

Spatially Distributed Normative Objects

Fabio Y. Okuyama¹, Rafael H. Bordini², and Antônio Carlos da Rocha Costa³

¹ Universidade Federal do Rio Grande do Sul, Brazil. okuyama@inf.ufrgs.br

² University of Durham, United Kingdom. R.Bordini@durham.ac.uk

³ Universidade Católica de Pelotas, Brazil. rocha@atlas.ucpel.tche.br

Abstract. Organisational structures for multi-agent systems are usually defined independently of any spatial or temporal structure. Therefore, when the multi-agent system is situated in a spatial environment, there is usually a conceptual gap between the definition of the system’s organisational structures and the definition of the environment. In this paper, we focus on a mechanism for the spatial distribution of an organization’s normative information. Spatially distributing the normative information over the environment is a natural way to simplify the definition of organisational structures and the development of large-scale multi-agent systems. By distributing the normative information in different spatial locations, we allow agents to directly access the relevant information needed in each environmental context. We extend our previous work on a language for modelling multi-agent environments in order to allow for the definition of spatially distributed norms in the form of *normative objects*.

1 Introduction

The environment is an important part of a Multi-Agent System (MAS), specially for systems of situated agents. Situated multi-agent systems are usually designed as a set of agents, together with the environment where they operate, their social structures, and the possible interactions among these components. In previous works, we introduced a language that allows MAS designers to describe, at a high conceptual level, environments for situated multi-agent systems [11, 1]. The language is called ELMS, and was created to be part of a platform for the development of (social) simulations based on multi-agent systems.

In this paper, we present extensions to the ELMS language which allow the distribution of normative information over an environment, construing what we call *situated norms*. In particular, we introduce here the notion of spatially distributed *normative objects*, which facilitates the modelling of various real-world situations, particularly for simulation, but more generally the coordination of large-scale multi-agent systems too, through situated norms.

To understand the notions of normative object and situated norm, consider the posters one typically sees in public places (such as libraries or bars) saying “Please be quiet” or “No smoking in this area”. Human societies often resort to this mechanism for decentralising the burden of regulating social behaviour; people then adopt such situated norms whenever they have visual access to such posters. This should be equally efficient for computational systems because it avoids the need for providing a complete,

In G. Boella et al. (eds), Proceedings of the International Workshop on Coordination, Organization, Institutions and Norms in Agent Systems (COIN), held with ECAI, Riva del Garda, Italy, 28 August 2006.

exhaustive representation of all social norms in a single public structure, known to all agents, as it is usually the case in current approaches to agent organisations.

Another extension we have introduced to our environment description language is the notion of normative places, which are zones where the normative objects and situated norms are relevant. As an example, consider a research group where there are agents with the role *principal researcher* whose main objective is to supervise the research of agents playing the *research student* role; such research can be conducted both at the laboratory or at the library. The interactions at the laboratory are to be outlined in the spatial scene of the laboratory space. The information about how to behave in a library is defined in the library spatial scene, where all researchers will also assume the role of *library users*. Normative information relevant for each such site (and each place within each site) can be made accessible to the agents with the help of normative objects.

In summary, the extensions we introduce here support situated norms and leaves the necessary room for the inclusion of group structures that are spatially situated within a (simulated) physical environment. This is done using two means: first, *normative objects*, which are objects that can contain normative information; and second, a normative principle for *situated norms*, conceived as a special form of conditional rule, where an explicit condition on an agent's perception of a normative object appears: 'When playing the relevant role and being physically situated within the confines referred by a situated norm \mathcal{N} expressed in a normative object previously perceived, the agent is required to reason about following norm \mathcal{N} ; otherwise, it is excused from reason about it'. Also, normative objects may be directed towards a specific role in a given organisation. We can thus model things such as a sign saying that students are not allowed beyond the library desk (while members of staff are).

In the next section, we briefly present our platform and the various component languages we use to model multi-agent systems. In Section 3, we briefly review how an environment should be modelled using our approach. In Sections 4 and 5, we present and discuss the normative extensions that we introduce in this paper. We then illustrate our approach with an example in Section 6; the example is based on the scenario presented in [4]. We discuss related work in Section 7, then conclude the paper.

