Social Commitment Policies as a New Formalism for the Autonomous Agent Social Metaphor[1]

Javier Soriano, Fernando Alonso, Genoveva López

Technical University of Madrid, Campus de Montegancedo S/N,
28660 Boadilla del Monte, Madrid, Spain
{jsoriano, falonso, glopez}@fi.upm.es

Abstract. In this paper, we propose a new domain-based organisational model as an extension of the traditional role-based models for formal agent society specification. The concept of social commitment policies, which govern various aspects of group behaviour, including penalties for individual acts in the context of group activities, is at the heart of this model. The concept of penalty is presented as an important concern in light of the assumption that, even conceived as society members, agents are still autonomous, so they can violate the established social commitments. We also show how deontic, action and temporal (DAT) modal logic will be used in the formal specification of the different types of policies and their related concept of violation. This formalisation will allow agents to reason deliberatively about the existing social norms, the responsibilities proper to their position and the consequences of violating these responsibilities. Finally, attention is paid to the formalisation of policy dynamics by means of illocutionary logic.

1. Introduction

Despite their intrinsic autonomy, agents tend not to be self-sufficient, which means that generally they will not be able to act exclusively on the basis of their needs and individual capabilities. On the contrary, they will have to relate to and cooperate with the other agents in their environment to achieve their own individual goals. If the environment is open, they will also need to address new goals, which we will term social goals, whose purpose is to assure that the multi-agent system (MAS) behaves coherently as a whole. The social goals will, in one way or another, influence the behaviour of the agents participating in the MAS.

This approach has led to the appearance of numerous proposals directed at easing and promoting interoperability between heterogeneous agents acting in open environments. A variety of abstractions and models have also been developed to capture the inherent sociability of autonomous agents participating in MAS, thereby facilitating the analysis and design of what are known as agent societies. Proposals founded on the extension of other traditional analysis and design techniques, based on object orientation [3], for example, are not very well suited for agent society conceptualisation because the abstractions involved are a sizeable distance apart. Compositional methods, traditionally used in object-oriented software architectures [4], are not very applicable to the definition of the communicative, organisational and role-oriented structures that characterise agent societies.

[1] This work is being supported by the Spanish Ministry of Science and Technology –CICYT (TIC2001-3451).

Over the last few years, a series of specific modelling techniques and development methodologies have been proposed for the agents paradigm [1], some of which began to conceive MAS as artificial organisations. Most of these modelling techniques define an organisation as a collection of roles [5] and a set of interaction protocols that occur in these roles [6], although they do not provide higher level organisational abstractions, like positions, organisational structures, norms, social laws or organic rules, useful for the complex task of modelling an agent society. Others, although they do make contributions in this respect [7], are not very suited for application in open environments [8].

Recently, there has been an increasingly widespread trend towards developing organisational abstractions that can be used to view multi-agent systems as open organisations. The social metaphor and role-based models have started to be generally accepted as a natural way of conceiving multi-agent systems. Additionally, the early proposals have matured tremendously and were well established in [10]. The conception of organisations made in [10] as work frames that accommodate a topology based on the concept of role has been a revulsive for the adoption of this metaphor. The reason is that such a model can be used to constrain the behaviour of the agents in an organisation through a series of rules, from which the structures existing in the organisation can be deduced and some behaviour patterns can even be extracted or exploited.

Less progress has been made with regard to the formalisation of these models. Although there are a variety of logical formalisms for modelling social behaviour [12][13], few researchers have applied them in their modelling proposals ([14][15] are noteworthy exceptions). Such formalisation is essential, however, as it will enable agents to be conscious of and reason about their position in the society.

