An Agent Oriented Ontology of Social Reality

Guido Boella[1]
Leendert van der Torre[2]
[1] *Dipartimento di Informatica - Università di Torino - Italy. E-mail: guido@di.unito.it*
[2] *CWI Amsterdam and TU Delft - The Netherlands. E-mail: torre@cwi.nl*

Abstract. In this paper we introduce an ontology based on the notion of agent to represent and reason about social reality. We model social constructions as agents, for example, groups, organizations, normative systems, and roles, and we attribute mental attitudes to them. Roughly, we define obligations or regulative norms as goals of the normative system, constitutive norms as beliefs of the normative system, joint, shared, mutual and social beliefs, desires and goal as beliefs, desires and goals of group, responsibilities of an agent as goals of the role he plays, and the required expertise of an agent as beliefs and actions of the role he plays. In this way, we achieve a uniform framework for a large variety of concepts using a small vocabulary, and, in particular, basing it on notions, like mental attitudes, which are commonly used in agent theories. The proposed ontology is modelled using a description logic.

1 Introduction

Concepts from social reality have been recently introduced in multiagent systems (MAS) to coordinate the behavior of the different and possibly heterogeneous agents and to control the behavior emerging from their interaction [1]. For example, [2, 3, 4] address the problem of coordinating the agents of a MAS who want to achieve the same goal by forming a group. Another solution to cope with the control problem is introducing obligations and permissions, i.e., regulative norms, to provide motivations to agents; in [5] we extend MAS to normative multiagent systems (NS) who monitor and sanction violations of agents and introduce a legal classification of reality by means of constitutive norms [6]. Moreover, normative systems are not static entities which do not change with time. Rather, they introduce new norms and repeal outdated ones and allow also ordinary agents to become "private legislators" [7]; in particular, agents can create new obligations by means of contracts [8]. Finally, organizational concepts can be introduced to deal with complex MAS where different responsibilities must be distributed to agents with different expertise: organizations can be articulated in sub-organizations (functional units, departments) and roles which describe the behavior expected from the actors playing those roles. Organizations build on normative multiagent systems since they use norms to control agents playing roles; in this sense, according to Ouchi [9], organizations can be considered as burocracies.

How to describe the concepts from MAS, from organizations and NS, and their relations? A MAS consists of agents, beliefs, desires and goals, *et cetera*, and organizations and normative systems consist of shared goals, norms, violations, sanctions, expertise, responsibilities, *et cetera*. Social reality in the multiagent field, however, still lacks of a precise conceptual

model. The traditional approach to model obligations and permissions is deontic logic which defines concepts and their relations by modal operators. For example, in standard deontic logic there are modal operators for obligation, permission, prohibition, which are related by axioms such as $F(p) = O(\neg p)$ and $P(p) = \neg F(p)$. An alternative approach is to model concepts as first order predicates. The relation between permissions and prohibitions can be given by $\forall x$: *permit*$(x) = \neg$*forbidden*(x), but formalizing the relation between *obliged* and *forbidden* is more problematic since 'the absence of x' has to be formalized, for example by a function *not*: *forbidden*$(x) =$ *obliged*(*not*(x)). A third approach studied in this paper is to use description logics [10], a family of decidable fragments of first order logic and the most popular formalism to describe concepts, used for example to analyze entity-relationship diagrams, represent ontologies in semantic web, *et cetera*. More precisely, description logics may either be seen as a branch of modal logic or as a branch of first order predicate logic.

In this paper, we study the use of description logic to build the ontology of social reality underlying a multiagent system we described in our previous papers, like [8, 11], *et cetera*, but at a level which abstracts away from the different formalisms we adopted in other papers, e.g., input/output logic [12], BOID architectures [13] or BDI$_{CTL}$ logic [14].

Our extension of MAS to social reality is based on the attribution of mental attitudes to normative systems, groups, roles, and organizations. Roughly, normative systems, groups, organizational structures including roles are all modelled as autonomous agents. For example, the expression "your wish is my command" states that what is commanded to MAS is related to what is wished by NS.