2 The MAS-SOC Platform

One of the main goals of the MAS-SOC simulation platform (**M**ulti-**A**gent **S**imulations for the **S**OCial Sciences) is to provide a framework for the creation of agent-based simulations which do not require too much experience in programming from users, yet allowing the use of state-of-the-art agent technologies. In particular, it should allow for the design and implementation of simulations with *cognitive* agents.

In our approach, an agent's individual reasoning is specified in an extended version of AgentSpeak [13], as interpreted by *Jason*, an *open source* agent platform⁴ based on Java [2]. The extensions allow, among other things, the use of speech-act based agent communication, and there is ongoing work to allow the use of ontologies and of organisational structures as part of a *Jason* multi-agent system.

⁴ Available at <http://jason.sf.net>.

The environments where agents are situated are specified in ELMS, a language we have designed for the description of multi-agent environments [11]. For more details on MAS-SOC, refer to [1]. We here concentrate on the ELMS extensions to describe basic organisational structures and social norms, and to relate an organisational structure and the relevant normative aspects to the spatial structures defined within the physical environment.

3 Modelling Physical Environments with ELMS

As presented in [11], we developed ELMS (Environment Description Language for Multi-Agent Simulation) as a means to describe environments and to execute simulated environments. Agents in a multi-agent system interact with the environment where they are situated and interact with each other (possibly through the shared environment). Therefore, the environment has an important role in a multi-agent system, whether the environment is the Internet, the real world, or some simulated environment.

We understand as environment modelling, the modelling of external aspects that an agent needs as input to its reasoning and for deciding on its course of action. Further, it is necessary to model explicitly the physical actions and perceptions that the agents are capable of in a given environment. Below we briefly review how a physical environment is described using this language.

To define an environment using ELMS, the following classes of constructs can be used:

- **Agent Body:** the agent’s characteristics that are perceptible to other agents. Agent “bodies” are defined by a set of properties that characterise it and are perceptible to other agents. Such properties are represented as *string*, *integer*, *float*, and *boolean* values. Each “body” is associated with a set of actions that the agent is allowed to perform and of environment properties that the agent can perceive.
- **Agent Sensorial Capabilities (Perception):** the environment properties that will be perceptible to each agent at a given time, and under given specific circumstances.
- **Agent Effective Capacities (Actions):** the actions that an agent is able to perform in order to change the current state of the environment. These actions are defined as assignments of values to the attributes of entities in the environment⁵. The production (i.e., instantiation) of previously defined resources (i.e., objects), and the consumption (i.e., deletion) of existing instances can also be part of an action description.
- **Physical Environment Objects (Resources):** the objects/resources that are present in the environment. Although objects and resources can have conceptual differences, they are represented by the same structure in ELMS. Agents interact with objects through their actions in the environment. Object structures are defined by a set of properties that are relevant to the modelling and may be perceived by an agent. In the same way as the properties of the “bodies” of the agents, the properties of objects are also represented by *string*, *integer*, *float*, and *boolean* values. Each object can also be associated with a set of reactions that may happen as consequence of an agent’s actions.

⁵ Note that agent bodies are also properties of the environment.

- **Object Reactions:** the objects can “react”, under specific circumstances, in order to respond to actions performed by the agents in the environment. Such reactions are given as the assignment of values to properties, the creation of previously defined object instances, and the deletion of existing object instances.
- **Space Structure (Grid):** the space is (optionally) divided into cells forming a grid that represents the spatial structure of the environment. When a grid is used, it can be defined in 2 or 3 dimensions. As for resources, each cell can have reactions associated to it. Although the specified set of reactions apply to all of the cells, this does not mean that all cells will behave equally, since they may be in different contexts (i.e., each cell has independent attributes, thus having different contents and, clearly, different positions, which can all affect the particular reactions).

3.1 Notes on Environment Descriptions

- **Perceptions:** agents do not normally have complete access to the environment. Perception of the environment will not normally give complete and accurate information about the whole environment and the other agents in it. However, since such restriction is not imposed by the ELMS model itself, designers can choose to create fully accessible environments if this is appropriate for a particular application.
- **Actions:** actions defined here are assumed to be atomic, as the action chaining or planning is meant to be part of the “mind” of the agent
- **Reactions:** all object reactions triggered by some change in the environment are executed in a single simulation cycle. This is different from agent actions, as each agent can execute only one action per cycle.

Additionally to the constructs mentioned above, the following operational constructs are used in our approach to model the (simulated) physical environment.

