MULTI-AGENT SYSTEMS IN INTELLIGENT PERVASIVE SPACES

Joaquim Filipe
Escola Superior de Tecnologia de Setúbal, Instituto Politécnico de Setúbal
Rua Vale de Chaves, Estefanilha, 2910-761 Setúbal, Portugal

Keywords: Multi-agent Systems, Theoretical Agent Models, Organisational Semiotics, Social Contexts.

Abstract: This paper describes an agent model based on social psychology and also on the concept of organisational semiotics information fields to provide a conceptual infrastructure for designing multi-agent systems in intelligent pervasive spaces. 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 adopt the EDA agent model (an acronym for its three component modules: Epistemic-Deontic-Axiological). Intelligent pervasive spaces are ICT-enhanced physical and social spaces, differing from traditional pervasive computing on the focus, which in pervasive spaces is essentially social instead of technological (Liu et al., 2010). Agents are often described in terms of their internal structure, emphasizing their autonomy even in social settings involving communication and coordination. In this paper we suggest that agents can be seen both as individual and social entities, simultaneously. The norm-based multi-agent social architecture defined in this paper is flexible enough to accommodate changes in social structure, including changes in role specification, and representation of inter-subjective social objects.

1 INTRODUCTION

Social groups can be seen as multi-agent systems, possibly including both human and artificial agents. If there is a strong social cohesiveness, then we may be in the presence of organisations, which can be modeled as multilayered Information Systems (IS), composed of an informal subsystem, a formal subsystem and a technical system as shown in figure 1. this structure is typical of the organisational semiotics perspective on information systems.
Organisational Semiotics is a particular branch of Semiotics, the formal doctrine of signs (Peirce, 1931-1935), concerned with analyzing
and modeling organisations as information systems. Core concepts such as *information* and *communication* are very complex and ill-defined concepts, which should be analysed 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 organisational activity is an information process based on the notion of responsible co-operative agents, we propose a model that accommodates both the social dimension in organisational agents behavior and the relative autonomy that individual agents exhibit in real organisations. The proposed model is an intentional model, based on three main components, each of them trying to capture particular relevant agent attitudes. Agents are seen as intelligent units of a larger distributed system, in the sense that each unit has an autonomous capacity to infer and act, based on a knowledge-based infrastructure (Filipe, 2002).

In our research work intelligent agents are placed in social information fields, or spaces, where they interact with other artificial agents or humans, on behalf of human users or human organisations, who must ultimately take responsible for the behavior of each artificial agent.

Organisational Semiotics, however, is not sufficiently developed to provide an analytical

Filipe J.
MULTI-AGENT SYSTEMS IN INTELLIGENT PERVASIVE SPACES.
DOI: 10.5220/0003258800090016
In *Proceedings of the Twelfth International Conference on Informatics and Semiotics in Organisations (ICISO 2010)*, page
ISBN: 978-989-8425-26-3

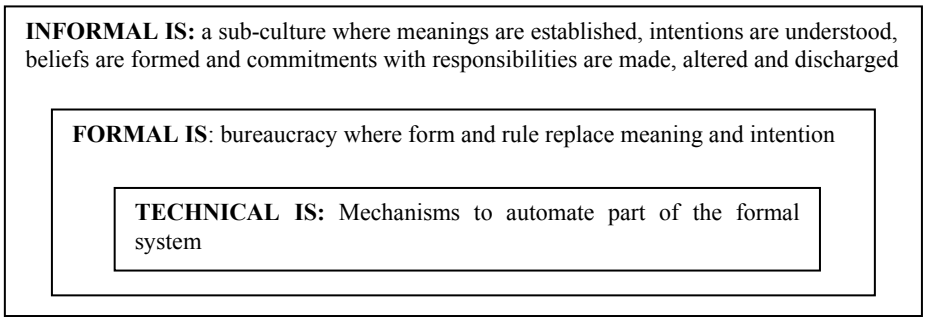

Figure 1: Three main layers of the real information system (Stamper 1996).

model for designing each agent in the organisational social system. This paper extends the work that has been done in semantic analysis, providing a way to clearly define the specification of individual agents at a pragmatic level, keeping a social and normative perspective.