The remainder of the paper is structured as follows. After introducing the most prominent aspects of our proposed organisational model in section 2, we formalise and classify the different types of social commitment policies in section 3 and discuss the interest of their application not only as an organisational abstraction but also as a mechanism for specifying the expected behaviour of the agent society. This section also shows how deontic, action and temporal modal logic is used to capture and formalise our proposed taxonomy of policies and the new concept of violation. Finally, in section 4 attention will be paid to policy dynamics, for which purpose we introduce the action of delegation of duty and propose its formalisation by means of speech act theory and illocutionary logic.

2. The Organisational Model

As stated above, not generally being self sufficient, agents have to relate to other agents. They are, therefore, thought of as members of an organisation whose model must be taken into consideration. This organisational model defines an agent society whose key concepts are (1) roles, (2) the relationships established between these roles, (3) the social commitment policies that govern these relationships and (4) the social structure that arises as a result. As the society is open, these concepts are dynamic. The role played by an agent reflects what that agent is expected to do in the organisation. It is defined in terms of the different activities, responsibilities and constraints that give an agent a well-defined position in the organisation with a set of associated expected behaviours (e.g. patterns of interaction). The organisational model is an aid for describing the configurations in which these roles act, thereby characterising agent societies.

This definition fits the concepts of classical role theory. However, our model extends this theory to cover concepts like job. Like a person, an agent can take more than one role to do its job within one organisation. The responsibilities associated with a job represent the sum of the responsibilities associated with each role required for correct job perform-

ance. This is not the case for example with permissions. Generally speaking, permissions should not be accumulated, as this could cause problems.

Each agent in the society assumes full responsibility for autonomously carrying out its individual tasks. However, social responsibilities are a different kettle of fish. The need to ensure that these social responsibilities are fulfilled means that both the relationships required for these actions to be completed and the social commitment policies governing such relationships (norms, obligations, permissions, etc.) have to be identified. These policies have to be respected and/or enforced by the society, through agency mechanisms and services, if it is to work properly according to the generally expected behaviour. It is these social-commitment policies that govern social relationships and ensure an organic and socially responsible behaviour in individualistically designed agents. The idea of ensuring organic and socially responsible behaviour during agent design is evidently not appropriate in a dynamic, heterogeneous environment like a society, where, just to give an example, interaction patterns are established dynamically. This is the reason why we propose in this paper a declarative approach to policy specification. For a further description of our organisational model, see for example [2].

Role-based models are insufficient for identifying agent groups covered by a given policy (rule, responsibility or authorisation) or between which there is a normative relationship (power, authorisation, etc.), as the same policy or normative relationship tends to affect more than one role. Therefore, we propose extending the role-based model to a domain-based model.

The domain concept extends the role concept in order to specify the organisational structure of the agent society. Accordingly, a domain identifies both functional and structural agent groups and, consequently, express concepts as important as (1) position, which can involve one or more roles between which there is a given functional relationship, (2) the organisational unit in which these positions occur and (3) the organisation itself.

For the purpose of combining domains and expressing the set of roles to which a policy applies, we introduce a declarative language ℓ_{dom} to build domain expressions (any domain expression is also a domain). This language should be semantically rich enough to be able to express the fact that a subdomain is not necessarily a subset of the parent domain, which means that the agents within this subdomain could be indirect, not necessarily direct, members of the parent domain. This is a key distinction for determining the scope of a given policy (a single subdomain or all subdomains, etc.). The definition of the language is:

$$\ell_{dom} = r \in \Re \mid d^{\circ} \mid d^{\bullet} \mid d^{n} \mid \delta_1 + \delta_2 \mid \delta_1 \oplus \delta_2 \mid \delta_1 \times \delta_2 \mid \delta_1 \otimes \delta_2 \mid \delta_1 - \delta_2 \qquad \textbf{(1)}$$

where *r* denotes a role, d° represents the set of all the members of a domain that are not themselves domains (domain entries representing subdomains), whereas $d^{\bullet}$ also includes the subdomain members recursively nested in *d* and d^{n} delimits nesting to depth n. $\delta_1 + \delta_2$ is the set that contains all the different members of the union of the sets resulting from evaluating the domain expressions δ_1 and δ_2. Similarly, $\delta_1 \times \delta_2$ represents the intersection, $\delta_1 \oplus \delta_2$ represents the disjoint union ($\delta_1 \oplus \delta_2 \equiv (\delta_1 + \delta_2) - (\delta_1 \times \delta_2)$), and $\delta_1 - \delta_2$ represents the difference ($\delta_1 - \delta_2 \equiv \delta_1 - (\delta_1 \times \delta_2)$).