- **Constructors:** Each agent and resource may need to be initialised at the moment of its instantiation. This is defined by a list of initial value assignment to its attributes.
- **Observables:** A list of environment properties whose values are to be displayed/logged; these are the specific properties of a simulation that the user wants to observe/analyse.

The simulation of the environment itself is done by a process that controls the access and changes made to the data structure that represents the environment; the process is called the *environment controller*. The data structure that represents the environment is generated by the ELMS interpreter from a specification in ELMS given as input. In each simulation cycle, the environment controller sends to all agents currently taking part in the simulation the percepts to which they have access (as specified in ELMS). Recall that ELMS environments are designed for cognitive agents, so perception is transmitted in messages as a list of ground logical facts. After sending perception, the process waits for the actions that the agents have chosen to perform in that simulation cycle and then execute the actions, changing the environment data structures accordingly.

4 Normative Objects and Situated Norms

Typically, environments will have some objects aimed at informing agents about norms, give some advice, or warn about potential dangers. For example, a poster fixed on a wall of a library asking for “silence” is an object of the environment, but also informs about a norm that should be respected within that space. Another example are traffic signs, which give advice about directions or regulate priorities in crossings. The existence of such signs, that we call *normative objects*, implies the existence of a regulating code in such context, that we call *situated norms*.

In the examples above, the norms are only meant to be followed within certain boundaries of space or time and lose their effect completely if those space and time restrictions are not met, which is the initial motivation for situated norms. Another important advantage of modelling some norms as situated norms is the fact that the spatial context where the norm is to be followed is immediately determined. Thus, the norm can be “pre-compiled” to its situated form, making it easier for the agents to operationalise the norm, and also facilitating the verification of norm compliance.

For example a norm that says “Be kind to the elderly”, may be quite hard to operationalise and verify, in general. However, in a fixed spatial context such as a bus or train, with the norm contextualised as “Give up your seat for the elderly”, or in a street crossing, with the norm contextualised as “Help elderly people to cross the street”, the norm would be much more easily interpreted by the agents, and verified by any norm compliance checking mechanism.

It is important to remark that the norm-abiding behaviour is not related to the existence of a normative object. Beyond the existence of such object, it is necessary for the agent to perceive the normative object, and autonomous agents will also reason about whether to follow or not the norm stated by the normative object.

4.1 Normative Objects and Situated Norms in ELMS

In the extended version of ELMS, normative objects are “readable” by agents under specific individual conditions: an agent is able to read a specific rule if it has the specific ability to perceive the type of object in which the rule is written at its given location. In the most typical case, the condition is simply being physically close to the object.

Normative objects can be defined before the simulation starts, or can be created dynamically during the simulation. Each normative object can be placed in a normative place (see below), in the spatial grid of the environment. The conditions under which the normative objects can be perceived are defined by the simulation designer using the usual ELMS constructs for defining perception capabilities and their conditions.

The normative information in a normative object is “read” by an agent through its perception ability. Besides the norm itself, it may contain meta-information, e.g., which agent or institution created the norm. In ELMS, normative objects should have at least the following properties:

- **Type:** the type of the normative information contained in the object; it determines the level of importance (e.g., a warning, an obligation, a direction);

- **Issued by:** where the power underlying the norm comes from (e.g., an agent, a group, an institution).
- **Norm:** a string that represents the normative information; this should be in the format of AgentSpeak predicates in the case of MAS-SOC environments, or whatever format the targeted agents will be able to understand.
- **Placement:** the set of normative spaces where the normative information applies. If omitted, the object is assumed to be accessible from anywhere, but normally under conditions determined by the designer; see the next item.
- **Condition:** conditions under which the normative information can be perceived. The conditions can be associated with groups, roles, abilities, and current physical placement and orientation of agents and objects.
- **Id:** identification string for eventual deletion/editing of the normative object.

We now briefly describe how the agents will receive normative information from normative objects. Whenever the agent position is such that access to the normative object is accessible, and the **Condition** is satisfied, the agent will receive perception of the form:

```
rule([PLACE],[GROUP],[ISSUED BY],[NORM])
Ex: rule(home, family, parents, obligation(child,play(TOY),tidy(TOY)))
```

The example above can be read as: “This is rule in group *family*, issued by the *parents*, with application at the normative place *home* (see below), that says: if the action *play(TOY)* is done by an agent of role *child*, then it is an obligation of that agent to do *tidy(TOY)* as well”.