2 ORGANISATIONAL SEMIOTICS

In this paper we approach the problem of constructing multi-agent systems in pervasive spaces using the Organisational Semiotics stance (Stamper, 1973; Liu, 2000), to provide adequate system requirements and a solid conceptual basis. Semiotics, traditionally divided into three areas – syntax, semantics and pragmatics – has been extended by Stamper 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 some modeling methods try to capture with the help of some software engineering approach – 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, 2000). Invariant behaviors available to an agent are called *affordances*. This philosophical stance ties every item of knowledge to an agent who is, in a sense, responsible for it.

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, organisations 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; Filipe, 2000). Social goals can be seen as norms therefore we hypothesize that the conceptualization and development of these intelligent pervasive systems require normative models.

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 (related to perception);
- *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.

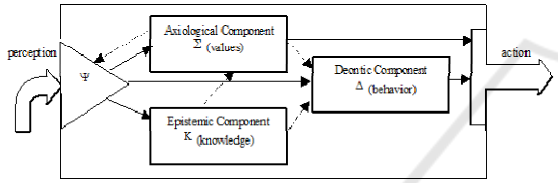

Figure 2: The EDA agent model.

Using this taxonomy of norms, and based on the assumption that an organisational agent behavior is determined by the evaluation of deontic norms given the agent epistemic state[1], we propose an intentional agent model, which is decomposed into three components: the epistemic, the deontic and the axiological (Figure 2).

Together, these components incorporate all the agent informational contents, where it is shown that information is a complex concept, and requires different viewpoints to be completely analyzed. The description and detailed analysis of each of the aforementioned components is provided in (Filipe and Liu, 2000).

4 RELATED WORK

Although inspired mainly in the semiotics stance, and the norms-attitudes relationships at different psycho-sociological levels, related to organisational 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 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.

Singh (1996) 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'.

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

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 organisations and business process models are based on dependency relationships among agents.

5 INTENTIONS AND SOCIAL NORMS IN THE PERVASIVE SPACES

Based on this agent model, we can create social structures composed of many interacting agents. The multi-agent system metaphor that we have adopted for modelling organisations implies that organisations are seen as goal-governed collective agents, which are composed of individual agents. This perspective comes in line with the principles of normative agents proposed in (Castelfranchi, 1993).

The social normative structure is essentially defined by agent roles and relationships. Roles can then be instantiated by one or more agents. Conceptually, a role is a set of Services and Policies, and a Policy is a set of Obligations and Authorizations. At the implementation level, agents are represented by objects and services are defined

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

by the object *interface*. Policies are sets of rules related to one or more EDA components, each of which includes at least one knowledge base (KB). When an agent is selected to perform a role, each of its EDA components downloads the adequate KB from an organisational role server.

Obligations are represented as particular goals whereas authorizations are represented using the same syntax as goals but in a pattern format, and are

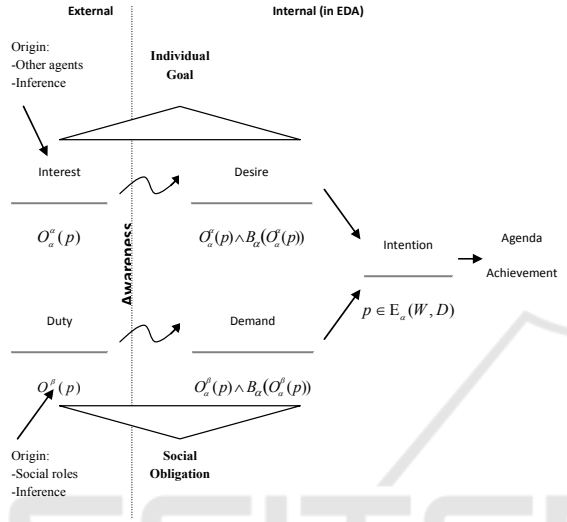

Figure 3: Social and Individual goals parallelism in the EDA model.

interpreted as potential action enabling/blocking devices.

