

Specification by refinement and agreement: designing agent interaction using landmarks and contracts

Hans Weigand¹, Virginia Dignum^{2,3}, John-Jules Meyer³, Frank Dignum³

¹Infolab, Tilburg University, PO Box 90153,
5000 LE Tilburg,, The Netherlands
weigand@kub.nl

²Achmea, P.O. Box 866,
3700 AW Zeist, The Netherlands
virginia.dignum@achmea.nl

³University Utrecht, PO Box 80089,
3508 TB Utrecht, The Netherlands
{virginia, jj, dignum}@cs.uu.nl

Abstract. In this paper, we argue that multi-agent systems developed to model and support organizations must on the one hand be able to describe and realize the goals and structure of that particular organization and on the other hand allow for autonomous behavior of participating actors. The agent autonomy should include collaboration autonomy, which means that agents can decide among each other how they want to cooperate. We present a model for agent societies and two techniques that can be used to achieve the described objectives: landmarks and contracts. The effects of these techniques on different agent society types are explored as well.

1. Introduction

Agent Societies emerge from the idea that agent interactions occur not just by accident but aim at achieving some desired global goals. That is, there are goals external to each individual agents that must be reached by the interaction of those agents. Desired behavior of a society is therefore external to the agents. Furthermore, we assume that in such a setting, social structure is determined by organizational design and not dependent on the agents themselves. We use the term Organizational Model (OM) to describe this global behavior and structure.

However, the behavior of individual agents is motivated from their own goals and capabilities, that is, agents bring in their own ways into the society as well. The actual behavior of the society emerges from the goal-pursuing behavior of the individual agents within the constraints set by the Organizational Model. From the society point of view this creates a need to check conformance of the actual behavior to the desired behavior and this has several consequences. Firstly, we have to make explicit the commitments between agents and the society. An actor is an agent performing one or more roles in the society. The objectives and capabilities of the role are described in

the OM and the actual agreements of the actor concerning its interpretation of the role are explicitly described in a social contract. We use the term Social Model to refer to the social contracts that hold at a given moment in a society. Secondly, actual interaction between actors must also be made explicit, which can be done through (bilateral) contracts as well. We call this the Interaction Model (IM). Checking the conformance of actual behavior to the desired behavior can now be realized in two steps:

- Checking whether the contract set up by two or more agents conforms to the OM. The OM can be more or less elaborate, depending on the type of society. Typically, it does not provide more than a few “landmarks” that describe the main features (conditions, obligations, and possibly partial plans) of interaction between roles.
- Monitoring whether the actual messaging behavior conforms to the contract. This is primarily a responsibility of the agents themselves. Depending on the type of society, the OM can constrain the monitoring or even assign it to the society.

In an Agent Society modeled from an organizational perspective, both the specific architecture of the individual agents and their motives to perform a role are ignored. In principle, an agent societies model is possible to be populated by any sort of agent (that is, first you have agents and then they decide to join a society), but the role description must be rich enough to allow for the design of agents specific to enact that role (that is, first you have a society and then you design agents to fulfil its roles).

The structure of this paper is as follows. In section 2, we describe the Agent Society model that we start from. Section 3 discusses the requirements on the specification of the interaction model that result from the society characteristics. Sections 4 and 5 present respectively, an Organizational Model and a contract concept that can meet the requirements. In section 6, the effects of these techniques on the different agent society types (markets, networks, hierarchies) are explored. Section 7 is the conclusion in which we evaluate the results and indicate some topics for future research.

2. Agent societies

The aim of agent societies is to describe multi-agent systems from an organizational (collective) perspective. Our work combines a society model similar to [2] and [12] with contracts, as a formalism to specify cooperation between agents and society, as proposed by [17], [6] and [1]. In [11] we introduce a model for agent societies, which takes an organizational perspective and describes a society in three levels:

- **Organizational model (OM):** describes the desired or intended behavior and overall structure of the society from the perspective of the organization in terms of roles, interaction scripts and social norms
- **Social model (SM):** populates the organizational model with specific agents mapped to roles through a social contract. Social contracts describe the agreed behavior for an agent within the society in terms of externally observable events. Agents enacting a role are called actors.

- **Interaction model (IM):** specifies interaction agreements between actors. Describes the actual behavior of the society.

Figure 1 depicts the different models. The model provides a separation between the intended, ideal behavior of the society as sought by the designer (organizational model) and the actual emergent behavior resulting from the autonomous action of participating agents (interaction model). We focus on designed societies (that is, societies that are explicitly developed for a certain aim, in opposition to societies that emerge from the local co-existence of a group of agents). In these societies the organizational aims are usually identified a priori and the society model must allow for predictability and verification. The organizational model describes the intended interactions between agents that from the perspective of the society designer will lead to the expected behavior of the society as a whole. However, actual interactions depend on the behavior and intentions of the actual agents involved at a given moment of the society life.

