

Multi-Society Organisations and Social Opacity: When Agents Play the Role of Observers

Nuno David^{1,2}, Jaime Simão Sichman² and Helder Coelho³

¹ Department of Information Science and Technology, ISCTE/DCTI, Lisbon, Portugal

² Intelligent Techniques Laboratory, University of São Paulo, Brazil

³ Department of Informatics, University of Lisbon, Portugal

This paper is submitted to another conference.

Abstract. Organisational models in MAS usually position the agents as plain actor-observers within environments shared by multiple agents and organisational structures at different levels of granularity. In this article we propose that the agent capacity to reason in environments with heterogeneous models of societies can be enhanced if the agent is positioned as an external observer to other organisational structures. To this end, we show that the delegation of the external observer's role to the agent is facilitated when one adopts organisational models that clearly circumscribe multiple opaque social spaces of interaction at the same level of abstraction.

Keywords: Multi-Agent Organisations, Multi-Agent Systems, Societies, Social Opacity, Observation

1 INTRODUCTION

The architecture of a multi-agent system (MAS) can naturally be seen as a computational organisation. The organisational description of a multi-agent society is useful to understand and improve the modularity and efficiency of the system, since the organisation constraints the agents' individual behaviour towards the system goals. To this end, several organisational abstractions are being increasingly proposed as a methodological tool to analyse, design and simulate MAS societies (see [20]). For instance, while some organisation models have put its emphasis on functional aspects of a society, like global plan specification, allocation of tasks, and coordination to execute plans (e.g.[14]); other models have emphasized the structural side of a society, in terms of specification of explicit organisational entities and relations among them, like groups, roles, authorities, permissions or obligations (e.g.[4]).

Meanwhile, while both research lines tend to commonly use the term *society* as an influential organisational metaphor to specify MAS (see [8]), its concept is rarely specified as an explicit structural and relational entity. Rather than an explicit entity, societies are implicitly defined in formally or informally terms of inclusiveness of multiple agents and other organisational structures, like communication languages, coalitions, norms and roles that agents are allowed to play (see, for instance [1]). This tendency comes from the general conceptual trend of conceiving autonomous agents as *internal actors* of societies, as opposed to the possibility of conceiving themselves also as neutral, *external observers*, creators, or even autonomous designers of one or more societies. Societies are then conceived as closed, possibly infinite, mutually opaque spaces, with a unique and omnipresent opaque observer in the person of the human designer. Although some alternative approaches have treated the problem of agent migration in models that explicitly define multiple societies [13], the underlying concept of society in those models is still reducible to one of a group, where agents are viewed simultaneously both as actors and non-neutral observers in a given society. Also in works with reactive agents [5] or simulation with cognitive agents (see [7,12]), where the stress is given to emergent organisational structures, the role of the observer is not assigned explicitly to agents, being exclusively and implicitly defined in the person of the system designer.

Nevertheless, in the real world, we have the ability to create explicit organisational structures and to reason about them, like other agents, institutions or even new societies (for example, artificial agent societies). Similarly, the artificial agent's ability to build topologies of multiple societies in a dynamic way can be very powerful. In some environments, especially in environments with cognitive agents, an important factor in the system dynamics is the agent's

beliefs and social reasoning mechanisms about other agents and the environment. The agent's skill to create and observe societies dynamically, possibly within a same or different level of abstraction than his own, corresponds to the ability to instantiate and observe *given* models of agents and societies in the world, allowing the agent to reason autonomously about the heterogeneity of different models at various levels of observation. Such capacity is especially important, for instance, in MAS models specified to observe and provide relevant visualization results of simulations of other self-motivated agent societies.

The problem of "agentified" autonomous design and observation is partially the problem of delegation of the human observer's role to the artificial agent. When an agent adopts the observer's role he should be able to create and observe dynamical aspects of organisational structures of other agents in other societies. In some situations, this means that the observer agent must have the ability to look inside the other agents' minds. In other situations it will even be useful to give the agent the ability to pro-actively influence or change the organisational structures and cognitive representations of other agents in other societies. But while the observed agents and societies must be visible to the observer at various dimensions, the observer must be *socially opaque* to the observed agents. The model that we propose in this paper characterizes an organisation of multiple societies that are explicitly defined and mutually visible, and where certain organisational configurations are able to dynamically manage different degrees of *social opacity* between these societies. A multi-society space is a multi-dimensional environment in which the agents are themselves capable of creating explicit organizational structures, like other agents and other societies. The problematic of *social opacity* configures the conditions under which the control of cognitive information transfer between these different societies is possible. While social opacity is straightforwardly attained when the observer and observed agent lie in distinct hierarchical levels of abstraction, it is especially important in multi-society spaces when the observer and observed agents are superimposed in social spaces at the same level of abstraction.