The research question of this paper is: how to use description logic to formalize social constructions as agents, and, in particular, the attribution of mental attitudes to these social constructions? We use here a simple description logic based on the OWL language, even if we are aware of the limitations of its representation power.

The attribution of mental attitudes to social constructions is based on our ongoing research while the idea of basing social reality on collective acceptance is inspired on the philosophical work of Searle [15] and Tuomela [16]. As we show in this paper, this intentional stance [17] leads to a crucial simplification of our ontology.

The layout of this paper is as follows. In Section 2 we present the main concepts from our papers. In Section 3 we introduce the ontology using a Description logic. In Section 4 we discuss related work. Finally, Conclusions end the paper.

2 Concepts

In this section we informally describe the concepts used in our papers on social reality:

A proposition is a description of part of the environment. We distinguish propositions whose truth value can be changed by the agent (actions, called decision variables) from parameters which can only be indirectly influenced by the agents.

A belief is an informational mental state expressing relations between propositions and it can be represented by a conditional rule.

A motivational attitude is a mental attitude expressing which propositions are the objectives of an agent. A preference relation is used by each agent to resolve the conflicts among motivational attitudes. Motivational attitudes, like desires and goals, are represented by conditional rules.

A rule is a description of the mental attitudes. A rule (in the simplified form we used in the other papers) links a set of propositions to a proposition.

An agent is an autonomous decision maker; it bases its decisions on its own beliefs, desires and goals. An agent is endowed with social abilities; first, it is able to adopt the goals of other agents as its own goals [18], e.g., it can adopt the goals of the normative systems if it is a respectful agent, the goals of the group it belongs to, if it is a cooperative agent or the goals of the role it is playing if it is a trusted agent; second, an agent is able to attribute mental attitudes to other agents, so both to have a profile of other agents and to create socially constructed agents like groups, normative systems, virtual communities and organizations; third, when it takes a decision it considers also the effects of its decision on the decisions of the other agents by using the profile they have of other agents: in this way, it can coordinate its action with those of the agents interfering in the same environment [19].

A group of agents coordinate their behavior in the achievement of their shared goals. In a group all the members take advantage from the satisfaction of the goals of the group, even if for different reasons. In our model a group is described as an agent constructed by its members, which attribute to it mental attitudes [4]. Its beliefs represent the conventions of the members, its desires the goals towards not incurring into costs and its goal the objective shared by the members. A group can achieve its purpose if its members are cooperative, i.e., it they adopts the desires and goals of the group as their own. Differently from an organization, a group is not a normative system and it is not structured into roles.

A normative system is a social construction a society of agents uses to achieve a certain social order. According to [20], a social order is a pattern of interactions among interfering agents "such that it allows the satisfaction of the interests of some agent". These interests can be a shared goal, a value that is good for everybody or for most of the members; for example, the interest may be to avoid accidents. So the agents attribute to the normative system, beliefs, desires and goals, and, in particular, the goal to autonomously enforce the conformity of the agents to the norms by means of sanctions [5, 21]. The need of sanctions is due to the fact that, differently with respect to a group, the agents of the society are not all assumed to have always the goals of achieving the social order the normative system is delegated to achieve.

An obligation is a goal of the normative system. A question has been raised why norms are usually not implemented explicitly in computer systems. An easy answer is that computer programs already model 'ideal' behavior. They must never violate the rules, just as they must never fail. This objection can be countered by Dignum's argument [22] that obligations can be violated because agents are autonomous. In a typical example, an agent has a desire to do otherwise and the desire is stronger than the obligation. Moreover, in order to deal with conflicts among norms, an agent must be able to drop some obligations in favor of others. Since agents must be motivated to stick to obligations, obligations are associated with other instrumental goals of the normative system: to consider behavior which does not conform to the obligation as a violation and to sanction violations [23].

A violation is recognized by the normative system when the behavior of an agent does not conform to the obligations addressed to it. As an action of the normative system, the recognition of a violation is not a consequence of the behavior of the agent, but an autonomous decision of the normative system. The normative system considers the behavior

of an agent as a violation since this is one of its goals: a goal which is instrumental to the maintenance of the social order.