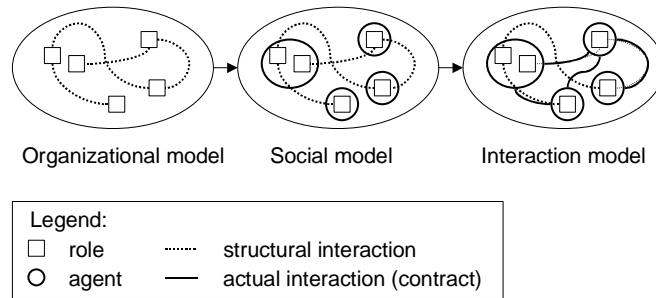


Fig. 1. Organizational framework for agent societies

We assume that the designer of the society will usually have no control over the design of the participating agents (open societies). Therefore, verification of society goals cannot be based on assumptions about the internal construction of agents. The social model describes explicitly the agreements between an individual agent and the society concerning the expected/allowed activity of the agent from the point of view of the society. These “labor contracts” will specify obligations and rights that an agent is entitled to as enactor of a society role (that was described in the organizational level). In principle, agents are free to pursue other of its own goals as far as those are not in violation of its social contract. In [5] we consider the implementation of agent societies in a given agent language (3APL, [13]), that is, how can the architecture of individual agents incorporate social concepts described in an organizational model of agent societies.

Actual interactions between agents are specified in the interaction model. These interactions can follow fixed protocols (e.g. the contract net protocol) but can be the result of communication between agents in which the interaction forms, results and duties for all parties have been negotiated. The type of interaction and the freedom for agents to negotiate specific interaction forms depends of the society model. In network models it is possible to create new interactions but market and hierarchical societies provide at least some fixed forms.

In order to develop a common understanding and a feeling of the possible applications of the agent society framework, it is necessary to develop scenarios in different areas, such as e-commerce, co-operative work and mediation. The agent society framework is especially suitable in situations where control and initiative are distributed and social order is an issue. In human organizations and societies, norms and conventions are used to specify the behavior that society members are expected to conform to. Such norms are usually backed by a variety of social institutions that enforce law and order, monitor for and respond to emergencies and prevent and recover from unanticipated disasters. In this way, citizens of civilized societies can utilize relatively simple and efficient rules of behavior, offloading the prevention and recovery of many problem types to social institutions that can handle them efficiently and effectively by virtue of their economies of scale and widely accepted legitimacy. The virtue of an institution resides in its potential to lend legitimacy to its members and security to its traders by establishing norms. Several researchers have recognized that the design of agent societies can benefit from abstractions analogous to those employed by our robust and relatively successful societies and organizations [4, 7, 12]. Electronic institutions can engender trust through certification of an agent and by the guarantees that it provides to back collaboration. Furthermore, the electronic institution can also function as the independent place in which all types of agent independent information about the interaction between the agents within the society is stored [8].

3. Requirements for interaction specification

In order to represent interactions between agents in such a context, a framework is needed that is able to specify sequences of interaction scenes

- independently from the internal design of the agent (internal autonomy requirement) and
- without fixing the interaction structures completely in advance (collaboration autonomy requirement).

Contracts are an adequate means for the explicit specification of interactions [17]. Contracts describe externally observable events and the contract specification language is independent from the internal agent specification.

As in human societies, existing norms and conventions do not cover all possible interactions between agents. That is, it is not possible to completely specify all interactions between agents a priori because we cannot account for environment and agent change during the life of the society. Therefore, there is a need to allow for the 'run-time' specification of new interaction scenes. If contracts are used as a specification mechanism, this implies that the society must be able to provide partial specifications of contracts. In principle, this can be done in several ways. For example, Kumar and Cohen describe interaction scenes using landmarks, which are sets of propositions holding in the state [14]. The landmarks are specified beforehand, but it is left to the agents how to achieve these landmarks. Some interaction scenes may allow for the negotiation of more specific interaction scenes, which involves the

definition of new interaction scenes and landmarks (specialization of a conversation protocol in terms of [14]).

Fundamentally, a tension exists between the goals of the society designer and the autonomy of the agents. When the society designer can specify the agent interactions in more detail, it is possible to check and guarantee more of the requirements at design time. Formal analysis of the interaction structures is possible, for example, to check basic audit principles as in [3]. It may be necessary to ensure the legal validity of the interactions that certain rules are followed. It can also be important for new coming agents that want to join the society that they know the norms and the interaction protocols used.