A rule like that would not normally be posted on a sign in a family home, but it illustrates the more general idea of situated norms as norms that apply within given environmental locations.

4.2 Normative Places in ELMS

Considering the ideas discussed above about normative objects and situated norms, ELMS descriptions of an environment (based on the concepts of agent bodies, objects, and an optional grid) need to be extended with the notion of *normative places*, i.e., a set of cells where an organisational activity occurs under the conditions of a set of situated norms.

A *normative place* can be defined in ELMS simply by its name and the set of cells that are part of it. A normative place may have intersections with other normative places, or even be contained within another normative place. For example, a normative place “school” may have a large set of cells where some of those cells refer to a normative place “classroom” and others to its “library”. A normative place allows for the definition of the spatial location where certain situated norms are valid and relevant, as it will be exemplified in the next section.

In order to facilitate the definition of repetitive normative place structures, classes of normative places can be first defined and then instantiated in specific positions of the grid. The place “home” in the previous section is an example of a normative place. Other examples of such definitions and instantiations are as follows:

```

<NORMATIVE-PLACE-TYPE NAME="library"/>
<NORMATIVE-PLACE-TYPE NAME="classroom"/>
<PLACE NAME="lib1" NORMATIVE-PLACE-TYPE="library">
  <CELL X="0" Y="0"/>
  <CELL X="0" Y="1"/>
</PLACE>
<PLACE NAME="cr1" NORMATIVE-PLACE-TYPE="classroom">
  <CELL X="2" Y="0"/>
</PLACE>

```

5 MAS-SOC Modelling of Organisations Governed by Situated Norms

As the MAS-SOC platform does not enforce a particular agent-oriented software engineering methodology, designers can use the one they prefer. It is possible to model a multi-agent system that will have an ELMS environment using any approach: starting from the system organisation (top-down), or starting from the agent interactions (bottom-up).

In both approaches, the modelling of the organisational structures and the agents' reasoning need fine tuning to achieve the desired results. To have a stable point on which to base the tuning-up of the agents' reasoning or the organisational model, we have suggested the use of an explicitly defined environment description written in the ELMS language and the concepts presented in the Section 3. The environment is an important part of a multi-agent system, and although it can be very dynamic, in regards to design it is usually the most "stable" part of the system.

Based on these observations, we suggest that the multi-agent system modelling starts with the environment definition, followed by the definition of the normative places. The environment modelling is achieved by:

1. Definition of which kinds of action each type of agent is able to perform in the environment. Actions typically produce effects over objects of the environment or other agents.
2. Based on the changes that the agents' effective capabilities are able to make in the environment and the objectives of the simulation, the size and granularity of the grid can be determined. For example, how many cells an agent can move within one action or simulation cycle, and in how many simulation cycles the agent would be able to traverse the simulated space.
3. Based on the granularity and size of the spatial environment, the sensorial capabilities of the agents can be modelled, defining for example in which range an agent can detect other agents or objects.
4. Based on an agent's sensorial capabilities and on its typical activities, it should be possible to define which attributes of that agent is important to declare as accessible to other agents. For example, if agents identify each other's role by the colour of their uniform, the "agent body" should have an attribute that represent the colour of the agent's uniform.

5. The types of objects or resources present in the environment should also be modelled based on which attributes will be perceptible by the agents and which actions can affect them.
6. Finally, instances of the agent and object classes should be placed in the environment, determining its initial state.

The definition of the environment should be followed by the definition of normative places and then by the definition of the spatially distributed normative objects, as follows:

1. Together with the object types placed in the environment, the types of normative places within the environment can also be defined.
2. By instantiating normative places into sets of cells, *normative places* are created.
3. Then, based the set of activities that can possibly be performed in each type of normative place, the norms that are relevant to that type of place can be defined.
4. Finally, the types of *normative objects* can be defined and instantiated in the normative places, defining the locations where situated norms can be perceived.

Using the environment as a basis, the agents' reasoning capabilities can then be defined so as to help agents achieve their objectives as well as the objective of the groups in which they participate. Also, the detailed definitions of possible organisational structures can be fine-tuned, in order to have the system achieving its overall objectives. In MAS-SOC, we use AgentSpeak to define the practical reasoning for each agent; in particular, we use the extended version of AgentSpeak as interpreted by *Jason*; for details, see [3].