This paper is structured as follows. In section 2 we will present our model of multi-society spaces. In section 3 we will analyse two different organisational abstractions that can be used to circumscribe opaque social spaces in the multi-society model. In section 4 we will present an application example with regard to a simulation infrastructure based on opaque multi-agent organisations. Finally, in section 5 we present some related work and conclusions.

2 MULTI-SOCIETY SPACES

From the observer's point of view, the concept of society encloses the vision of a common interaction space that allows the agents to coexist and interact, and generates the conditions for the explicit or emergent design of organizational structures, like groups, role-playing or teamwork. Since a society may contain any number of such structures our concept of society belongs to a higher level of abstraction than those structures. Nevertheless, some of the social features of computational MAS are ultimately specified by a minimal set of organizational structures. In this sense, the explicit consideration of a society as an organizational entity is instrumental to generalize models of one society to models of many societies. The organisational model that we present next is somewhat influenced by the organisational model of [4], which is based on the core concepts of *agents*, *roles* and *groups*.

2.2 Multiple Society Spaces

Our Multi-society Space (MSS) model is based on four explicit organizational ingredients: *societies*, *agents*, *roles* and *role-players*. The components of a MSS are as follows:

- A set AGT of agents. Agents are active entities that are able to play a set of roles.
- A set ROL of roles. A role is an abstract function that may be exercised by agents, like different abilities, identifications or obligations.
- A set SOC of societies. A society is an interaction space that authorizes the playing of certain roles. An agent is allowed to enter a society and play a specific role if that role is *authorized* in that society. The partial functions $agtsoc: SOC \rightarrow P(AGT)$ and $rolsoc: SOC \rightarrow P(ROL)$ map a

given society, respectively, on the set of agents that are *resident* in the society and the set of *authorized* roles in the society.

- A set RPY of role-players. We distinguish roles from role-players. Role-players are the actual entities through which all agents act in the MSS. Each role-player is able to play a single role, but multiple role-players in the MSS can represent a same agent. For example, if the MSS is the planet earth and societies are nations, a possible situation for an agent with three role-players is to have a *Professor* role-player and a *Father* role-player in Portugal, and another *Professor* role-player in Brazil as well. In addition, every role-player holds a set of *delegable* roles that may be ascribed to other role-players upon its creation. We represent a role-player as a quadruple $rpy_i = (soc_i, agt_i, rol_i, R_i)$ composed by a society $soc_i \in SOC$, an agent $agt_i \in agtsoc(soc_i)$, a playing role $rol_i \in rolsoc(soc_i)$ and a set of *delegable* roles $R_i \in P(ROL)$. The partial function $delrol: RPY \rightarrow P(ROL)$ maps a given role-player on his set of delegable roles.

Definition 1. A MSS is a 7-tuple, $\langle AGT, ROL, SOC, agtsoc, rolsoc, RPY, delrol \rangle$, with components as above¹.

Agents interact in the MSS with others through *social events*, like message passing and creating new societies or other agents, as will be defined shortly. Agents can also be created and operate on behalf of *external applications*. An external application (EA) is an entity capable of creating agents or societies in the MSS but that is not explicitly represented in the MSS, such as the agents launching shell or a package for meta-analysis in a simulation environment.

One may see EAs represented in a different level of abstraction from the MSS. As a result, the transfer of information between agents can occur *explicitly* and internally to the MSS, through social events, or *implicitly* and externally to the MSS, via arbitrary interactions between agents, EAs, and again agents. For most of this paper we assume that *implicit* transfer of information does not take place. This is not always the case and we will refer to it when appropriate. Nevertheless, it will serve the purpose of concentrating the focus of the paper on the problem of observation and social opacity between different social spaces represented in the same level of abstraction.

2.2 Social Space Dynamics

Agents and EAs can modify the state of the MSS along the time through *social events*. Social events can take place by the initiative of agents or EAs. If a social event is on an agent's initiative, it must occur by means of his role-players. We call this role-player the *invoker* role-player. External applications originate social events when they wish to launch agents and societies in the MSS.