3. Social Commitment Policies

In the process of conceptualising an agent society, it is necessary to capture and model the social conscience of the agents that participate in the society. Our proposal for modelling this social conscience focuses on the concept of social commitment policy. In the context

of this paper, a policy is defined as a formal and declarative specification, derived from the global goals of an agent society, of a rule that defines alternative behaviours of the agents that interpret them. Four interrelated deontological categories have been used to formalise this concept: norms, obligations, responsibilities and authorisations.

3.1. Norms

Generally, a norm represents a specific behaviour that is seen as beneficial by either a particular organisation or the society as a whole. The formal specification of a set of norms will make it possible to influence the behaviour of agents that adhere to these norms. This will help to standardise and coordinate the behaviour of the different agents participating in an open society, enabling them to anticipate the behaviour of the others, thereby improving the processes of agent cooperation, negotiation and interaction in heterogeneous environments [16].

Norms do not restrict agent autonomy, although they can influence their intentions, improving coordination and making it more effective and efficient. An agent can decide whether or not to adopt a norm solely on the basis of the benefits it sees in doing so, behaving in a manner that we will call socially responsible.

A new modal operator N is introduced to specify norms[2]: $N^D(\alpha \mid \psi)$ to indicate that the norm in domain D is to perform action α when ψ is true.

A typical example of a norm within an agent society would be:

$$\forall a_1, a_2 \in S \quad [request(a_1, a_2, \_)]N^S(response(a_2, a_1, \_) \mid true) \qquad \textbf{(2)}$$

which makes use of dynamic logic [17] and the new operator *N* to express that all agents should respond to the requests made by another agent whenever possible. S is a domain expression and represents the society as a whole.

3.2. Obligations

Obligations represent responsibility assignation policies for action (non-) performance. They are, therefore, an explicit mechanism for influencing agent behaviour without cancelling out their autonomy. The explicit representation of obligations and the existence of a suitable framework for their performance will allow agents to interpret the responsibilities they have towards the organisation in which they act and reason deliberatively about these responsibilities.

Like norms, obligations can also influence agent behaviour. However, a distinction should be made between the two concepts: whereas norms place no constraints on agent autonomy, as they merely report standard behaviours that are socially beneficial, obligations do restrict agent autonomy, insofar as they represent their responsibilities within the organisation in which they perform their activities. This loss of autonomy is for the benefit of a more coherent and balanced behaviour of the organisation as a whole, which we will call organic behaviour.

A new modal operator *O* is introduced to express obligations: $O^D_{t,s}(\alpha \mid \psi)$ to indicate that the role t is obliged to perform action α when ψ is true. Each role is associated with the domain D in which it participates to make it easier to organise and analyse both the roles and their associated obligations. s denotes the role that has originated the obligation (sub-

[2] The full model includes another acceptation for each modal operator, e.g. $N^D(\phi \mid \psi)$ to express the fact that the norm in domain D is that the formula ϕ is true when ψ is true.

ject). If the obligation is imposed by the organisation that domain D represents or by the society itself (D=S), only $O_i^D(\alpha \mid \psi)$ will appear. D is a domain expression that represents the organisation/society responsible for enforcing the penalty in the event of violation.