6 Example

Below we give an example showing how normative objects are defined using our approach. It is based on the scenario presented in [4], a scenario in which the agents are placed on an environment where they may eat the food they find, challenge other agents for their food, or move in search of food.

In this scenario, an agent owns any food item that is near to itself (at a distance of up to 2 cells). The agents can "see" food and other agents in a radius of 1 cell, but can sense food in a radius of 2 cells. The physical space is represented by a grid of 10×10 cells.

The norms used in that scenario essentially concern the respect for the ownership of a food item, which means they prescribe non-aggressive behaviour. In the original scenario, the norms were valid throughout the grid, but in this example norms are valid only within normative places, as indicated by normative objects.

A shortened version of the physical environment description in ELMS is given below.

```
<!DOCTYPE ENVIRONMENT SYSTEM "elms.dtd">
<ENVIRONMENT NAME = "NORMATIVE">
  <DEFGRID SIZEX="10" SIZEY = "10"/>

  <RESOURCE NAME="food">
```

```

    <STRING ownner = "none">
  </RESOURCE>

  <AGENT_BODY NAME="agent">
    <INTEGER NAME = "id"> "SELF" </INTEGER>
    <PERCEPTIONS>
      <ITEM NAME = "vision"/>
      <ITEM NAME = "sense_food">
    </PERCEPTIONS>
    <ACTIONS>
      <ITEM NAME = "walk"/>
      <ITEM NAME = "attack"/>
      <ITEM NAME = "eat"/>
    </ACTIONS>
  </AGENT_BODY>

  <PERCEPTION NAME="vision">
    <CELL_ATT ATTRIBUTE="CONTENTS" ABSOLUTE="TRUE">
      <X> +0</X>
      <Y> +0</Y>
    </CELL_ATT>
    <CELL_ATT ATTRIBUTE="CONTENTS" ABSOLUTE="TRUE">
      <X> +1</X>
      <Y> +0</Y>
    </CELL_ATT>
    <!-- shortened-->
  </PERCEPTION>

  <PERCEPTION NAME="sense_food">
    <!-- shortened-->
  </PERCEPTION>

  <ACTION NAME="eat">
    <PARAMETER NAME = "FOOD_ID" TYPE="INTEGER" />
    <!-- shortened-->
  </ACTION>

  <ACTION NAME="walk">
    <!-- shortened-->
  </ACTION>

  <ACTION NAME="attack">
    <!-- shortened-->
  </ACTION>

  <INITIALIZATION>
    <!-- instantiation and placement of
      food and agents -->
  </INITIALIZATION>
</ENVIRONMENT>

```

In the code excerpt above, the grid size is defined, then `food` is defined as an environment resource, then a generic type of agent body is defined. The agent body is defined as being capable of two types of perception — vision and food sensing – and being able to perform three types of actions: walk, attack, and eat. The `vision` perception allows the agent to perceive the contents of the current cell and the 4 neighbouring cells, while `sense_food` allows it to perceive food within a 2-cell radius.

For this example, the grid is partitioned in four normative places of equal sizes, and the normative objects are defined and placed only in the upper-left and upper-right quadrants, as shown in the code excerpt below:

```
<NORMATIVE-PLACE-TYPE NAME="food-protected"/>
```

```

<PLACE NAME="upper-left" NORMATIVE-PLACE-TYPE="food-protected">
  <CELL X="0" Y="0"/><CELL X="1" Y="0"/>
  <!-- shortened-->
  <CELL X="3" Y="4"/><CELL X="4" Y="4"/>
</PLACE>

<PLACE NAME="upper-right" NORMATIVE-PLACE-TYPE="food-protected">
  <CELL X="5" Y="0"/><CELL X="6" Y="0"/>
  <!-- shortened-->
  <CELL X="8" Y="4"/><CELL X="9" Y="4"/>
</PLACE>

<PLACE NAME="lower-left" NORMATIVE-PLACE-TYPE="null">
  <CELL X="0" Y="5"/><CELL X="1" Y="5"/>
  <!-- shortened-->
  <CELL X="3" Y="9"/><CELL X="4" Y="9"/>
</PLACE>

<PLACE NAME="lower-right" NORMATIVE-PLACE-TYPE="null">
  <CELL X="5" Y="4"/><CELL X="6" Y="4"/>
  <!-- shortened-->
  <CELL X="8" Y="9"/><CELL X="9" Y="9"/>
</PLACE>

<NORMATIVE_OBJECT ID="norm1" TYPE="prohibition" PLACE = "upper-left">
  <NORM>prohibited(true, attack(SELF, AGENT)) </NORM>
</NORM_OBJ>

<NORMATIVE_OBJECT ID="norm2" TYPE="prohibition" PLACE = "upper-right">
  <NORM>prohibited(not(owned(FOOD, SELF)), eat(SELF, FOOD)) </NORM>
</NORM_OBJ>