Given a MSS in state k , the occurrence of a social event will modify its state. We record the state of the MSS with a superscript like MSS^k when referring to the MSS in state k . The occurrence of a social event $MSS^k \rightarrow MSS^{k+1}$ depends on a set of pre-conditions that have to be satisfied. In the following, the character $*$ next to a pre-condition denotes that is not applicable if the event is originated by EAs. In this paper, we define four social events $SE1, \dots, SE4$ as follows:

SE1: Society creation. Role-players or EAs can invoke the creation of societies. Given a set of intended authorized roles in the society, it may be the case that these roles are not yet defined in the MSS. The creation of a society that authorizes the playing of a set of roles R_j , will create a new society $soc_j \notin SOC^k$ and eventually a new set of roles in the MSS:

$$MSS^k \xrightarrow{SE1} MSS^{k+1} \mid \begin{array}{l} agtsoc^{k+1}(soc_j) = \emptyset, rolsoc^{k+1}(soc_j) = R_j, \\ SOC^{k+1} = SOC^k \cup \{soc_j\}, ROL^{k+1} = ROL^k \cup R_j \end{array}$$

¹ Notice that our concept of *society* is different from a *group* as specified by [5], since we do not pre-specify the set of interactions that can appear within each society. Additionally, the organisational model of [5] does not consider the concept of *role-player*, which will be instrumental to analyse the conditions under which the control of cognitive information transfer between these different societies is possible.

SE2: Agent creation / SE3: Role-player creation. Agent creation refers to the instantiation of new agents in the MSS, invoked by role-players or EAs. The instantiation of a new agent requires the creation of a new role-player in some target society. However, if an agent is already instantiated in the MSS, a similar event is the creation of additional role-players in target societies, which cannot be invoked by EAs. This event occurs when agents want to join additional societies with new role-players or want to be represented in a same society with additional role-players. We restrict our specification to agent creation, since role-player creation has a similar specification. We use the subscript i when referring to the creator agent and the subscript j when referring to the new agent. If the social event is on an agent's initiative, consider its invoker role-player rp_{y_i} . The creation of a new agent, playing the target role rol_j , with delegable roles R_j , in the target society $soc_j \in SOC^k$, generates a new agent $agt_j \notin AGT^k$, a new role-player $rp_{y_j} = (soc_j, agt_j, rol_j, R_j)$ and, possibly, a new set of roles in the MSS:

- if, (c1) $rol_j \in rol_{soc}^k(soc_j)$, the target role rol_j must be authorized in the target society soc_j ;
 (c2*) $rol_j \in del_{rol}^k(rp_{y_i})$, the target role rol_j must be delegable by the invoker role-player rp_{y_i} ;
 (c3*) $R_j \subseteq del_{rol}^k(rp_{y_i})$, the target set of delegable roles R_j must be a subset of the invoker role-player rp_{y_i} delegable roles.

$$MSS^k \xrightarrow{SE2} MSS^{k+1} \mid \quad AGT^{k+1} = AGT^k \cup \{agt_j\}, \quad ROL^{k+1} = ROL^k \cup R_j, \\ RPY^{k+1} = RPY^k \cup \{rp_{y_j}\}, \quad del_{rol}^{k+1}(rp_{y_j}) = R_j, \quad agt_{soc}^{k+1}(soc_j) = agt_{soc}^k(soc_j) \cup \{agt_j\}$$

SE4: Message passing in a society. Only role-players can originate this social event, therefore excluding any EA. Message passing in the MSS does not alter its structure, but the sender and receiver role-players must operate in the same society.

The particularity of a multi-society space is the possibility of creating multiple societies in the same level of abstraction: an agent may be the creator of a society and also its member; and a member of the created society can be a member of the creator agent's society². In effect, while role-players are only allowed to communicate between each other if they share a same society, a same agent can act with multiple role-players across multiple societies. As a result, societies are not *opaque* relative to each other, in terms of information transfer between agents residing in different societies.

3 SOCIAL SPACES AND OPACITY

The set of available social events and pre-conditions for its invocation determines the initial conditions to analyse the opacity between different societies. Opacity is also fundamentally dependent on the organisational dynamics. Ultimately, if an agent ever resides in more than one society during his life cycle, opacity will depend on the internal agent architecture with respect to the playing of its different roles in different societies. In general, we characterize the opacity of a society according to information transfer conditions from the *inside* to the *outside* of a society. Of course, the dilemma here is to identify what the inside and the outside of a society are.

3.1 Visibility Dimensions

To begin with, we analyse the opacity of a society along three dimensions:

Organisational visibility – relative to the access, from the outside of a society, to organizational properties of the society in the global context of the multi-society space, like its physical location or shape. For instance, a valley that appears to be the environment of an isolated tribe in the Amazon may become identifiable by a satellite photograph, even though we may have no relevant information coming from inside the tribe. In our MSS, this is inherently obtainable through the invocation of social events that create organizational structures. The identification of societies, agents, roles and role-players is always visible to its creator, and may become visible by others through message passing.

² Other social events may be defined: for instance, a role-player may be able to define new roles in the MSS and add these roles to his set of delegable roles.

Openness – relative to organisational conditions, prescribed by the MSS designer, or subjective conditions, prescribed by the agents inside the society, restricting agents in the outside from entering the inside. These conditions may vary extensively, for instance, according to some qualified institutional owner (an human or artificial agent), which decides if some given agent may or may not enter the society. In our MSS, the openness of a society will ultimately depend on the level of convergence between the set of authorized roles in that society and the set of delegable roles accessible to each agent's role-player.