A sanction is an action of the normative system which is applied to an agent when the normative system recognizes its violation. Sanctions can be positive (rewards) or negative: the precondition of their efficacy is that the addressee of the norm desires/fears them. Like violations, sanctions are the result of an autonomous decision of the normative system, which can be influenced by the violator not to apply them.

A permission is an exception to some obligations of the normative system which consider a behavior as a violation. As obligations are modelled in terms of desires and goals of the normative system, and in particular, of the goal to consider a given behavior as a violation, permissions are modelled as goals of the normative system not to consider the given behavior as a violation. In this sense, permissions presuppose obligations to be meaningful [24].

An institutional fact is a concept introduced by a normative system to give a legal classifications of reality; for example, marriages, properties and money are institutional facts. An institutional fact is made true by a constitutive rule.

A constitutive rule is a belief of the normative system about the legal consequences of propositions. Searle [25] argues that there is a distinction between two types of rules, a distinction which also holds for formal rules like those composing normative systems:

> Some rules regulate antecedently existing forms of behaviour. For example, the rules of polite table behaviour regulate eating, but eating exists independently of these rules. Some rules, on the other hand, do not merely regulate an antecedently existing activity called playing chess; they, as it were, create the possibility of or define that activity.The institutions of marriage, money, and promising are like the institutions of baseball and chess in that they are systems of such constitutive rules or conventions ([25], p. 131).

For Searle, institutional facts like marriage, money and private property emerge from an independent ontology of "brute" physical facts through constitutive rules of the form "such and such an X counts as Y in context C" where X is any object satisfying certain conditions and Y is a label that qualifies X as being something of an entirely new sort. E.g., "X counts as a presiding official in a wedding ceremony", "this bit of paper counts as a five euro bill" and "this piece of land counts as somebody's private property".

In our model, constitutive norms are represented as beliefs of socially constructed agents. Since beliefs are described by conditional rules, they express the relations between propositions and the legal classification of reality in terms of institutional facts [6].

Moreover, in [6] we consider also constitutive rules which do not create only institutional facts, but which have an effect on the norms of a normative system; this is possible since they express what counts as an action which modifies the mental state of the normative agent.

A power of an agent in an organization or normative system is an action of the agent which creates new obligations or even new constitutive rules by modifying the mental state of the organization or the normative system. A power achieve its effect by means of some constitutive rule. Powers are at the basis of the definition of contracts [8].

An organization is an agent socially constructed by the members of a MAS to coordinate their behavior. The organization is collectively attributed by them mental attitudes: its beliefs represent the statute of the organization, its desires and goals its objectives, for example to maximize profits. The organization is structured in suborganizations (functional areas, departments) and roles assigned to agents. In order to make those agents stick to their role they are subject by the organization to obligations. The organization is, thus, a normative system. At the same time, the organization, as an agent, can be subject to obligations too: in juridical terminology, it is a legal person [8].

A role is a description of the expected behavior of an agent in an organization and thus it is defined as an agent constructed by an organization by attributing to it mental attitudes. The set of roles constitute the structure of an organization. The beliefs of the role represent its expertise and the goals its responsibilities. Roles are played by agents which have to act as if they had the beliefs attributed to the role and to adopt the desires and goals of the role. An agent can play several roles. Roles specify also the obligations which are posed on the agent playing them [26].

Finally, roles are used also to structure normative systems to keep separate the executive, legislative, and juridical functions [27].

3 Formalization of the ontology

We have chosen to formalize our ontology using the language OWL based on the DAML+OIL language. For space reasons, we present it here in a standard description logic like syntax [10].[1]

- If ϕ and ψ are classes, then so are $\phi \sqcap \psi$ and $\phi \sqcup \psi$. $\phi \sqcap \psi$ is the application of the *intersectionOf* constructor of OWL to the two classes and $\phi \sqcup \psi$ is the *unionOf* constructor.
- If ϕ is a class and r is a property, then $\exists r.\phi$ is a class and $\forall r.\phi$ is a class. $\exists r.\phi$ is the application of the *someValuesFrom* constructor to r with scope ϕ, and $\forall r.\phi$ is the *allValuesFrom* constructor.
- If r is a property and n an integer, then $\leq\ n\ r$ is a class and $\geq\ n\ r$ is a class. If r is a property and x an instance, then $x \in r$ is a class. $\leq\ n\ r$ is the application of the *maxCardinality* constructor to r and $\geq\ n\ r$ the application of *minCardinality*. If r is a property and x an instance, then $x \in r$ is the application of the *hasValue* constructor.
- If r is a property r^{-1} is a property. r^{-1} is an application of the *inverseOf* constructor.
- If r is a property, *domain*(r) and *range*(r) must be classes. *domain*(r) and *range*(r) are the *domain* and *range* operators.