```

The norms in the above example are very simple, and are given simply to illustrate how they can be modelled in our approach. For instance, `norm1` says that an agent ought not to attack (steal food from) another agent, while `norm2` says that the agent ought not to eat a food item that is not owned by itself.

Clearly, the agents' behaviour will be different in the four quadrants of the environment:

- in the upper-left quadrant, an agent is barred from eating food that belongs to another agent (since the situated norm states that an agent is prohibited from stealing food);
- in the upper-right quadrant, agents are supposedly prohibited of doing that, but not effectively, since the situated norm only prohibits the eating of food that is not owned by the agent itself rather than the stealing of food, so an agent can eat food that previously belonged to another agent if it first manages to steal that food;
- the lower quadrants (both left and right) are lawless areas, where agents are completely free to attack each other and to eat anyone else's food.

Notice that `prohibited` is used as a conditional deontic operator, with two arguments: the first argument is a condition to be tested, the second argument is the action that is prohibited.

7 Related Work

The notion of artifacts [16] and coordination artifacts [12] resembles our notion of *normative objects*. As defined in [12], coordination artifacts are abstractions meant to

improve the automation of coordination activities, being the building blocks to create effective shared collaborative working environments. They are defined as runtime abstractions that encapsulate and provide a coordination service to the agents. Artifacts [16] were presented as a generalisation of coordination artifacts. Artifacts are an abstraction to represent tools, services, objects and entities in a multi-agent environment.

As building blocks for environment modelling, artifacts encapsulate the features of the environment as services to be used by the agents. The main objective of a coordination artifact is to be used as an abstraction of an environmental coordination service provided to the agents. However, coordination artifacts express normative rules only implicitly, through their practical effects on the actions of the agents, and so their normative impact does not require any normative reasoning from the part of the agents. In our work, rather than having a general notion of objects that by their (physical) properties facilitate coordination, *normative objects* are objects used specifically to store *symbolic* information that can be interpreted by agents, so that they can become aware of norms that should be followed within a well-defined location.

Our choice has the advantage of keeping open the possibility of agent autonomy, as suggested in [5]. Agents are, in principle, able to decide whether to follow the norms or not, when trying to be effective in the pursuit of their goals. This is something that is not possible if an agent's action can only happen if in accordance to norms enforced by coordination mechanisms.

Another important difference is that *normative objects* are spatially distributed over a physical environment, with a spatial scope where they apply, and closely tied to the part of the organisation that is physically located in that space. While the objective of the coordination artifacts is to remove the burden of coordination from the agents, our work tries to simplify the way designers can guide the behaviour of each individual agent as they move around an environment where organisations are spatially located; this allows agents to adapt the way they behave in different social contexts.

In [8], the authors present the AGRE model, an extension to the previous AGR model. These latest extensions allow the definition of structures that represents the physical space. The approach defines organisational structures (i.e., groups) and the physical structures (i.e., areas) as “specialisations” of a generic space. The social structures are not contextualised in the space as they are in our work, leaving the social and physical structures quite unrelated.

In ELMS, however, it is not possible to explicitly define social structures, even though it would be possible to implicitly define them through the norms. This is because the aim of ELMS is, as mentioned earlier, to allow for environmental infrastructures compatible with existing approaches to organisational modelling, not for the modelling of organisations as such; the combination of ELMS with existing approaches to modelling organisations is planned as future work.

Another important series of related work is that on Electronic Institutions [9]. The internal working of an electronic institutions is given (in a simplified view) as a state-machine where each state is called a “scene”. Each scene specifies the set of roles that agents may perform in it, and a “conversation protocol” that the agents should follow when interacting in the scene. To traverse the series of scenes that constitute

the operation of the electronic institution, agents must do a sequence of actions in each scene, and also to commit to certain actions in certain scenes, as the result of their having performed certain other actions in certain other scenes. Our notion of normative space was inspired by such notion of scene, through giving it a physical, spatial content.