Generally, organic behaviour is also very much influenced by social norms, which means that it has a socially responsible behaviour component. Indeed, our model defines an organic agent as an agent whose behaviour is influenced by a set of obligations that represent its responsibilities within the organisation in which it participates and whose design predisposes it to adhere to the norms established by the society in which it acts. In this respect, a typical social norm in any agent society is a norm that advises that the obligations should be met whenever possible, as expressed in (3):

$$\forall a \in S, d \in D(S) \quad N^S\left(\alpha \mid O_a^d(\alpha \mid \psi) \wedge \psi\right) \tag{3}$$

According to traditional deontic logic [18], there are three categories of obligations: the obligation, which represents that an agent is responsible for performing a given action, interdiction $I_{t,s}^D(\alpha \mid \psi)$, which expresses that an agent is responsible for not performing a given action, and permission $P_{t,s}^D(\alpha \mid \psi)$, which explicitly expresses the absence of a given interdiction. Interdiction and permission can be expressed in terms of obligation using the concept of action negation taken from action logic. Accordingly, we have that $I(\alpha) \leftrightarrow O(\overline{\alpha})$ and $P(\alpha) \leftrightarrow \neg I(\alpha) \leftrightarrow \neg O(\overline{\alpha})$, but a new modal operator is introduced for each one with a view to improved expressivity. Some authors even introduce another operator to represent the concept of discretionality or freedom $D(\alpha) \leftrightarrow \neg O(\alpha)$.

The separation of duties technique is a paradigmatic example of an obligatory policy that involves the joint use of permissions and interdictions to specify normative behaviours in the event of potentially conflicting activities. In the dynamic separation of duties, all agents acting in the same position are initially permitted to perform a set of potentially conflicting actions. However, after one of these actions has been taken, they are prohibited from taking the other conflicting actions. The following constraints show how the agents acting as managers in the accountants domain (A) are permitted to issue (4) and approve (5) cheques (C), but one and the same agent is prohibited from signing and approving the same cheque.

$$P_M^{Accountants}\left(issue(m:M, c:C) \mid m.id \neq c.approverID\right) \tag{4}$$

$$P_M^{Accountants}\left(approvePayment(m:M, c:C) \mid m.id \neq c.issuerID\right) \tag{5}$$

3.2.1 Violation of obligations

If the autonomy of an organic agent is to be preserved, situations in which the obligatory normative constraints can be violated should be accounted for. The deontological approach affords some degree of freedom of choice, and consequences can be explicitly associated with decision making. The decision to obey a norm is subject to a deliberative decision-making process driven exclusively by the weighting of the cost/benefit ratio (individual and social) associated with adhering to the norm. However, the fulfilment of an obligation is subject to a deliberative decision-making process also driven by a weighting of the consequences derived from violating this obligation, that is, the associated cost of the penalty incurred by breaching the obligation. On the other hand, the interpretation of authorisations (positive and negative) is not subject to any deliberative decision-making process, as they can never be violated.

The weighting ratio of social benefit and the weighting ratio of consequences derived from the violation of obligations will define a partial order between norms and obligations, respectively. Obligation-based reasoning is simpler than norm-based reasoning.

This is because the obligations have explicit penalties associated with their violation, whereas the consequences derived from not adhering to a norm are indirect and longer term.

By reducing deontic logic to dynamic logic [11], the concept of violation can be introduced as a new predicate in the formalisation of deontological operators, as follows:

- $O_{ij}(\alpha(i)) = [\overline{\alpha(i)}]V_{ij}^{k}(\alpha)$ expresses that an action is obligatory if its non-performance does not lead to a violation,
- $I_{ij}(\alpha(i)) = [\alpha(i)]V_{ij}(\alpha) \equiv O_{ij}(\overline{\alpha(i)})$ expresses that an action is interdicted if its performance leads to a violation and, finally,
- $P_{ij}(\alpha(i)) = \neg[\alpha(i)]V_{ij}(\alpha) \equiv \neg O[\overline{\alpha(i)}] \equiv \neg I_{ij}(\overline{\alpha(i)})$ expresses that an action is permitted if its performance does not lead to a violation.