We distinguish between the basic model (Figure 1), our work which models social constructions as agents and, in particular, normative systems (in dashed lines in the left hand side of Figure 2), and the dynamics of socially constructed agents (in grey in the right hand side of Figure 2). In Figure 1 and 2 we follow the convention that a box is a class, an oval is an instance, a line or an arrow with full triangle head is an association which can be interpreted as any property, an arrow with unfilled triangle head is the "is-a" or subsumption relation.

[1] The OWL version of the ontology can be found at `http://normas.di.unito.it/ontology`

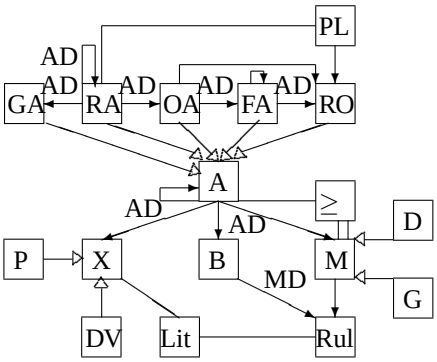

Figure 1: The basic conceptual model.

The top classes of our basic ontology are $\{A, X, B, M, \textit{Rul}, \textit{Lit}, \geq\}$. A is the set of agents, X is the set of propositions. B is the set of beliefs. M is the set of motivations. *Rul* is the set of rules. Lit is the set of literals made out of propositions (in the following we only consider propositions for simplicity). $\geq$ is the reification of a relation associating for each agent a preference relation on motivations. The main properties are $\{AD, MD, \textit{Input}, \textit{Output}\}$:

AD is the agent description relation associating to each agent its description in terms of agent profiles it attributes to other agents, mental attitudes and actions: $\textit{domain}(AD) = A, \textit{range}(AD) = A \sqcup X \sqcup B \sqcup M$.

MD is the mental attitude description relation associating to each belief or motivation its description in terms of rules: $\textit{domain}(MD) = B \sqcup M, \textit{range}(MD) = \textit{Rul}$.

Input, *Output* connect rules (*Rul*) to their antecedents and consequents: $\textit{domain}(\textit{Input}) = \textit{Rul}, \textit{range}(\textit{Input}) = X$, $\textit{domain}(\textit{Output}) = \textit{Rul}, \textit{range}(\textit{Output}) = X$.

It must be noted that real agents attribute mental attitudes also to other real agents. This attribution, even if it shares the same metaphor as attributing mental attitudes to social entities, has a different meaning, in that it represents the fact that an agent has a profile also of other real agents which interact with it in the same environment. In [4], for example, we use this attribution to model how an agent predicts by means of recursive modelling [19] the behavior of its partners using the mental attitudes attributed to them. Another reason agents have a profile of other agents is that, only by knowing other agents' goals, they can adopt some of them [18] and, thus, be respectful, trustful or cooperative agents (see agent types in [13]).

Agents can be a member of a group (GA), play a role (RO), be a member of an organization (OA). The characterizing feature of our model is that normative systems and organizations (OA), functional areas (FA), roles (RO) and groups (GA) are modelled as agents (A): they are subtypes of agents and they are distinguished from the real agents (RA) by the fact that they exist only as long as they are attributed mental attitudes by the other agents.

This collective acceptance of social entities modelled by means of the intentional stance [17] is expressed by the agent description AD property; it associates to each agent another agent in A which specifies its profile in terms of actions, beliefs and motivations by means of the AD property.