Behavioural and Cognitive visibility – relative to the access, from the outside of the society, to behaviours or cognitive representations of agents in the inside. Behavioural visibility concerns the observation of social events; for instance, a spy satellite may try to scout the transmission of messages between agents of a competitor country. Cognitive visibility concerns the observation of internal representations of agents, such as its goals and beliefs. In our MSS, behavioural and/or cognitive visibility implies the superposition of agents in the inside and the outside of a society. This is a necessary condition, but not a sufficient condition. Moreover, the agents' ability to observe internal representations of other agents may be *prima-facie* counter-intuitive, since that is not the usual idea of a multi-agent system. As we will soon show, other mechanisms must be designed in order to provide cognitive visibility in the MSS.

The three dimensions are not independent from each other. Suppose we have an MSS with two societies and there is not a single agent residing simultaneously in both societies. The organizational and cognitive visibility of one society relative to agents in the other society will vary according to the existence of a potential *bridge* agent in the latter society able to join the former society. Note also that an agent can join a society by accident, even if the society is not known or visible to the agent. Our interest is to identify circumscriptions of spaces in the MSS so as to fix some of these dimensions and analyse independently the other dimensions. In this sense, the concept of opacity is related to the difficult problem of circumscribing the *internal* from the *external* environment of a society.

The circumscription of the internal environment depends essentially on two factors: (1) objective organizational conditions associated with the dynamic structure of the MSS and independent from the agents' internal representations, like communication or role playing conditions, and (2) different internal representations emerging cognitively [2] within each member, relative to his own individual perception about the range of its social environment, like for instance dependence relations [18,3,17,]. Our interest is to fix circumscriptions along the first factor so as to control the range of possible circumscriptions based on the second factor. We classify the internal space of a society along two vectors, respectively, communication and role-playing conditions.

3.2 Communication Opacity

The internal communication space of a society is defined according to communication conditions between agents that are resident and agents that are not resident in that society. To begin with, consider the *Plane Communication Space* (PCS) of a society. The PCS circumscribes all role-players that are able to communicate directly with each other using message passing inside the society, that is, inside the society plane boundaries.

Plane Communication Space. The PCS of a society $\text{soc}_j \in \text{SOC}$ is the set of all role-players in that society:

$$\text{PCS}(\text{soc}_j) = \{ (\text{soc}_i, \text{agt}_i, \text{rol}_i, \text{R}_i) \in \text{RPY} \mid \text{soc}_j = \text{soc}_i \}$$

Agents playing roles in a society may also play roles outside that society. The *Internal Communication Space* (ICS) of a society superset the PCS, by including role-players in the outside of the society if the corresponding agents are members of that society.

Internal Communication Space. The ICS of a society $\text{soc}_j \in \text{SOC}$ is the set of all role-players in the MSS controlled by agents who are members of that society:

$$\text{ICS}(\text{soc}_j) = \{ (\text{soc}_i, \text{agt}_i, \text{rol}_i, \text{R}_i) \in \text{RPY} \mid \text{agt}_i \in \text{agtsoc}(\text{soc}_j) \}$$

Pure Internal Communication Space. The ICS of a society $\text{soc}_j \in \text{SOC}^k$ in state k is *pure* if for any state i , with $i \leq k$, the ICS coincides with the PCS.

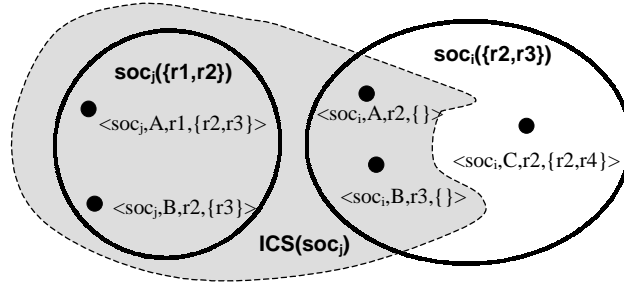


Figure 1. Non-pure ICS of society soc_j .

In figure 1 we represent a non-pure ICS relative to society soc_j . There are two societies – soc_j and soc_i – and three agents – A, B and C. Each point represents an agent's role-player, for several points may represent an agent. For example, the role-player $\langle \text{soc}_j, A, r1, \{r2, r3\} \rangle$ is the agent A in society soc_j playing role $r1$ with delegable roles $\{r2, r3\}$. Society soc_j authorizes roles $r1$ and $r2$, and society soc_i authorizes roles $r2$ and $r3$. The PCS of society soc_j is $\text{PCS}(\text{soc}_j) = \{ \langle \text{soc}_j, A, r1, \{r2, r3\} \rangle, \langle \text{soc}_j, B, r2, \{r3\} \rangle \}$. The ICS is non-pure because agents A and B are playing roles in both societies.

