed the existence of a single observable reality, external to the agent, which requirements modeling methods try to capture with the help of entity-relationship models and data-flow diagrams – it adopts a social subjectivist stance. This means that for all practical purposes nothing exists without a perceiving agent nor without an agent engaging in action (Stamper et al., 2000). This philosophical stance ties every item of knowledge to an agent who is, in a sense, responsible for it. Thus, adopting this stance,

(…) the information model avoids the concept of truths that have been delivered from on high, independently of a knowing agent. Truth is something that the individual agents and group agents have to decide upon and for the consequence of which they have to take responsibility (Stamper, 1996 p.371).

The distinction between entity, attribute and relationship is therefore replaced by the concepts of agent, affordance and norm (a socially defined pattern). These concepts are related in a way that indicates ontological dependency[2].

The social world level, especially in what concerns social norms and commitments, is particularly relevant for information systems analysis and simulation, because organizational activity is mainly the result of the coordinated behavior of several organizational units/agents. Using this semiotics framework to cover the main stages of systems lifecycle, a collection of methods has been defined, which can be applied to all the systems development activities along the systems development lifecycle. These were developed under the MEASUR research program – Method for Eliciting, Analyzing and Specifying User Requirements (Stamper et al. 1988).

This semiotic framework was recommended by FRISCO (Falkenberg, 1998) as a philosophical foundation for information systems analysis.

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 order to achieve social goals, in a cooperative manner (Wooldridge and Jennings, 1995; Singh, 1996).

3 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;

[2] The word ontology is used here in a different sense than the one commonly found in AI. By ontology dependency we mean an existential relationship, i.e. that a concept's existence depends on others' existence.

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[3], 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 2, where it is shown that information is a complex concept, and requires different viewpoints to be completely analyzed.

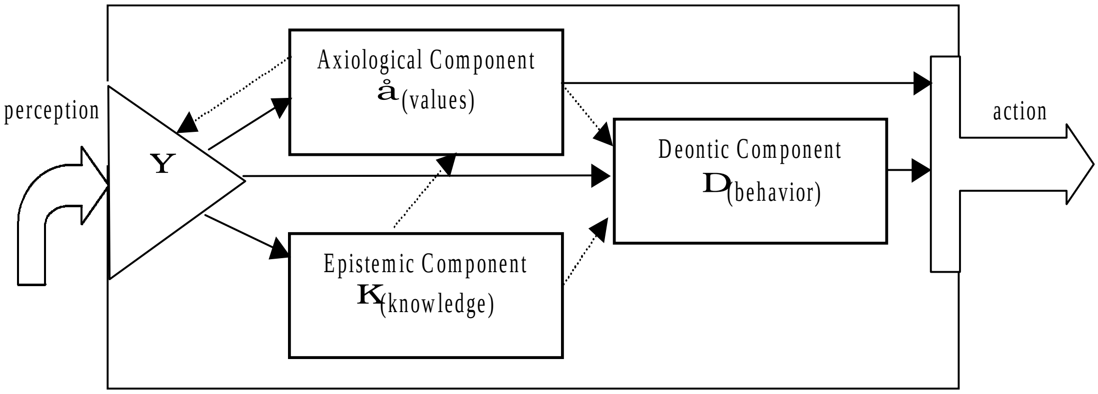

Figure 2: The EDA agent model.

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

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

K 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.

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

4 KNOWLEDGE REPRESENTATION

To encode, symbolically, the agent architecture conceptually described in the previous section, we need to use higher abstract level constructs than those found in mainstream programming languages, including Java. Since agents are intentional, since the EDA model ascribes them mental attitudes, it seems adequate to use Artificial Intelligence knowledge representation paradigms. All model components require both procedural and declarative knowledge. Whereas procedural knowledge is more efficient, declarative knowledge provides grater

[3] von Wright (1968) suggests that the study of deontic concepts and the study of the notions of agency and activity are intertwined.

flexibility and adaptability, thus we decided to use a hybrid paradigm, combining frame systems with rule-based systems, which allows the separation of declarative knowledge from procedural general purpose inference machine and furthermore provides the capability of using procedural attachments where needed.