Similar to the electronic institutions approach, there is work on *computational institutions* [14], which are defined as virtual organisations ruled by constitutive norms and regulative norms. In computational institutions, organisational modelling uses the abstraction of coordination artifacts as building blocks, in a way that is very similar to our use of normative objects in spatially distributed organizations, but still keeping implicit in coordination artifacts the normative content imposed on the agents.

8 Conclusions

In this paper we have extended the ELMS language for describing environments with the means to define normative structures that make part of an environment representation. There are currently many approaches to modelling and implementing multi-agent systems: some are top-down approaches with focus on the organisations, while bottom-up approaches focus on the agents. We believe that including environment modelling at the initial stages of both approaches would help the modelling and implementation of multi-agent systems. To help such modelling, we have proposed an approach with an explicit environment description which now also includes the notions of *situated norms*, *normative places*, and (spatially distributed) *normative objects*.

It is important to note that our work is not an approach for modelling the organisational dimension of a multi-agent system. With the definition of *normative places*, where group structures would be inserted, we intend to fill a conceptual gap between the usual ways in which organisations and physical environments are modelled. In future work, with the integration of current means for defining organisational structures with ELMS, and thus with the possibility of associating them to normative places, we hope to contribute to a more integrated approach to designing and implementing the various aspects of multi-agent systems: concentrating on one particular organisation section at a time, specially if it is an organisation section attached to a spatial location, makes it easier for designers to define the groups, roles and agent behaviour that should operate in that particular organisation section.

By distributing the normative information in the environment, it is possible to partition the environment in a functional way, thus helping the structured definition of large simulations, norms being associated only with the places where they are meant to be followed. It is also more efficient (by taking advantage of natural distribution) to have norms spread in an environment than having them in a repository made available for the whole society, as it is usually the case.

We believe that an explicit environment description is an important part of a multi-agent system because it is a stable point from where the agent reasoning and the organisational structures can be fine-tuned so as to facilitate the development of agents and organisations that can achieve their goals. The notion of *spatially distributed normative objects* that we have introduced here can be a good solution connecting definitions of

organisations and definitions of environments. Additionally, distributing the organisational/normative information can facilitate the modelling of large organisations.

It is interesting to note that, being conditioned on the possibility of checking the existence of a normative object, the normative reasoning required from agents that deal with normative objects is necessarily of a non-monotonic nature, and the experience of programming such reasoning in AgentSpeak is something we plan to experiment with in the future. Also as future work, we intend to allow a normative place to be associated with group structures, creating a connection between the organisational structures and the physical environment. We plan to make possible such association for any existing approach to agent organisations, such as $\mathcal{M}OISE^+$ [10], OperA/OMNI [15], GAIA [17], or approaches based on electronic institutions [6, 7]. The recursive nature of normative places may not be compatible, however, with some of such approaches to organisation, where the (possibly implicit) system of normative rules has no provision for a recursive structure in its operation.

Acknowledgements

This work was partially supported by CNPq, CAPES, and FAPERGS. Rafael Bordini gratefully acknowledges the support of The Nuffield Foundation (grant number NAL/01065/G). The authors greatly benefited from the comments and suggestions that arose in the discussions during the presentation of the paper at the COIN@ECAI2006 workshop in Riva del Garda.