If for some state an agent resides simultaneously in two societies, the ICS of either society will be circumscribed outside the boundaries of the PCS, encompassing role-players of both societies. On the contrary, the ICS of a society is *pure* if there was never an agent with role-players in the society that has ever had role-players in any other society. Consequently, there was never an agent with role-players outside the society that has ever communicated with agents that are or have been resident inside the society. Note, however, that a *pure* ICS is not a sufficient condition to ensure the opacity of a society, at least in terms of organizational visibility. To this end a set of organizational conditions must be set in order to preclude agents outside the society to identify it and eventually create new agents within it.

Consider a society and a set of resident agents created by an EA. Suppose that (1) *implicit* information transfer between agents through the EA is not possible, in other words, the society is not visible to the outside right after its creation; (2) the agents inside the society are benevolent and cannot join other societies according to their design specification; and (3) the organizational conditions do not ever allow for role-players outside the society to create role-players inside the society, in other words, the society is closed on the outside. The last condition can be achieved if all delegable roles outside the society are different from all authorized roles in the society. Since no agent in the outside will ever reside simultaneously inside and outside the society, the society ICS will be *pure* and the society opacity will not depend on cognitive information transfer through the agents' internal architectures. The society organizational visibility will exclusively depend on *implicit* information transfer through the EAs³, but still preventing agents outside the society to enter the inside. Nevertheless, it is precisely the closeness character of the society and the impossibility of explicit information transfer in the MSS from the inside to the outside of such a society that makes its range of practical applications limited, restricted to systems where the agents are explicitly designed to co-operatively achieve a given set of goals.

3.3 Role-Playing Opacity

Another way of circumscribing a space of a society is to make use of role playing conditions. Differently from the previous case, the composition of communication and role-playing conditions allows an agent to be represented simultaneously through different role-players in the internal and external *pure* space of a society. The purpose of using role-playing conditions is to

³ Or possibly through the underlying MSS internals, i.e., the system. Such security problems are not our concern here. We are thus assuming that the system is opaque to the MSS.

control the society's opacity through the agents' internal architectures, but nevertheless easily manageable through the agents' internal mechanisms that control the agents' role-playing activity. The *Internal Role-playing Space* (IRpS) of a society **subsets** the ICS by excluding role-players outside the society that do not have its playing roles authorized in the society:

Internal Role-playing Space. The IRpS of a society $\text{soc}_j \in \text{SOC}$ is the set of all role-players in the ICS that have its playing roles authorized in that society:

$$\text{IRpS}(\text{soc}_j) = \{ (\text{soc}_i, \text{agt}_i, \text{rol}_i, R_i) \in \text{ICS}(\text{soc}_j) \mid \text{rol}_i \in \text{rolsoc}(\text{soc}_j) \}$$

Pure Internal Role-playing Space. The IRpS of a society $\text{soc}_j \in \text{SOC}^k$ in state k is *pure* if for any state i , with $i \leq k$, the IRpS coincides with the PCS.

Figure 2a illustrates a non-pure ICS and a non-pure IRpS relative to society soc_j . The IRpS is non-pure because agent A is playing role $r2$ in society soc_i , whereas role $r2$ is also authorized in society soc_j . The difference between a *non-pure* IRpS and a *pure* IRpS is that in the first case an agent can play a same role inside and outside the society. An IRpS of a society is kept *pure* if the set of agents that have role-players in that society do not control other role-players in the outside whose playing roles are authorized in the inside. But differently from a *pure* ICS, opacity will now depend on the agents' internal architectures, relative to the transferability of cognitive representations between the playing of an agent's *different* roles. Figure 2b illustrates a possible state for a *pure* IRpS.

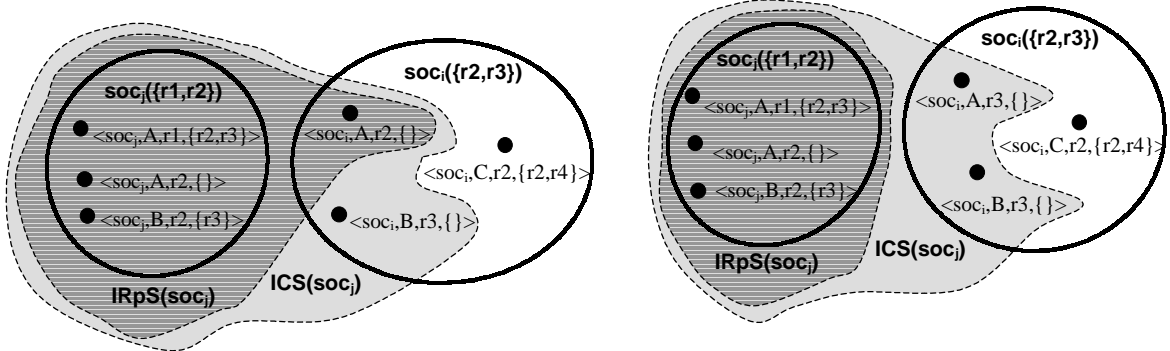


Figure 2a. Non-pure ICS and IRpS of society soc_j .