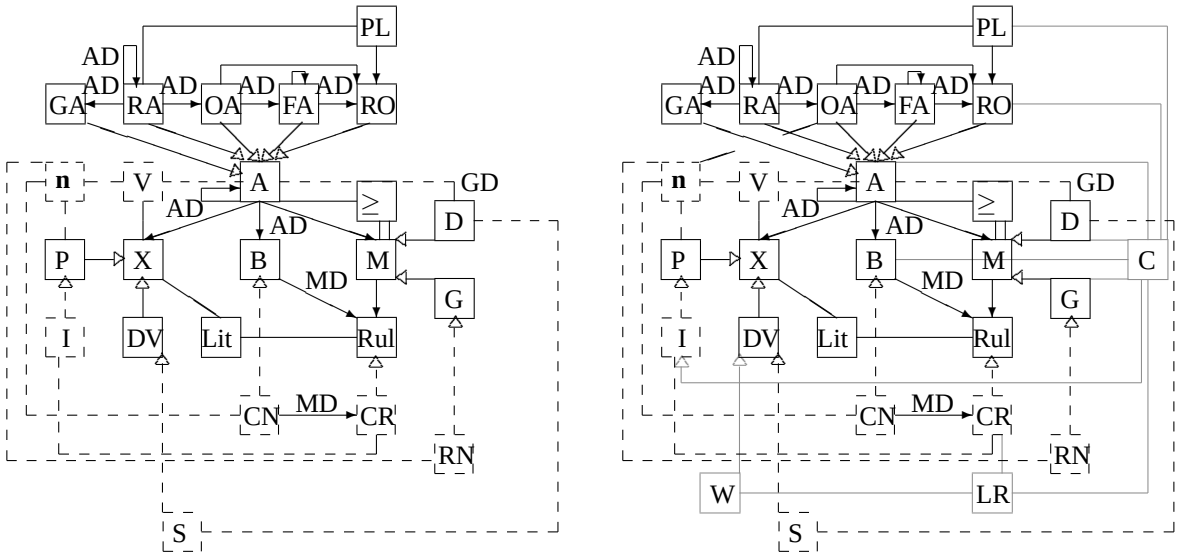

Figure 2: The social reality and its dynamics.

Note that the term "agent" does not mean the same as "person". Rather it has a wider application since it can be used to denote also intelligent programs (software agents). It is used to denote entities whose complex behavior can be described or implemented by means of beliefs, desires and goals. Hence, there is no inconsistence in using the term "agent" for social entities or collectives, since agents do not have the same properties as persons.

What distinguishes groups from organizations is that the latter are organized as burocracies [9] which have a control structure composed of obligations, prohibitions and permissions. This introduces the notion of normative systems. Normative systems contain norms which are described by violations (V) and sanctions (S).

These classes in the ontology are $\{RA, OA, FA, RO, GA, D, G, DV, P, I, PL, V, S, CN, RN, CR\}$:

RA, OA, FA, RO, GA are, respectively, real agents, organizations, functional areas and roles, and groups. They are all subsets of agents A but they are differentiated by restrictions on the AD property representing the agents which they attribute profiles to. Social entities exist only as long as they are attributed mental entities by other agents:

$$RA \sqsubseteq A \sqcap \forall AD.(RA \sqcup GA \sqcup OA \sqcup X \sqcup B \sqcup M)$$

$$OA \sqsubseteq A \sqcap \forall AD.(FA \sqcup RO \sqcup X \sqcup B \sqcup M) \sqcap \exists AD^{-1}.RA$$

$$FA \sqsubseteq A \sqcap \forall AD.(FA \sqcup RO \sqcup X \sqcup B \sqcup M) \sqcap \exists AD^{-1}.(OA \sqcup FA)$$

$$RO \sqsubseteq A \sqcap \forall AD.(X \sqcup B \sqcup M) \sqcap \exists AD^{-1}.(OA \sqcup FA)$$

$$GA \sqsubseteq A \sqcap \forall AD.(X \sqcup B \sqcup M) \sqcap \exists AD^{-1}.RA$$

In the following, with $\mathbf{n} \in OA$ we represent an instance of normative system.

D, G the desires and goals are subset of motivations; further properties must be added to distinguish one from the other, e.g., about their place in the preference relation : $D \sqsubseteq M, G \sqsubseteq M$.