As can be seen from the above definitions, one or more states of violation $V^{k}(\alpha)$, one for each of the reasons why the (non-)performance of the action is incorrect, are defined for each action. The violation states are defined and the reasons for the violation are represented in the predicate in order to express the corrective action.

The following equations illustrate the use of the violation predicate:

$$[borrow(L, r:R, b:B)]O_{L,R}^{Library}(return(r:R, L, b:B))_{before(21days)} \quad (6)$$

$$V_{L,R}^{Library}: return(r:R, B) \rightarrow I_{L,R}^{Library}[borrow(r:R, B)] \quad (7)$$

$$V_{L,R}^{Library}: return(r:R, b:B) \rightarrow [return(r:R, b:B)]P_{L,R}^{Library}[borrow(r:R, B)] \quad (8)$$

L represents the role Librarian, R the role Reader and B the domain Book. (6) expresses that, after borrowing a book, the reader will have to return it within 21 days. If the reader does not return the book by the return date, (7) determines that the reader will commit a violation (on the basis of the axiomatic definition of obligation) and will be prohibited from requesting further loans until he or she has returned the book, at which point the penalty will be lifted as shown in (8).

This example also shows how, despite the fact that obligations refer to roles or domains, they often need to be linked to a particular agent or element belonging to this role or domain. The operator ':' represents this link.

Even though, as we have just seen, the specification of some obligations and violations calls for some of the roles involved to be linked to particular instances of these roles, obligations should, nevertheless, be referred to domains in order to improve their distribution.

3.2.2 Obligations and Responsibility

One concept that plays an important role in modelling the social conscience of agents is responsibility. Within our agent work frame, we equate the concepts of responsibility and obligation, i.e., if an agent has the obligation to perform an action, it is said to be responsible for this action. When an agent fails to perform an action that it was under obligation to (i.e. for which it is responsible), it is its mission to put the situation right, which means that this agent becomes the subject of the violation, i.e. of the action or situation caused by the violation.

Nevertheless, the concept of responsibility becomes of special interest when modelling a situation in which an agent has entered into an obligation but really has another subordinated agent performing the action. Without the concept of responsibility, the obligation entered into by the subordinated agent always refers to this agent, which means that it would be the subject of the penalty derived from the non-performance of the action in question. As we will see later when studying the dynamics of obligations, the concept of responsibility can be used to separate the subject of the penalty from the subject of the ob-

ligation and associate it with the subject of the responsibility. The first agent is then said to have delegated the duty (i.e. the obligation) for performing this action, but has not delegated the responsibility (i.e. maintains the control).

3.3 Authorisations

Authorisation policies define the set of requests that can be accepted by a system element on the basis of the identity of the applicant and the circumstances (conditions) in which the request is made.

Despite appearances, the concepts of positive and negative authorisation are not equivalent to the concepts of permission and interdiction, respectively. The main difference lies in the fact that the obligations (i.e. the permissions and interdictions) are interpreted by the subjects and, therefore, can be violated, whereas authorisations are interpreted by access control mechanisms linked to the target and cannot be violated. Additionally, authorisations do not restrict agent autonomy, although they do influence decision-making processes, as they are prerequisites for performing given actions.

A negative authorisation can be interpreted as an interdiction with infinite cost associated to its violation or as an interdiction with a fail action associated to its penalty.

Two new modal operators A+ and A- are introduced to represent the concept of authorisation. $A+_s^D(\alpha \mid \psi)$ expresses that the role s is authorised to request an object belonging to the domain expressed by D to perform action α when ψ is true. There will be a domain controller associated with D in charge of validating requests, which will, therefore, be responsible for enforcing the authorisations. $A-_s^D(\alpha \mid \psi)$ expresses a negative authorisation.