4.1 EDA Model Components

We provide here only a description of the knowledge representation chosen for each model component. A more detailed syntactic and semantic analysis of each component is provided in (Filipe and Liu, 2000).

Epistemic component

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 $\mathcal{A} = \{\alpha_1, \alpha_2, \ldots, \alpha_k\}$ be the set of agents and let $\mathcal{F} = \{\varphi_1, \varphi_2, \ldots, \varphi_m\}$ be the set of affordances, represented in the ontology chart; Let $\mathcal{R} = \{\rho_1, \rho_2, \ldots, \rho_n\}$ 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: $\rho_i = \rho_i(\alpha_k \mid \alpha_j, \alpha_k \mid \varphi_k \mid \varphi_j, \varphi_k)$.

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.

Entities in our ontology are represented by frame objects; properties and special relationships are represented by frame slots. This kind of representation is particularly adequate for classification hierarchies where more general types subsume more specific ones and property inheritance is a useful inference mechanism. Other kinds of knowledge, not easily amenable to hierarchical class-subclass relationships, are represented using production rules and deductive inference.

Following the recommendation of Liu (1993) we use a semantic temporal database concept to keep track of affordances existence, although we have preferred to represent affordances as beliefs in the system’s knowledge base.

Cognitive norms are represented as conditional beliefs, i.e. rules. We use a temporal qualifier to identify the time interval when the norm applies.

Deontic component

Norms, commitments and plans are represented in an unified way: as goals, according to the modal operator proposed in the EDA model description:

$$G_i = O_\alpha^\theta(P, \tau, \sigma) = O\ ([\alpha \text{ stit } P] \text{ in-time-window } \Delta\Im \text{ with-utility } [\ \mu\ \ \sigma\ ] \text{ commited-to } \theta)$$

where O is the standard deontic operator ‘ought-to-be’ and [α stit P] is an agency statement, saying that agent α sees-to-it-that proposition P becomes true. This means that α will perform a plan to bring about P in time window $\Delta\Im$, where $\Delta\Im$ is a time expression, specifying the time window during which proposition P is supposed to be satisfied. μ and σ are places to be filled in by the axiological component, with the utility of doing and of not-doing, respectively, the action [α stit P]. Although the exact values must be calculated by the axiological component, based on the current situation, the deontic component may fill in default values, eventually using common knowledge, i.e. norms, about the domain. The default value of μ indicates the expected value of [α stit P] for α, i.e. the expected benefit of α fulfilling its obligation; σ indicates the expected value of the sanction, i.e. the expected violation cost.

Planning becomes a goal-directed behavior simulation. We use the rule-based paradigm for representing means-ends process, which constitutes the base of our planning activity, and is based on a process of goal decomposition and sequencing. This is made using a backward chaining inference process, over rule sets where rule antecedents represent less general goals and consequents represent more general goals.

This basic concept is very similar to the *production rules* paradigm, used by most expert systems. However, two changes had to be incorporate into the framework before it worked as desired:

- Firstly, it is necessary to avoid the automatic insertion of goals into the system's agenda. Goals can only be introduced into the agenda after being filtered by the axiological component, for value assignment and filtering;
- Secondly, time must be taken into account, both to trigger events that signal the start or finish of some affordance and to identify the existence of affordances during hypothetical reasoning, during the plan generation process.

According to the model dynamics defined in the previous section, the start and finish of affordances is registered in the semantic temporal database, by deontic norms and agents.

Axiological component

Agent preferences are represented as meta-norms, i.e. as rules in a meta-language outside the domain representation language. Objects in this meta-language define norm priority as a partial order relation between norms. This process has been represented as a knowledge-based system to be used exclusively by the axiological component of an EDA agent.

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 r_1 has priority over r_2 iff $(r1, r2) \in p$, or $r1 \; p \; r2$.

Given a set of formulae *E*, a default rule $a \to b \in D$ 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 $i \geq 0$:

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

This component permanently computes utilities associated to deontic norms and, using a 'any-time algorithm' is able to suggest the best agent's next action, whenever required to do so.