Figure 3 describes the parallelism between mental and social constructs that lead to setting a goal in the agent's agenda, and which justifies the adoption of an obligation. Here, p represents a proposition (world state). $B_{\alpha}(p)$ represents p as one of agent α's beliefs. $O^{\beta}_{\alpha}(p)$ represents the obligation that α must see to it that p is true for β. $O^{\alpha}_{\alpha}(p)$ represents the interest that α has on seeing to it that p is true for itself – a kind of self-imposed obligation. In this diagram $p \in \mathrm{E}_{\alpha}(W, D)$ means, intuitively, that proposition p is one of the goals on α's agenda.

Interest is one of the key notions that are represented in the EDA model, based on the combination of the deontic operator 'ought-to-be' (von Wright, 1951) and the agentive 'see-to-it-that' stit operator (Belnap, 1991).

Interests and *Desires* are manifestations of *Individual Goals*. The differences between them are the following:

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

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

Not all interests become desires but all desires are agent interests. This may seem contradictory with a situation commonly seen in human societies of agents acting in *others' best interests*, sometimes even against their desires: that's what parents do for their children. However, this does not mean that the agent desires are not seen as positive by the agent; it only shows that the agent may have a deficient axiologic system (by its information field standards) and in that case the social group may give other

agents the right to override that agent. In the case of artificial agents such a discrepancy would typically cause the agent to be banned from the information field (no access to social resources) and eventually repaired or discontinued by human supervisors, due to social pressure (*e.g.* software viruses).

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

- *Duties* are social goals that are attached to the particular roles that the agent is assigned to, whether the agent is aware that they exist or not. The statement $O_\alpha^\beta(p)$ means that agent α ought to do p on behalf of another agent β. Agent β may be another individual agent or a collective agent, such as the society to which α belongs. Besides the obligations that are explicitly indicated in social roles, there are additional implicit obligations. These are inferred from conditional social norms and typically depend on circumstances. Additionally, all specific commitments that the agent may agree to enter also become duties; however, in this case, the agent is necessarily aware of them.
- *Demands* are duties that the agent is aware of[2]. This notion is formalised by the following logical statement: $O_\alpha^\beta(p) \wedge B_\alpha(O_\alpha^\beta(p))$. Social demands motivate the agent to act but they may not be achievable and may even conflict with other agent duties; being autonomous, the agent may also decide that, according to circumstances, it is better not to fulfill a social demand and rather accept the corresponding sanction. Demands become intentions only if they are part of the preferred extension of the normative agent EDA model – see (Filipe, 2000 section 5.7) for details.
- *Intentions*: Whatever their origin (individual or social) intentions constitute a non-conflicting set of goals that are believed to offer the highest possible value for the concerned agent. Intentions are designated by some authors (Singh, 1990) as psychological commitments (to act). However, intentions may eventually (despite the agent sincerity) not actually be placed in the agenda, for several reasons:
  - They may be too abstract to become directly executed, thus requiring further means-end analysis and planning.
  - They may need to wait for their appropriate time of execution.
  - They may be overridden by higher priority intentions.
  - Required resources may not be ready.

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

6 INFORMATION FIELDS FOR INTER-SUBJECTIVE REPRESENTATION OF SOCIAL OBJECTS

Following Habermas (1984) we postulate the existence of a shared ontology or inter-subjective reality that defines the social context (*information field*) where agents are situated (Filipe, 2003). This kind of social shared knowledge is not reducible to individual mental objects (Conte and Castelfranchi, 1995). For example, in the case of a commitment violation, sanction enforcement is explicitly or tacitly supported by the social group to which the agents belong, otherwise the stronger agent would have no reason to accept the sanction. This demonstrates the inadequacy of the reductionist view.