4 Obligations dynamics

There are persistent obligations, i.e. obligations that do not disappear after performance of the obligatory action. This is the case of the obligations assigned to their roles by an organisation that model intrinsic behaviours of those roles. Accordingly, an organisation TA acting as a travel agency may require all the agents acting as its sellers to respond to all the queries made by trusted client agents:

$$[query(Client, Seller, \_)]O^{TA}_{Seller}(reply(Seller, Client, \_) \mid trusted(s)) \quad \textbf{(9)}$$

This is an example of an obligation that does not disappear after responding to the first query, but continues as an intrinsic responsibility of the seller role.

However, many obligations do disappear after the performance of the obligatory action, as is the case of contextual obligations created by an agent with respect to another through a relationship of power or temporary authorisation. These two relationships have to be modelled to be able to formalise the obligations dynamics.

The relationship of power is the formal mechanism that enables new obligations (and beliefs) to be created and deleted. There is a relationship of power between a role i and a role j with respect to an action α, if an agent acting as i is empowered to order the agent acting as j to perform α. The predicate $\Pi^{\alpha}_{i,j}$ is introduced to model this relationship.

The relationship of power between two roles is usually persistent, as it represents the normative structure of the organisation and this does not tend to change much over time.

Let us look at one use of the relationship of power to express an obligation entered into as a result of a relationship of power between two roles:

$$\Pi_{s,t}^{request(pay(t,s,q))} \rightarrow [request(s,t,pay(t,s,q)]O_t^D(pay(t,s,q) \mid trusted(s)) \quad \textbf{(10)}$$

(10) expresses that the role s will have to pay the amount q required by t every time it is asked to by t, as there is a relationship of power between t and s for making payment requests.

However, if what (11) should express is that t could ask s to pay the amount q once for given services provided, the above representation would be incorrect, as it would then be necessary to annul the relationship of power, and this is a persistent relationship as discussed above.

$$O_t^D(pay(t,s,q) \mid trusted(s)) \rightarrow ([pay(t,s,q)] \rightarrow \neg\Pi_{s,t}^{request(pay(t,s,q))}) \quad \textbf{(11)}$$

It is better to introduce a temporary relationship of power that we will call relationship of temporary authorisation. This relationship is established for a definite time period, upon mutual agreement and subject to certain constraints. It is modelled by the predicate $\Delta(i,\alpha)$ to indicate that i is authorised to perform α. The semantics of this relationship differs from the semantics of the relationship of power in that it is temporary. The relationship of authority is established for a given time period upon mutual agreement (i.e. with some constraints). The main quality of this relationship is that it establishes the obligations dynamics on the basis of possible withdrawal. The above example would be modelled using the relationship of authorisation as follows:

$$\mathrm{A}_{s,t}^{request(pay(t,s,q))} \rightarrow [request(s,t,pay(t,s,q)]O_t^D(pay(t,s,q) \mid trusted(s)) \quad \textbf{(12)}$$

$$O_t^D(pay(t,s,q) \mid trusted(s)) \rightarrow ([pay(t,s,q)] \rightarrow \neg\mathrm{A}_{s,t}^{request(pay(t,s,q))}) \quad \textbf{(13)}$$

4.1 Creation, Derogation and Distribution of Policies by means of Speech Acts

According to the original ideas of speech act theory [9][15], the communication between agents must not be seen merely as information transmission, but also as actions that alter the mental state of the agents involved in the system (i.e. its beliefs, desires and intentions).

An important aspect of speech act theory is what can be inferred as a result of a speech act. In this respect, and for the purposes of this article, we will focus on speech acts that are performed to say something, for example, ask or reply to a question, provide information, etc. [9] decomposes each speech act into an utterance act, which specifies the type of action (requesting, warning, informing, etc.) performed, and propositional act, which specifies the details of the action (what is being asked, what information is being given, etc.). [15] further decomposes the propositional act into an act of predication and a reference act, although this distinction is not important in the context of this article.