DV, P represent respectively the actions of agents and the propositions which can only be indirectly influenced by agents. The two classes are disjoint since they have incompatible restrictions on the AD property $DV \equiv (X \sqcap (\geq 1\ AD^{-1}))$ and $P \equiv (X \sqcap (= 0\ AD^{-1}))$.

I the institutional facts are a subset of the parameters P, hence, they are not directly controlled by an agent. $I \sqsubseteq P$.

PL is the reification of the role playing relation: $\textit{domain}(PL) = RA, \textit{range}(PL) = RO$. It is reified to allow creation actions to change its extension.

V the violation variables are a subset of the decision variables DV (of the normative system **n**); a violation variable is related to a literal which represents which situation is considered as a violation and an agent who is the responsible for the violation (these properties are not included for space reasons, see Figure 1): $V \sqsubseteq (DV \sqcap \mathbf{n} \in AD^{-1})$.

S the sanctions are a subset of the decision variables DV (of the normative system $\mathbf{n}$); a decision variable is a sanction only if it is not desired by the sanctioned agent (but this constraint is too complex to be represented in OWL): $S \sqsubseteq (DV \sqcap \mathbf{n} \in AD^{-1})$

CN constitutive norms are beliefs of the normative system $\mathbf{n} \in OA$ which are described by rules having as *Output* an institutional fact (or more precisely a literal built out of an institutional fact): $CN \sqsubseteq (B \sqcap (\mathbf{n} \in AD^{-1}) \sqcap \forall MD.CR)$ and $CR \sqsubseteq (\textit{Rul} \sqcap \forall \textit{Output}.I)$.

RN the regulative norms are goals of the normative system $\mathbf{n}$; a more complete formalization requires adding sanctions and violations, but this requires constraints on the fillers of properties: $RN \sqsubseteq (G \sqcap \mathbf{n} \in AD^{-1})$.

One fundamental feature of social reality is its dynamics. In this paper we consider, in particular, how normative systems specify the possibilities of changing themselves and that other agents change their norms.

Regulative norms like obligations and permissions are defined as goals of the normative systems and constitutive norms are defined as their beliefs; moreover, normative systems are defined as autonomous agents; hence, the only possibility to change norms is that the normative system autonomously changes its own goals and beliefs. For this reason we add to the conceptual model a set of institutional facts called creation actions $C \in I$ which change beliefs, desires and goals of an agent, and, in particular, of the normative system. They also change the agents which compose an organizations (in this way it is possible to change roles, as they are defined as agents) and the role playing relation between the real agents and the role they play. These institutional facts can be made true only if the normative system believes they are true since some other fact "counts as" them by means of some constitutive rules: i.e., by means of a constitutive norm. Hence, only the normative system can decide who has the power to change its own beliefs, desires and goals, and, as a consequence, its norms.

These classes are $\{C, W, LR\}$:

C the creation actions are institutional facts which change beliefs and motivations, roles and role playing relations; we do not address in this paper how to model such effects: $C \sqsubseteq I$.

W the powers of agents are actions (decision variables) which can change the mental attitudes of agents by means of constitutive norms relating them to creation actions: $W \sqsubseteq (DV \sqcap LR)$ and $LR \sqsubseteq (CR \sqcap \exists MD^{-1}CN \sqcap \forall \textit{Output}.C)$.

4 Related work

Breuker *et al.* [28] describe a 'functional modelling' of legal systems. The model is based on the idea that:

> "To a large extent, law is also an artifact, so a functional perspective is strongly implied already by the objects (concepts) it uses [...] that is, the Law is analyzed and interpreted through a functional point of view. [...] It is assumed that the legal system as a whole (and therefore each of its components) exists to accomplish a certain function, in order to obtain certain social goals. The legal system is, thus, viewed as an entity with a certain internal structure, behaving in an environment, and that was designed to work in a certain way in order to be able to accomplish specified social goals. The main function of the legal system is to change or influence society in specific directions, determined by certain social goals."