Once again, we look at human organisational models for designing multi-agent systems; for example, contracts in human societies are often written and publicly registered in order to ensure the existence of socially accepted, and trusted, witnesses that would enable the control of possible violations at a social level. Non-registered contracts and commitments are often dealt with at a bilateral level only and each concerned agent has its internal contract copy. This observation suggests two representational models:

- A distributed model: Every agent keeps track of social objects in which that agent is involved and may also be a witness of some social objects involving other agents.
- A centralised model: There is an Information Field Server (IFS) that has a social objects database, including shared beliefs, norms, agent roles, social commitments, and institutions.

The distributed model is more robust to failure, given the implicit redundancy. For example, a

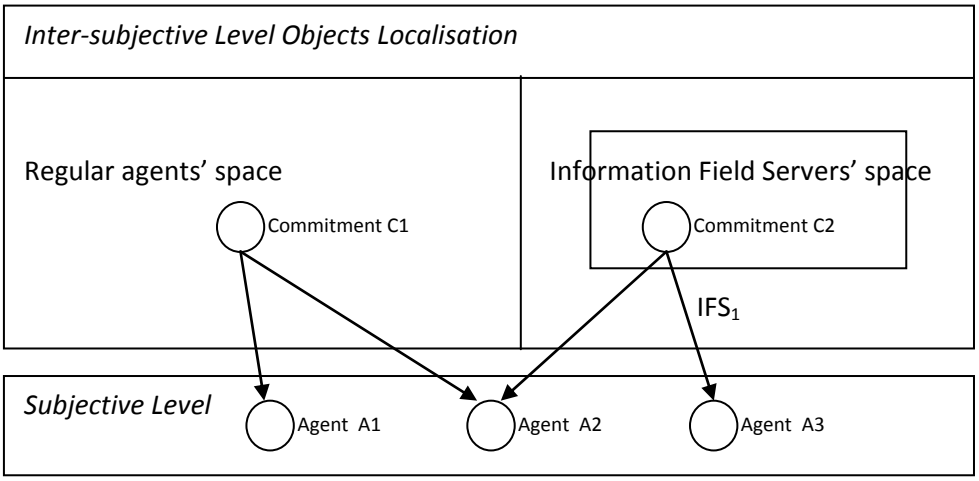

Figure 4: Social objects representation at inter-subjective level and their usage.

contract where a number of parties are involved is kept in all concerned agents' knowledge bases, therefore if an agent collapses the others can still provide copies of the contract. It is also more efficient assuming that all agents are honest and sincere; for example, commitment creation and termination involved in business transactions would not need to be officially recorded – a simple representation of a social commitment at the concerned agents EDA model would suffice.

However, since these assumptions are often unrealistic, the distributed model cannot completely replace the role of certified agents, trusted by society to keep record of shared beliefs and social commitments. We assume here that these social notions are part of the ontology that is shared by all members of an information field; that's why we call these trusted repositories of the shared ontology "Information Field Servers". These servers have the following characteristics:

- Different information fields must have different IFS because the shared ontology may differ among specific information fields.
- Each information field may have several non-redundant IFS, each representing a small part of the shared ontology.
- The robustness problems of IFS are minimized by reliable backup (redundant) agents.

Considering the empirical semiotics level, communication bandwidth is another relevant factor to consider: if all social objects were placed in central IFS agents these might become system bottlenecks.

A conceptual problem that exists but is not in the scope of this paper is related to the representation of social objects resulting from the interaction of agents belonging to different information fields. Possible solutions range from the unification of the different conceptual frameworks to the creation of new information fields where the ontology is constructed from a continuous meaning negotiation process via the interaction of the concerned agents.

In figure 4 the architecture of the inter-subjective level is depicted with respect to the localisation of social objects in addition to an example showing how social objects are used at the subjective level.