4.2 Organizational Roles Representation

The representation of roles is composed of two distinct knowledge entities: services and policies (deontic statements, including obligations and authorizations).

$$Roles = \{Services, Policies\}$$

$$Policies = \{Obligations, Authorizations\}$$

Whereas services represent the role know-how and are represented procedurally in the Epistemic component, policies are deontic statements that define the obligations and authorizations that come with the role, as normative knowledge, and are represented in the Deontic component. In the latter case we have chosen a declarative paradigm and represent policies as prioritized default rules. The declarative representation enables the usage of a (non-monotonic) deductive inference process.

Hierarchical and other power relationships are defined by sets of institutional norms represented essentially in the deontic components, e.g. 'if a lecturer wants a book then he must send his request to the department's library representative'; 'the head of department must authorize all requests previously to the representative sending them to the library'; 'The head of department must see to it that each request does not exceed 1% of the total budget for the department'; 'Lecturers may ask for resources directly from the head of department, unless an alternative procedure exists'.

These example statements are not to be stored in a single organizational knowledge-based but rather in separated role knowledge bases; they are uploaded when needed into an agent knowledge base. This perspective avoids the existence of highly centralized and monolithic pieces of the information system, in favor of a distributed more flexible information system.

A role is encoded in an organizational COTS ("component-off-the-shelf") and can be "up-loaded" to any agent that is required to play it, typically as a result of the need to engage in a conversation with another agent.

4.3 Multi-Agent Interaction

Based on the theoretical work of Winograd and Flores (1986) several conversation-based business process models have been developed, under the designation of 'Language Action Perspective' including Action Workflow (Medina-Mora et al., 1992), DEMO (Dietz, 1990; van Reijswoud, 1996) and others.

Figure 3 depicts a typical *conversation for action*. Step 3, in this figure, is *the* action execution step, requiring agent access to the EDA deontic component, where the services the agent is able to perform are stored under a procedural representation.

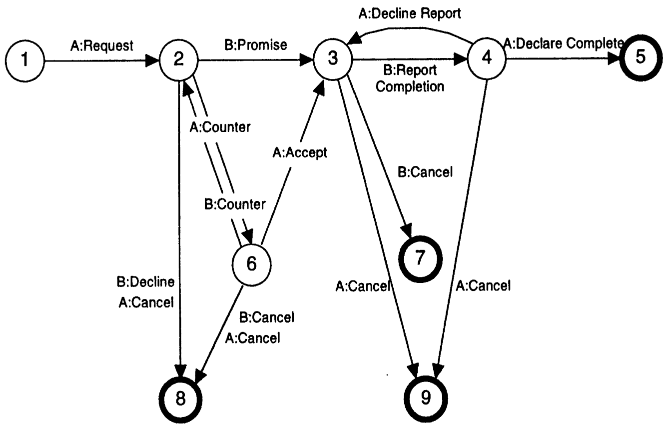

Figure 3: A conversation for action (Winograd and Flores, 1986)

The design of conversations, in our case, is made using syntax-directed translation schema and attributed grammars (Fred and Filipe, 2000).

5 SOFTWARE FRAMEWORK

The framework developed to support EDA agents is based on two Java environments: JINI and JESS. JINI provides the communication infrastructure, while JESS provides the knowledge-base support as a frame and rule based knowledge representation language with an inference machine capable of forward chaining and backward chaining (figure 4).