We share this idea that the legal system has a function (achieving a social order) and that it is conceived to achieve this function. What is different is that our work bases the social ontology on the concept of agent, which, by using the intentional stance, can be seen as a further methodology besides functional models to describe artifacts like social entities are.

Our approach has important consequences for the notion of social role. Steimann [29] proposes a detailed analysis of roles which considers the possibility that roles can be played by other roles, e.g., an employee can be a project director. In our model it is possible to extend the PL property to cope with this case, but it is necessary first to consider which is the decision process of an agent playing a role which plays in turn another role.

Masolo *et al.* [30] provide an ontological analysis of social roles. They reify roles and attribute them the properties of anti-rigidity (playing a role is not a necessity), relational nature (their instances depend on instances of other concepts) and temporal extension. In this paper, roles are considered as entities, too. In contrast we attribute to roles the status of agent, since roles are usually defined as the expected behavior of their actors and we describe complex behavior by means of agents with beliefs, desires and goals. In [26], we address the problem of how roles described as agents are used by the actors when they play a role.

Moreover, roles in our model can be considered anti-rigid and temporally extended, in that, even if we do not currently model time explicitly, the relation connecting agents and roles and the set of roles themselves can be changed in organizations by means of creation actions. The relational nature of roles is explained by the fact that they exist only if they are attributed mental attitudes by an organization. Restrictions on role compatibility and requirements (e.g., to be president of a state an agent must be a citizen of that state) are expressed by the constraints on the constitutive rules creating relations between agents and the roles they play: a precondition for declaring an agent president is that it is a citizen of the state. On the other hand, we do not consider yet the problem of descriptions defining social concepts.

Works on multiagent systems offer as aggregation methods the notion of group or of organization. According to Zambonelli *et al.* [31] "an organization is more than simply a collection of roles (as most methodologies assume) [...] further organization-oriented abstractions need to be devised and placed in the context of a methodology [...] As soon as the complexity increases, modularity and encapsulation principles suggest dividing the system into different suborganizations". According to Jennings [32], however, most current approaches "posses insufficient mechanisms for dealing with organisational structure". Moreover, what is the semantic principle which allows decomposing organizations into suborganizations must be still made precise.

5 Conclusions

In this paper we present in a unified way the ontology underlying our approach to normative systems, organizations, groups and roles we described in several papers (e.g., [4, 5, 8, 11, 21, 26]). The general methodology is to explain social entities by attributing them mental attitudes, like beliefs, desires and goals: i.e., by considering social entities as agents. This metaphor leads to a simplification of the ontology and allows modelling social reality in terms of notions commonly used by agent theories; in this way, it can be used to structure multiagent systems.

In this paper we focus on the relations between the concepts populating our ontology rather than on providing a justification of them, since this is discussed in the different papers. Moreover, here we do not address the discussion of behavior of agents which respect this conceptual model, for example, how they take decisions and recursively model the behavior of other agents, how they act when they play roles in organizations, *et cetera*: again this is discussed in related papers; moreover, this issue requires a complex ontology to describe processes.

Future work is discussing the properties of roles, comparing them with the requirements posed by [29] and extending the ontology to include the notion of coalitions seen at different levels of details as discussed in [33]. Finally, it should be studied how the basic ontology can be integrated with some well founded top ontology like DOLCE [34].

References

[1] V. Dignum, J-J. Meyer, and H. Weigand. Towards an organizational-oriented model for agent societies using contracts. In *Procs. of AAMAS'02*, pages 694–695, Bologna, 2002. ACM Press.

[2] P. R. Cohen and H. J. Levesque. Confirmation and joint action. In *Procs. of IJCAI'91*, pages 951–957, Sydney, 1991.

[3] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[4] G. Boella and L. van der Torre. Groups as agents with mental attitudes. In *Procs. of AAMAS'04*, pages 964–971, New York, 2004.

[5] G. Boella and L. van der Torre. Attributing mental attitudes to normative systems. In *Procs. of AAMAS'03*, pages 942–943, Melbourne, 2003. ACM Press.

[6] G. Boella and L. van der Torre. Regulative and constitutive norms in normative multiagent systems. In *Procs. of KR'04*, pages 255–265, Whistler (CA), 2004.