The relationships of power and temporary authorisation, along with illocutionary logic, represent the basis for entering new obligations, beliefs and permissions into the system by means of communication acts between agents. The existence of a relationship of power or authorisation is the condition under which a prescriptive speech act is considered to have directive force. In this respect, obligations can be created, eliminated, distributed and delegated and authorisations can be applied for, granted and delegated, thereby creating a dynamic environment for the establishment and derogation of authorised norms.

For the sake of brevity, we will only consider directive illocutionary force, as this will be used to formalise the concept of delegation of duty. Directives (DIR) serve to request the performance of an action. They will be used, in the context of a relationship of power or temporary authorisation, to create and report an obligation to the speech act target.

$$([DIR_{\Pi}(i,j,\alpha)]O_{j,i}(\alpha)) \leftarrow \Pi_{i,j}^{\alpha} \quad (14)$$

$$([DIR_{\Delta}(i,j,\alpha)]O_{j,i}(\alpha)) \leftarrow \Delta_{i,j}^{DIR(i,j,\alpha)} \quad (15)$$

Axiom (14) expresses that the existence of a relationship of power between i and j is sufficient for a directive speech act issued by i to generate an obligation in j. (15) expresses that the existence of a relationship of authorisation between the roles i and j requesting the performance of an action α is sufficient for this request to create the obligation of performing this action in j.

4.1.1 Delegation of duty

The delegation of duty models the act of entrusting the performance of an obligatory task to another agent. This section shows how the dynamics of delegating an obligatory task to another agent can be modelled by interchanging one or more imperative communication acts and correctly managing the violation predicate.

The delegation of an obligation by an agent means that this agent trusts the delegated agent to act in its benefit, which means that, for this to take place, there must be some sort of link between the delegating agent and the agent to which the task is delegated. In our proposal, this link is specified by the relationships of power and trust: the relationship of power is necessary to create an obligation in the agent to which the task is delegated, and the relationship of trust is used to choose this agent from possible candidates (an agent in respect of which there is a relationship of power). It follows from this approach that there must be an explicit link between the imperative communication acts and the normative positions established in the society role model.

After having delegated the duty, the agent maintains the responsibility, which means the agent delegating the duty must assure that its delegate agent acts appropriately. Despite of this, the delegating agent is liberated from the obligation because, in other case, it should intend to perform the task delegated.

In our proposal, the concept of responsibility is modelled by specifying to whom the penalty incurred as a result of the violation of the obligation to which this duty refers applies.

If the concept of responsibility were obviated, we could only express independent obligations like $O(k,i,\alpha) \rightarrow [\overline{\alpha}]V_{k,i}^{\alpha}$. This would mean that if i is an agent to which j has delegated an obligation $O(i,j,\alpha) \rightarrow [\overline{\alpha}]V_{i,j}^{\alpha}$, the agent that performs that task could be the agent held responsible in the event that the tasks were not completed successfully. The obligation of agent i disappears, but its responsibility needs to be maintained.

The concept of responsibility and its management by means of the delegation of duty can express $O(i,j,\alpha) \rightarrow [\overline{\alpha}]V_{j,k}^{\alpha}$, which is closer to reality, as the agent originally under obligation continues to be responsible for the successful completion of the task (is the subject of the violation, i.e. of the associated penalty), whereas the agent under obligation may not be responsible for performing this action. Otherwise, it would always be responsible with respect to the agent that delegated the task to it.

5 Conclusions

This paper has discussed a number of issues related to the specification and dynamics of multi-agent systems in open environments. The organisational concepts of agent roles and role models have now become an important research area in the field of agent-based systems. In this paper, however, we have introduced further organisational abstractions and

formalisms. We have presented the concept of domain as a means of specifying the organisational structures that arise as a result of the structural and behavioural relationships within an agent society (i.e. organisations, departments, positions, roles, etc.). We have formalised the different categories of social commitment policies that govern various aspects of society behaviour, including penalties for individual acts in the context of domain activities. The concept of penalty has been presented as an important concern in light of the assumption that agents, even conceived as society members, continue to be autonomous, so they can violate the established social commitments. We have shown how deontic, action and temporal modal logic will deal with our proposed taxonomy of policies and with the new concept of violation. Finally, attention has been paid to the policy dynamics, and its formalisation by means of illocutionary logic.