Figure 2b. Non-pure ICS and pure IRpS of society soc_j .

The purpose of circumscribing role-playing spaces is to produce a flexible mechanism to design different organizational topologies of opaque *observation* spaces and non-opaque spaces, according to role-playing conditions, that can be autonomously prescribed by the observer agent. Since the agents themselves can create other agents, roles and societies, the compartmentalization topology of social spaces may assume different configurations in a dynamic way. This means that the MSS itself can assume an emerging autonomous character from the human designer with respect to its own topology, as well as to its different points for opaque observation of social spaces.

4 MOSCA: AN OPAQUE ORGANISATION

The example that we illustrate in this paper is motivated by the field of MAS simulations (see [7,12]), especially simulation of cognitive agents. In such simulations it is often the case that the simulated organisational setting and the agents' behavioural rules or internal cognitive representations have to be observed, or even enforcedly modified, during the simulation. The goal is to design such a simulator based on MAS organisations. Regarding the motivation and operational examples for simulators based on MAS organisations we refer to [11,19,6,10]. The design is conceived to run on the top of our MAS development platform [15].

A society that should desirably have an opaque IRpS to other societies is the multi-agent organization that plays simultaneously the role of computing the simulation of another multi-agent organization *outside* the IRpS, and the role of observing and controlling the computation of the latter from the *inside* of the IRpS. In MOSCA (Meta-organisation for Simulation of Cognitive Agents) an opaque multi-agent organization must be able to reproduce a target multi-agent organization of self-contained agents in controllable conditions. The simulation of MAS

societies requires one MOSCA agent and two basic roles for *each* target agent intended as object of simulation: the *Control* and *Generic* role.

The *Control* role is played exclusively within a society or set of societies (a region) called *S_Control*, with a *pure* IRpS, whereby MOSCA agents co-operate for a common goal: to reproduce in a MAS distributed environment the behaviours of agents that are the real targets of simulation in a controllable way and **outside** the IRpS of the *S_Control* society. The set of societies outside the IRpS of *S_Control* is called the *Arena*. Hence, each MOSCA agent plays at least two roles expressing distinct behaviours: (i) the behaviour of a benevolent agent that cooperates with other MOSCA agents in the *S_Control* society, expressed exclusively through the *Control* role-player, in order to observe and maintain a consistent world state in the *Arena*, and (ii) a given arbitrary behaviour, expressed exclusively through the *Generic* role-player in the *Arena*, which is the effective target of simulation. To this end, the MOSCA agent architecture implements an internal state machine that reproduces the target agent's social events through the *Generic* role. Hence, the computation of the target agent will exclusively evolve according to the occurrence of social events in the *Arena*.

Besides reproducing the target agent's social events in the *Arena*, the MOSCA agent must respond to the users' observation and intervention requests, such as observing and reporting social events or changing the target's internal states (beliefs, goals, intentions) throughout the computation. Owing to the distributed character of the environment and to observation and intervention activities, each social event invoked by the *Generic* role-player will imply a contingency set of social events invoked by the *Control* role-player within the *S_Control* society.

The key to maintain observation and intervention activities opaque to the *Arena* is an adequate manipulation of organisational visibility, authorized roles in the *S_Control* society, and delegable roles in the *Arena*. While the *S_Control* society must be kept not visible in the *Arena*, its IRpS must be kept *pure* for the duration of the simulation. Suppose the goal is to simulate a particular MAS organization, which we call the *target application*. Initially, the MSS is initially empty and MOSCA is an external application (EA). The simulation proceeds as follows:

Stage A. Launching MOSCA

- (1) The MOSCA EA loads the target application script that specifies the target agents (a set of Java classes), the society and delegable/authorized roles that must be launched to start the target application. We call *S_Arena* to the target society, *Generic* to every target role, and *generic-player* to role-players playing the role *Generic*.
- (2) Subsequently, the MOSCA EA invokes the creation of a society called *S_Control* with a single authorized role called *Control*. As a result, the *S_Control* society will be liable to the playing of a single role: the *Control* role. We use the name *control-player* when referring to role-players that play the *Control* role.
- (3) The MOSCA EA creates an agent called *Guardian* in the society *S_Control*. The *Guardian control-role* includes in his set of delegable roles the *Control* and *Generic* roles. The purpose of the *Guardian* is to coordinate the simulation with other MOSCA agents, while safeguarding the *S_Control* opacity to the *Arena*.