[7] H.L.A. Hart. *The Concept of Law*. Clarendon Press, Oxford, 1961.

[8] G. Boella and L. van der Torre. Contracts as legal institutions in organizations of autonomous agents. In *Procs. of AAMAS'04*, pages 948–955, New York, 2004.

[9] W.G. Ouchi. A conceptual framework for the design of organizational control mechanisms. *Management Science*, 25(9):833–848, 1979.

[10] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook. Theory, Implementation and Applications*. Cambridge University Press, 2003.

[11] G. Boella and L. van der Torre. Local policies for the control of virtual communities. In *Procs. of IEEE/WIC WI'03*, pages 161–167, Halifax (CA), 2003. IEEE Press.

[12] D. Makinson and L. van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.

[13] J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly*, 2(3-4):428–447, 2002.

[14] A. S. Rao and M.P. Georgeff. Deliberation and its role in the formation of intentions. In *Uncertainty in Artificial Intelligence: Proceedings of the Seventh Conference (UAI-1991)*, pages 300–307, San Mateo, CA, 1991. Morgan Kaufmann Publishers.

[15] J.R. Searle. *The Construction of Social Reality*. The Free Press, New York, 1995.

[16] R. Tuomela. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press, 1995.

[17] D. Dennett. *The intentional stance*. Bradford Books/MIT Press, Cambridge (MA), 1987.

[18] C. Castelfranchi. Modeling social action for AI agents. *Artificial Intelligence*, 103:157–182, 1998.

[19] P. J. Gmytrasiewicz and E. H. Durfee. Formalization of recursive modeling. In *Procs. of ICMAS'95*, pages 125–132, 1995.

[20] C. Castelfranchi. Engineering social order. In *Procs. of ESAW'00*, pages 1–18. Springer Verlag, 2000.

[21] G. Boella and L. van der Torre. Norm governed multiagent systems: The delegation of control to autonomous agents. In *Procs. of IEEE/WIC IAT'03*, pages 329–335, Halifax (CA), 2003. IEEE Press.

[22] F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7:69–79, 1999.

[23] G. Boella and L. van der Torre. Normative multiagent systems and trust. In *Procs. of Trust in Agent Societies Workshop at AAMAS'04*, New York, 2004.

[24] G. Boella and L. van der Torre. Permissions and obligations in hierarchical normative systems. In *Procs. of ICAIL'03*, pages 109–118, Edinburgh, 2003. ACM Press.

[25] J.R. Searle. *Speech Acts: an Essay in the Philosophy of Language*. Cambridge University Press, Cambridge (UK), 1969.

[26] G. Boella and L. van der Torre. Attributing mental attitudes to roles: The agent metaphor applied to organizational design. In *Procs. of ICEC'04*, Delft (NL), 2004.

[27] G. Boella and L. van der Torre. Game specification in normative multiagent system: the trias politica. In *Procs. of IAT'04*, Beijing, 2004.

[28] J. Breuker, A. Valente, and R. Winkels. Legal ontologes: A functional view. In *Procs. of 1st LegOnt Workshop on Legal Ontologies*, pages 23–36, 1997.

[29] F. Steimann. On the representation of roles in object-oriented and conceptual modelling. *Data and Knowledge Engineering*, 35:83–848, 2000.

[30] C. Masolo, L. Vieu, E. Bottazzi, C. Catenacci, R. Ferrario, A. Gangemi, and N. Guarino. Social roles and their descriptions. In *Procs. of KR'04*, Whistler (CA), 2004.

[31] F. Zambonelli, N.R. Jennings, and M. Wooldridge. Developing multiagent systems: The Gaia methodology. *IEEE Transactions of Software Engineering and Methodology*, 12(3):317–370, 2003.

[32] N. R. Jennings. On agent-based software engineering. *Artificial Intelligence*, 117(2):277–296, 2000.

[33] G. Boella, L. Sauro, and L. van der Torre. Social viewpoints on multiagent systems. In *Procs. of AAMAS'04*, pages 1358–1359, New York, 2004.

[34] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, and L. Schneider. Sweetening ontologies with dolce. In *Proc. EKAW 2002*, Siguenza (SP), 2002.