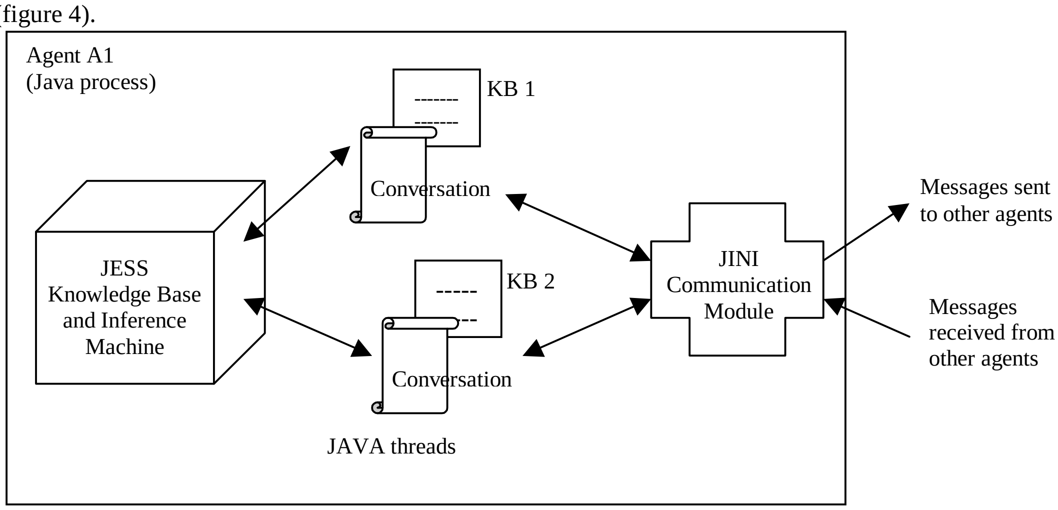

Figure 4: Software Agent Architecture

5.1 Communication Infrastructure

The communication infrastructure is provided by JINI™. Jini is a network-centric computing architecture, reflecting the computing paradigm change described above, based on social concerns. A key aspect of this new paradigm is that social systems are not static: agents and affordances have a dynamics, where start and finish events are determined either by responsible agents or automatically by social norms. Since agents are autonomous they may decide to form groups dynamically, and Jini facilitates this by enabling a "plug and play" network where clients access services by their Java interface, and by promoting the concept of *federation of services*, where *service* may be e.g. a hardware or software agent, a communication channel, or the user. Once part of a federation, a service can be used by other services or clients. Services are accessed through lookup services, which act as facilitators to locate services. There are basically three processes defined in Jini specifications:

Discovery: Devices/Clients locate and obtain references to lookup services.

Join: Devices register their services with lookup services.

Lookup: Clients locate and contact services

A new service makes itself available by dropping a multicast UDP “presence announcement” while lookup services listen for presence announcements, and contact senders directly to confirm discovery, sending them a service registrar object. Services can then register (via join) through the service registrar

Once a service is registered, clients may access it. A service object may then define the access protocol. Although service protocols are private, maximum flexibility is achieved when the service is implemented transparently and the client just needs to invoke methods declared in the Java interface.

Jini programming provides some new distributed control mechanisms and data structures, including: distributed leasing, distributed events and distributed transactions.

There is a service that is embedded in the Jini framework and makes use of these programming constructs: it is JavaSpaces™ - Sun's first Jini service. JavaSpaces is based on the concept of blackboard, developed in the AI field, and it assists in distributed persistence and the design of distributed algorithms. A JavaSpaces server holds entries: typed groups of objects. A client may look up entries with templates: entry objects that have some or all fields set to values that must be matched exactly. It is possible to request a space to notify when an entry that matches a template is written (uses Jini distributed events). All operations that modify a space are done in a transactionally secure manner with respect to that space and clients can use transactions to do multi-operation or multi-space updates atomically (uses Jini transactions). Entries written to a space are leased (uses Jini leasing).

5.2 Conversation Management

Agents are Java programs and conversations are Java threads running within agents, using the JINI infrastructure to establish communication channels with other agents. The content of messages exchanged through these communication channels are encoded in an agent communication language (ACL). Several efforts have been made to specify a standard ACL – an *inter-lingua* to be used by heterogeneous agents, independently developed. KQML (Finin et al., 1994) and FIPA (*Foundation for Intelligent Physical Agents*) are two examples of such efforts. However, since no ACL has been widely adopted as a standard, so far, we decided to use our own specification for message content, based on Speech Act Theory (Austin, 1962; Searle, 1969). The result is similar to a subset of the KQML ACL.

JESS, briefly described below, is the software environment where the message interpreter has been developed. It is Java-based and it provides both frames and production rules, including an inference machine that can work both in forward and backward chaining.