Commitments are first class objects, which can be represented either in the agents' EDA models (which we designate as the agents' space) or in the IFS' EDA model (which we designate as the Information Field Server's space). In the example above, agents A1 and A2 have only an internal representation (in each EDA model) of a shared commitment C1, whereas Agents A2 and A3 do not have an internal representation of commitment C2 because this commitment is represented in IFS1. All agents A2 and A3 need is a reference (i.e. a pointer) to that shared commitment, although for implementation reasons related to communication bandwidth and efficiency copies may be kept internally.

7 CONCLUSIONS AND FUTURE WORK

The EDA model described here is based on the organisational 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 toward a norm-based social interaction in organisations.

The EDA architecture integrates a number of ideas gathered 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 social modeling is the notion of 'commitment'. 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 (Brewka, 1994).

In a multi-agent environment the mutual update of agents' EDA models is essential as a result of perceptual events, such as message exchange. There is also the possibility of using shared spaces such as the information fields mentioned in section 6, which exist at an inter-subjective level. 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

Belnap, N., 1991. *Backwards and Forwards in the Modal Logic of Agency*. Philosophy and Phenomenological Research, vol. 51.

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.

Castelfranchi, C., 1993. Commitments: from Individual Intentions to Groups and Organisations. In Working Notes of AAAI'93 Workshop on AI and Theory of Groups and Organisations: Conceptual and Empirical Research, pp.35-41.

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

Habermas, J. 1984. The Theory of Communicative Action: Reason and Rationalization of Society. Polity Press. Cambridge.

Filipe, J., 2000. Normative Organisational Modelling using Intelligent Multi-Agent Systems. Ph.D. thesis, University of Staffordshire, UK.

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

Filipe, J., 2002. A Normative and Intentional Agent Model for Organisation Modelling. Proceedings of the 3rd Workshop on Engineering Societies in the Agents World, ESAW 2002. Madrid, Spain, também publicado como post-proceedings em: Petta, Tolksdorf and Zambonelli (Eds), Engineering Societies in the Agents World III, Springer-Verlag, LNAI 2577, pp. 39-52, ISBN 3-540-14009-3.

Filipe, J., 2003. Information Fields in Organization Modeling using an EDA Multi-Agent Architecture. AAAI Spring Symposium 2003 on Agent-Mediated Knowledge Management (AMKM-2003). Stanford, USA.

Jennings, N., 1994. *Cooperation in Industrial Multi-Agent Systems*, World Scientific Publishing, Singapore.

Liu, K., 2000. Semiotics Information Systems Engineering. Cambridge University Press. Cambridge.

Liu, K., K. Nakata and C. Harty, 2010, *Pervasive informatics: theory, practice and future directions* Intelligent Buildings International, Volume 2, Number 1, 2010 , pp. 5-19(15).

Peirce, C. (1931-1935) Collected papers of Ch. S. Peirce, 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.

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.

Stamper, R., 1973. Information in Business and Administrative Systems. John Wiley & Sons.

Stamper, R., 1996. Signs, Information, Norms and Systems. In Holmqvist et al. (Eds.), Signs of Work,

Semiosis and Information Processing in Organisations, Walter de Gruyter, Berlin, New York.

Stamper, R., K. Liu, M. Hafkamp and Y. Ades, 2000. Understanding the Roles of Signs and Norms in Organisations – a Semiotic Approach to Information Systems Design, Behavior and Information Technolog, 19(1), 15-27.

Stamper, R. (2000). Information Systems as a Social Science: An Alternative to the FRISCO Formalism. In *Information System Concepts - An Integrated Discipline Emerging,* E.D. Falkenberg, K. Lyttinen, A. A. Verrijn-Stuart (eds). Proc. ISCO 4 Conference, Kluwer Academic Publishers, pages 1-51.

von Wright, G., 1951. Deontic Logic. Mind, 60, pp.1-15.

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

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 Organisational Issues for Enterprise Integration. In Proceedings of ICEIMT'97, 529-538.