Stage B. Launching the Target Application

- (3) The *Guardian* creates the target society *S_Arena* where the simulation will initially take place, with authorized role *Generic*. Subsequently the *Guardian* creates a set of agents in the society *S_Control* that we call *Monitors*. Each *Monitor control-player* includes in his set of delegable roles the *Generic* but it does not include the *Control* role. This means that they are not able to create other *control-players*. Nevertheless, the *Monitors* are benevolent agents with a well-defined specification: to cooperate with the *Guardian* and other *Monitors* in order to reproduce in a controlled way the target application in the society *S_Arena*.
- (4) In the *S_Control* society, the *Guardian* notifies each *Monitor* about the target agents, delegable roles and the target society where the targets will be created.

Stage C. Running the Simulation

(5) At this point, the *Monitors* are ready to create and execute the target agents, expressing its social events through the society *S_Arena*, or any society that may be created during the simulation. Each social event invoked by the *generic-players* may be preceded by a contingency set of social events in the society *S_Control*. The *Monitors* cooperate between themselves in the *S_Control* society, in order to maintain a controlled and synchronized simulation world in the *S_Arena* societies.

According to these conditions the IRpS of *S_Control* will be pure: since the *Control* role is not delegable by role-players outside the society, the target agents in the *Arena* will never be able to join the society. Moreover, since the *Control* role is not visible to the computation of the *Generic* role, the *S_Control* society will be opaque.

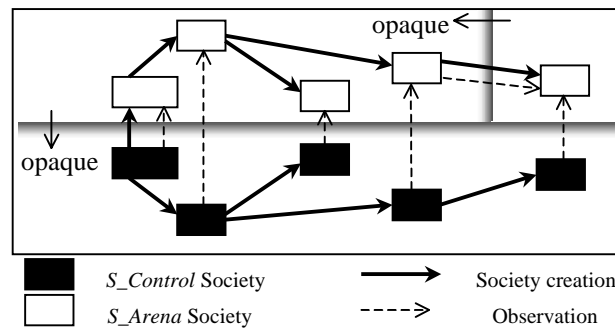


Figure 3. Mirror topology.

Another important point that should be pointed out is that while the algorithm illustrates the creation of a single *S_Control* society, role-playing opaque societies can be easily generalised to opaque regions that encompass a set of mutually visible *S_Control* societies. This is useful if one wants to distribute various points of observation according to the emergent topology of multiple societies in the *Arena*. Different points of observation are particularly important in large simulations, for reasons of computational and bandwidth resource distribution. Modularity is also an issue. One can distribute different control societies according to different groups of targets, associated with an independent logical or physical pattern of execution, like different simulation step algorithms (discrete time, event based...). In figure 3 we illustrate an example with a control region that strictly follows a *mirror* topology: for every new society created in the *Arena* a society and a corresponding *Guardian control-player* are created in the control region. Notice a second opaque region in the *Arena*: target agents in the *Arena* can create recursively their own observation spaces, but they still will be liable to observation in the control region.

5 SUMMARY AND RELATED WORK

Organisational models in MAS usually position the agents as plain actor-observers within environments shared by multiple agents and organisational structures at different levels of granularity. In this article we have proposed that the agent capacity to reason about heterogeneous models of societies can be enhanced if the agent is positioned as an opaque external observer to other organisational structures. To this end, we have showed that the delegation of the observer's role to the agent is facilitated when one adopts organisational models that explicitly circumscribe spaces of interaction at the same level of abstraction. Nevertheless, we have also shown that the right set of organisational conditions must be found in order to elect the agent as a socially opaque observer.

We have exemplified how the model can be applied to the design of MAS simulators based on MAS organizations. Regarding this example, a related work that deserves special attention is the Swarm [9] simulation system. The Swarm model accommodates multi-level modelling approaches in which agents can be composed of *swarms* of other agents in nested hierarchies of abstraction. A parent swarm can observe other swarms and swarms in the same nested level can interact *implicitly* if they share a same parent swarm, avoiding any explicit visibility between societies in the same level of abstraction. This is partly because the observer agent is

represented within a different level of granularity from the observed agent. In contrast with the flexibility of our model, the interaction between different societies is therefore not transversal, since swarm agents cannot create other agents in their own or other societies by themselves. However, a limitation in our model in the context of MAS simulation is that flexibility may be gained at the cost of efficiency loss, partly because the simulator infrastructure must be based on MAS organisations.