5.3 Normative Framework: Roles, Services and Policies

The organization normative structure is essentially defined by the organizational roles and their relationships. The set of external functions and procedures (methods) that an object offers constitutes its *interface*. An agent role consists of a Java object presenting an interface reflecting the agent know-how plus a set of policies. A policy is a set of rules related to one or more EDA components. Therefore, each EDA component may receive a set of rules as a result of the uploading of a single agent role.

An EDA policy may represent an obligation or an authorization. Obligations are represented as goals, using the goal definition indicated in section 4. Authorizations are represented using the same syntax as goals but in a template format, i.e. where some of the elements are wildcards (match any constant that may be in their place).

6 CASE STUDY

This case study reflects an organizational process of the Polytechnic Institute of Setúbal concerning the book acquisition process, initiated by a request from a lecturer, as depicted in figure 5.

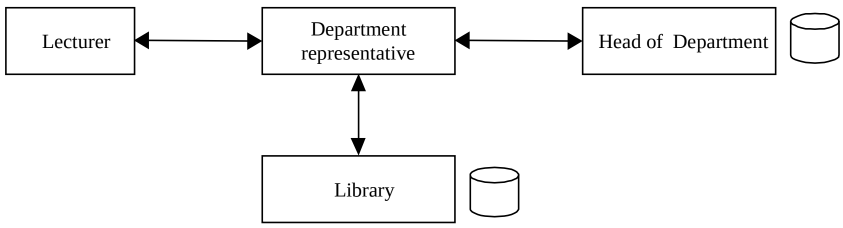

Figure 5: Book acquisition process model (brief)

6.1 Informal Description

The book acquisition process is started by the lecturer, who merely states the desire to obtain a certain book. This is picked up by the person that performs the mediator role between the department to which the lecturer belongs and the library. This mediator must enquire if the president of the department considers the book worth being acquired and has enough money to buy it. If the answer is positive then the library will purchase the book and will inform the lecturer as soon as the book becomes available.

6.2 Domain Ontology

Each organizational agent may be a human agent or an artificial agent. However, artificial agents are acting in the organizational network on behalf of users who are ultimately responsible for these agents. In order to achieve a general understanding of the organizational processes, independently of the agent who performs which role, it is necessary to establish the relevant domain ontology. This can be done using an ontology chart. The ontology chart for this case study is depicted in figure 1.

6.3 Roles

There are four roles identified in the case study. Each role corresponds to a set of services plus a set of policies and it is represented as facts and rules in a knowledge base. To illustrate the process of construction of a role, we will focus on the department representative agent.

Services

Services are advertised using beliefs asserted in the role knowledge base. An example, using JESS is:

```
(assert (service (name book-Acquired) (client DSI-lecturer) (supplier Lib)))
```

this means that this agent performs a service to achieve the goal ‘book-Acquired’ but only if the client is a DSI-lecturer; the service may be delegated to a supplier agent ‘Lib’.

A service may be an atomic task, represented procedurally by a method, or it may be a composition of tasks, represented declaratively.

Policies

Policies are sets of Obligations and Authorizations. Role obligations are default norms imposed by the organization to which the agent belongs. In our example, the agent role is ‘Rep’, and the organization is ‘EST’. An example in JESS is:

```
IF  (request (from ?agent)(service book-acquired)(details ?book1)) and
    (lecturer (name ?agent) (type DSI-lecturer))
THEN
(obligation (agent Rep) (stit book-acquired ?book1) (time now now+300)(utility 0)(cost 200)(sanction EST)))
```

This means that if the Rep agent receives a request from a DSI-lecturer then it must accept the commitment to see to it that the book is acquired, although the actual satisfaction of the commitment may not occur, e.g. if the Rep agent does not get the required authorization from the head of the department. In that case a second obligation comes into existence that forces the violation of the first. In this case the sanction cost will be zero because the violation has been ‘accepted’ by the organization that controls the sanction application.