References

1. Rafael H. Bordini, Antônio Carlos da Rocha Costa, Jomi F. Hübner, Álvaro F. Moreira, Fabio Y. Okuyama, and Renata Vieira, ‘MAS-SOC: a social simulation platform based on agent-oriented programming’, *Journal of Artificial Societies and Social Simulation*, **8**(3), (2005).
2. Rafael H. Bordini, Jomi F. Hübner, et al., *Jason: A Java-based Interpreter for an Extended Version of AgentSpeak*, manual, release version 0.9 edn., July 2006. <http://jason.sourceforge.net/>.
3. Rafael H. Bordini, Jomi F. Hübner, and Renata Vieira, ‘Jason and the Golden Fleece of agent-oriented programming’, in *Multi-Agent Programming: Languages, Platforms and Applications*, eds., Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni, chapter 1, Springer-Verlag, (2005).
4. C. Castelfranchi, R. Conte, and M. Paolucci, ‘Normative reputation and the costs of compliance’, *Journal of Artificial Societies and Social Simulation*, **1**(3), (1998). <<http://www.soc.surrey.ac.uk/JASSS/1/3/3.html>>.
5. Cristiano Castelfranchi, Frank Dignum, Catholijn M. Jonker, and Jan Treur, ‘Deliberative normative agents: Principles and architecture’, in *6th International Workshop on Intelligent Agents VI, Agent Theories, Architectures, and Languages (ATAL)*, Lecture Notes In Computer Science, Vol. 1757, pp. 364–378, Londo, (1999). Springer-Verlag.
6. Marc Esteva, David de la Cruz, and Carles Sierra, ‘Islander: an electronic institutions editor’, in *AAMAS*, pp. 1045–1052. ACM, (2002).

7. Marc Esteva, Bruno Rosell, Juan A. Rodríguez-Aguilar, and Josep Lluís Arcos, 'Ameli: An agent-based middleware for electronic institutions.', in *AAMAS*, pp. 236–243. IEEE Computer Society, (2004).
8. Jacques Ferber, Fabien Michel, and José-Antonio Báez-Barranco, 'Agre: Integrating environments with organizations.', in *E4MAS*, pp. 48–56, (2004).
9. Andrés García-Camino, Pablo Noriega, and Juan A. Rodríguez-Aguilar, 'Implementing norms in electronic institutions.', in *AAMAS*, eds., Frank Dignum, Virginia Dignum, Sven Koenig, Sarit Kraus, Munindar P. Singh, and Michael Wooldridge, pp. 667–673. ACM, (2005).
10. Jomi Fred Hübner, Jaime Simão Sichman, and Olivier Boissier, 'MOISE⁺: Towards a structural, functional, and deontic model for MAS organization', in *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'2002)*, Bologna, Italy, (2002).
11. Fabio Y. Okuyama, Rafael H. Bordini, and Antônio Carlos da Rocha Costa, 'ELMS: An environment description language for multi-agent simulations', in *Proceedings of the First International Workshop on Environments for Multiagent Systems (E4MAS)*, held with *AAMAS-04*, 19th of July, eds., Danny Weyns, H. van Dyke Parunak, and Fabien Michel, number 3374 in Lecture Notes In Artificial Intelligence, pp. 91–108, Berlin, (2005). Springer-Verlag.
12. A. Omicini, A. Ricci, M. Viroli, C. Castelfranchi, and L. Tummolini, 'Coordination artifacts: Environment-based coordination for intelligent agents', in *AAMAS'04*, (2004).
13. Anand S. Rao, 'AgentSpeak(L): BDI agents speak out in a logical computable language', in *Proceedings of the Seventh Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96)*, 22–25 January, Eindhoven, The Netherlands, eds., Walter Van de Velde and John Perram, number 1038 in Lecture Notes in Artificial Intelligence, pp. 42–55, London, (1996). Springer-Verlag.
14. Rossella Rubino, Andrea Omicini, and Enrico Denti, 'Computational institutions for modelling norm-regulated MAS: An approach based on coordination artifacts', in *1st International Workshop "Agents, Norms and Institutions for Regulated Multi-Agent Systems" (ANI@REM 2005)*, eds., Gabriela Lindemann, Sascha Ossowski, Julian Padget, and Javier Vázquez-Salceda, AAMAS 2005, Utrecht, The Netherlands, (25 July 2005).
15. Javier Vázquez-Salceda, Virginia Dignum, and Frank Dignum, 'Organizing multiagent systems.', *Autonomous Agents and Multi-Agent Systems*, **11**(3), 307–360, (2005).
16. Mirko Viroli, Andrea Omicini, and Alessandro Ricci, 'Engineering MAS environment with artifacts', in *2nd International Workshop "Environments for Multi-Agent Systems" (E4MAS 2005)*, eds., Danny Weyns, H. Van Dyke Parunak, and Fabien Michel, pp. 62–77, AAMAS 2005, Utrecht, The Netherlands, (26 July 2005).
17. Michael Wooldridge, Nicholas R. Jennings, and David Kinny, 'The GAIA methodology for agent-oriented analysis and design', *Autonomous Agents and Multi-Agent Systems*, **3**(3), 285–312, (2000).