The idea of multi-society spaces has elsewhere been proposed in somewhat of a different approach in [16]. The authors do not provide an operational framework, but speculate around the convenience of creating societies in the conceptualisation context of emergence and multiple viewpoints analysis. For this purpose they hypothesize the usefulness of creating apprehensible micro-macro links in MAS, by giving the agents the means to become aware of their mutual interaction, giving birth to new types of agents and societies out of their collective activity. As the latter authors we believe that the answer to build interesting MAS is the creation of environments capable of showing spontaneous emergence along multiple levels of abstraction, while being liable to the construction of explicit organisational structures in order to actively observe, and eventually manipulate, such emergent structures in the *same* level of abstraction. The model that we have presented here is a valuable and original starting point to this end. However, many questions are still left to answer. For instance, the organisational conditions to achieve social opacity in our model were relatively straightforward to find. But what about more complex models that can possibly have a higher number of different social events? While we believe that an alternative axiomatic definition to derive opacity conditions is the answer to analyse more complex models, we also think that it will be unfeasible in very high complex and dynamical domains. In the future we plan to investigate alternative methodologies to access the problem of observation and social opacity in complex dynamic organisations.

5 REFERENCES

- [1] Artikis A. and Pitt J., *A Formal Model of Open Agent Societies*. In Proceedings of Agents 2001.
- [2] Castelfranchi C., *The Theory of Social Functions: challenges for computational social science and Multi-Agent learning*. Journal of Cognitive Systems Research 2, 5-38, 2001.
- [3] David N., Sichman J.S. and Coelho H., *Agent-Based Social Simulation with Coalitions in Social Reasoning*. In [12], pp.244-265..
- [4] Ferber J. and Gutknecht O., *A Meta-model for the Analysis and Design of Organizations in Multi-Agent Systems*. In ICMAS'98, IEEE Computer Society, 1998.
- [5] Ferber J., *Reactive Distributed Artificial Intelligence: Principles and Applications*. In G. O'Hare and Jennings N., editors, Foundations of DAI, 1996.
- [6] Gasser L. and Kakugawa K., MACE3J: Fast Flexible Distributed Simulation of Large, Large Grain MAS. Accepted to AAMAS'02.
- [7] Gilbert N., Sichman J.S. and Conte R. (eds), *Multi-Agent Systems and Agent-Based Simulation*, Springer Verlag, LNAI1534, 1998.
- [8] Huhns M. and Stephens L.M., *Multiagent Systems and Societies of Agents*. In Weiss G (editor), *Multi-Agent Systems – A Modern Approach to AI*, MIT press, 79-114, 1999.
- [9] Langton C., Minar N., and Burkhart R., *The Swarm Simulation System: A Tool for studying complex systems*.
- [10] Marietto M.G., David N., Sichman J.S. and Coelho H. (2002), *Multi-Agent Simulation Platforms: State of the Art and New Prospects*, Accepted to MABS/AAMAS2002, see <http://www.lti.pcs.usp.br/SimCog>.
- [11] Michel F., Ferber J., and Gutknecht O., *Generic Simulation Tools Based on MAS Organization*. Proceedings of MAAMAW'01, 2001.
- [12] Moss S. and Davidsson P. (eds), *Multi-Agent-Based Simulation*, Springer-Verlag, LNAI1979, 2000.
- [13] Norbert G. and Philippe M., *The Reorganization of Societies of Autonomous Agents, Multi-Agent Rationality*, Proceedings of MAAMAW'97, 1997.
- [14] Prasad M., Decker K., Garvey A. and Lesser V., *Organisational Design with TAEMS: A case study of distributed data processing*. In Proceedings of ICMAS'98, 283-290, 1998.
- [15] SACI, Simple Agent Communication Infrastructure. <http://www.lti.pcs.usp.br/saci/>
- [16] Servat D., Perrier E., Treuil J.P. and Drogoul A., *When Agents Emerge from Agents: Introducing Scale Viewpoints in Multi-agent Simulations*. In [7], 183-198, 1998.
- [17] Sichman J.S. and Conte R., *Multi-Agent Dependence by Dependence Graphs*, Accepted to AAMAS'02.
- [18] Sichman J.S., Conte R., Demazeau Y. and Castelfranchi C., *A Social Reasoning Mechanism based on Dependence Networks*. In Proceedings of ECAI, 188-192, 1994.
- [19] Uhrmacher A., Fishwick P. and Zeigler B. (eds), *Agents in Modeling and Simulation: Exploiting the Metaphor*, Proceedings of IEEE, Vol.89, No.2, 127-213, 2001.
- [20] Zambonelli F., Jennings N. and Wooldridge M., *Organisational Abstractions for the Analysis and Design of MAS*, LNCS1957, Springer Verlag, 2001.