Every agent goal has an implicit precondition, to be satisfied before the goal is placed in the agent's agenda: that the agent owns an authorization for achieving that goal. An authorization is a goal template, in which some places are filled in by wildcards, such as in the example:

(assert (authorized (agent Rep) (stit book-Acquired :title "*" :author "*" :cost "<1000") (time *)(utility *)(cost *)(sanction *)))

6.4 Conversations

A conversation for action, of the type indicated in figure 3, is represented by two Java threads: one connected to the client agent (the initiator) and one connected to the service provider. The Jini framework provides a direct message-based communication infrastructure and also a blackboard service (called JavaSpaces™) associated to a lookup server, which is used as a facilitator agent, providing the mechanisms for clients to match their service requests with registered service providers.

Conversations develop accordingly to the transition rules specified by a finite state machine, which we implemented declaratively in JESS using a set of rules. Although conceptually equivalent to a finite state machine this method allows for dynamic modification of the transition rules, thus enabling on-line machine learning.

7 RELATED WORK

Although inspired mainly in the semiotics stance, and the norms-attitudes relationships at different psycho-sociological 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 (1991). 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 i*nformation 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.

With respect to enterprise modeling methods, we recognize the importance of comprehensive frameworks such as CIMOSA or GERAM, which provide an elaborate description of the enterprise, including functional, informational, resource and organizational aspects. However, these models are not specifically oriented towards distributed knowledge representation and social organizational relationships, and furthermore do not provide a good way to understand the structure of business processes. This explanatory kind of analysis is however one the strengths of rule-based architectures, like the one we proposed for representing our agents' knowledge in general and their goals in particular. Yu and Mylopoulos (1997) also recognized the importance of explicitly representing and dealing with goals, in terms of means-ends reasoning, and they have proposed the i* modeling framework, in which organizations and business process models are based on dependency relationships among agents.

8 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.

Austin, J., 1962. *How To Do Things with Words*. Oxford University Press: Oxford, England.

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.

Conte, R., and C. Castelfranchi, 1995. *Cognitive and Social Action*, UCL Press, 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.

Falkenberg, E., W. Hesse, P. Lindgreen, E. Nilsson, H. Oei, C. Rolland, R. Stamper, F. Assche, A. Verrijn-Stuart, K. Voss, 1998. FRISCO: A Framework of Information System Concepts, (web edition), IFIP, ftp://ftp.leidenuniv.nl/pub/rul/fri-full.zip.

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, 1997. A Semiotic Approach to Knowledge Sharing using Intelligent Agents. In *IWCS'97 (First International Workshop on Computational Semiotics)*, Paris, France.

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.

Filipe, J. and K. Liu, 2000. The EDA Model: An Organizational Semiotics Perspective To Norm-Based Agent Design. Submitted to the Workshop on Norms and Institutions in Multi-Agent Systems, Barcelona, Spain.

Finin, T. D. McKay, R. Fritzon and R. McEntire 1994. KQML: an Information and Knowledge Exchange Protocol. In K. Fuchi and T. Yokoi (Eds) Knowledge Building and Knowledge Sharing. Ohmsha and IOS Press.

Fred A. and J. Filipe, 2000. Syntax-Directed Translation Schemes For Multi-Agent Systems Conversation Modeling. Submitted to ICEIS 2000, Stafford, UK.

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.

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.

Medina-Mora, R., T. Winograd, R. Flores and F. Flores, 1992. The Action Workflow Approach to Workflow Management Technology. In: Proceedings of the 4th Conference on Computer Supported Cooperative Work. ACM, New York, pp. 281-288.

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.

Rao, A. S., and Georgeff, M. P., 1991. Modeling rational agents within a BDI-architecture. In (J. Allen, R. Fikes, and E. Sandewall, eds.) *Proceedings of the Second Int'l Conference on Principles of Knowledge Representation and Reasoning*, 473-484, Morgan Kaufmann, 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.

Searle, J., 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press: Cambridge, England.

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.

von Wright, G., 1968. An Essay in Deontic Logic and the General Theory of Action. Acta Philosophica Fennica, 21. North-Holland.

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.

Yu, E. and J. Mylopoulos, 1997. Modelling Organizational Issues for Enterprise Integration. In Proceedings of ICEIMT'97, 529-538.