6 References

1. C. Iglesias, M. Garijo, and J. Gonzales. A survey of agent-oriented methodologies. In A. S. Rao, J.P. Muller and M. P. Singh, editors, Intelligent Agents IV (ATAL98), LNAI. Springer-Verlag, 1999.
2. F. Alonso, F. Fernández, G. López , P. Rojas and J. Soriano. Intelligent Virtual Agent Societies on the Internet. Intelligent Virtual Agents, LNAI. Springer-Verlag. 2001.
3. E. A. Kendall. Role modelling for agent system analysis, design, and implementation. In 1st International Symposium on Agent Systems and Applications. IEEE CS Press, Oct. 1999.
4. J. Odell, H. Van Dyke Parunak, and C. Bock. Representing agent interaction protocols in UML. In OMG Document ad/99-12-01. Intellicorp Inc., 1999.
5. J. Ferber and O. Gutknecht. A meta-model for the analysis and design of organizations in multi-agent systems. In Proceedings of the 3rd International Conference on Multi-Agent Systems (ICMAS 98). IEEE CS Press, June 1998.
6. Y. Demazeau and A. C. Rocha Costa. Populations and organizations in open multi-agent systems. In 1st National Symposium on Parallel and Distributed AI, 1996.
7. M. Wooldridge, N. R. Jennings, and D. Kinny. The Gaia methodology for agent-oriented analysis and design. Journal of Autonomous Agents and Multi-Agent Systems, 3(3):285-312, 2000.
8. F. Zambonelli, N. Jennings, A. Omicini, and M. Wooldridge. Agent-oriented software engineering for Internet applications. In A. Omicini, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors, Coordination of Internet Agents: Models, Technologies and Applications. Springer-Verlag, 2000.
9. J. A. Austin. How to Do Things with Words. Harvard University Press, 1962.
10. F. Zambonelli, N. R. Jennings and M. Wooldridge. Organisational Abstractions for the Analysis and Design of Multi-Agent Systems. In Agent Oriented Software Engineering, pp 185-193, 2000.
11. J. –J. Ch. Meyer. A different approach to deontic logic: Deontic Logic viewed as a variant of dynamic logic', Notre Dame Journal of Formal Logic 29: 109-136. 1989.
12. R. J. Wieringa and J.-J. Ch. Meyer. Applications of Deontic Logic in Computer Science: A concise Overview. In Deontic Logic in Computer Science, pp 17-40, John Wiley & Sons, Chichester, England, 1993.
13. I. Pörn. Action Theory and Social Science: Some Formal Models. Synthese Library, 120, Dordrecht: D. Reidel, 1970.
14. O. Pacheco and J. Carmo. A Role Based Model for the Normative Specification of Organized Collective Agency and Agents Interaction. To appear in Journal of Autonomous Agents and Multi-Agent Systems. Kluwer Academic Publishers. 2002.
15. J. R. Searle. A classification of illocutionary acts. Language in Society, 5:1-23, 1976.
16. F. Dignum, D. Morley, E. A. Sonenberg and L. Cavedon, “Towards socially sophisticated BDI agents”. In Proceedings of the ICMAS 2000, pages 111-118, 2000
17. D. Harel. Dynamic Logic. In D. M. Gabbay and F. Guenther, editors, Handbook of Philosophical Logic, 2, Reidel, 1984.
18. J. –J. Ch. Meyer and R. J. Wieringa. Deontic Logic: A Concise Overview. In Deontic Logic in Computer Science, pp 17-40, John Wiley & Sons, England, 1993.
19. G. H. von Wright. Deontic Logic. Mind, 60:1-15, 1951.