On the other hand, there are good reasons to leave the agents with some degree of freedom. We call this the collaboration autonomy requirement, which basically says that the agents who want to collaborate are free in choosing their own way of achieving this collaboration. Advantages of collaboration autonomy are:

- **Extensibility:** not everything needs to be, and often cannot be known beforehand. The society can evolve smoothly when new interactions or ways to perform these interactions are being developed.
- **Flexibility:** specifying a certain interaction structure beforehand often means that design choices are made that are not really necessary. Forestalling these choices gives the agents the possibility to make these choices themselves, geared to their particular situation.
- **Perusability:** not only is it possible to develop new interaction structures within the society, but it is also easier to reuse interaction structures across societies.

It must be clear that the collaboration autonomy (as any autonomy requirement) is always relative. Certain rules or structures must be present, so that the agents can build on top of that. For example, in a market society it must be possible to arrive at a transaction. It is often not important how the transaction is achieved, but it must be clear at some point that a business transaction has been closed. Sometimes it is important how the transaction is achieved. If the number of suppliers and/or buyers is very large and efficiency is critical, an auction mechanism is most effective, and this means that the interaction structure is highly fixed.

In the agent society model, two levels of specification for interactions are provided. The OM provides a script for interaction scenes according to the organizational aims and requirements and the IM, realized in the form of contracts, provides the interaction scenes such as agreed upon by the agents. It is the responsibility of the agents to check that the IM is a refinement of the OM (this can be supported by formal verification). It is also the responsibility of the agents to ensure that their actual behavior is in accordance with the contracts (the IM – they can use a monitoring agent or notary services provided by the society for that). However, it is the responsibility of the society to check that the agents fulfill these responsibilities (this is realized by means of social contracts, cf. section 5).

4. Organizational model

The organizational model specifies the structure of an agent society in terms of externally observed behavior and components. An organizational model is a tuple $OM = (R, CF, I, N)$, where R is the set of role descriptions, CF is the communicative framework, I is the interaction structure and N is the set of society norms or rules. The elements of OM can be referred to by:

$$\begin{aligned} roles(OM) &= R \\ comm-framework(OM) &= CF \\ interactions(OM) &= I \\ norms(OM) &= N \end{aligned}$$

Society goals are not part of the organization model but form the background to the society definition and are represented by the goals of the roles. That is, the goals of a society are specified in terms of roles that correspond to the different stakeholders in the domain. Moreover, the organizational model can be seen as split into two parts: facilitation and operation. The facilitation layer provides the backbone of the society and consists of institutional agent roles, which are designed to enforce the social behavior of agents in the society and assure the global activity of the society. The operational layer models the overall objectives and intended action of the society and consists of domain related roles. The operational layer is always domain and application dependent whereas the facilitation layer depends on the cooperation characteristics of the environment and can be reused across domains (this is one reason why the recognition of this layer is useful during a design process). Typically, facilitation issues are suitable for outsourcing to an external party ([8]). The organizational model does not need to make any formal difference between facilitating agents and operational agents (we can use the same models and logics), but usually, the facilitating agents will be more controlled by the institution.

In the following, we will not describe all the components (see [9] for a more elaborate treatment), but focus on the Interaction Model using the example of conference organization. The global goal of the conference society is to realize a conference. This goal can be split into a facilitation component that aims at the organization of the conference, and an operational part where the conference program is generated and consumed. *Operational* stakeholders in this society, specified as roles in the operational component of the model, are authors, PC members and participants. The goal of author is to get papers accepted, the PC member aims at assuring the quality of the program and the participant hopes for a high quality to consume. *Facilitation* activities can be described in terms of an organizer role that administrates the conference and a chairperson role responsible to regulate conference sessions.

4.1. Interaction structures

The possible actions of a role determine interactions with other agents. These interactions are articulated into (institutionalized) interaction patterns between agents that establish the possible dialogues and interchanges between agents. Such patterns

are designed in such a way that the society aims are achieved. Examples are the negotiation dialogue needed to close a business transaction, or the workflow supporting a delivery process. We define an **interaction structure** as a pair $I = (S, T)$ where S is a set of scenes (represented by interaction scripts) and relation $T: S \times S$ gives the scene transformation matrix. For $s, s' \in S$, $(s, s') \in T$, means that s' can be reached from s .

The scene transformation induces a partial ordering T^* on the set of scenes. Scene transformation T has consequences for *roles(OM)* in the sense that actors will follow different paths according to their roles. Furthermore, the performance of an actor in a given scene will have consequences to the roles that actors will play in other states. For example, the interaction structure of the conference society is depicted in Fig. 2. Note that the graphical representation is very similar to (and in fact inspired by) the work of Sierra [12]. However, an important difference is that protocols are fixed in that work, while our framework is based on the principles of refinement and collaboration autonomy.

An **interaction script** describes a scenario of activity, that is, how roles interact and evolve in the context of a scene. Interaction scripts serve as a blueprint for the actual interactions between actors. Landmarks are logical expressions that describe the characteristics (for instance, goals and action plans) of the scene. The level of specification of landmarks determines the degree of freedom the actors have about their performance. Norms are deontic expressions that regulate the activity of actors.

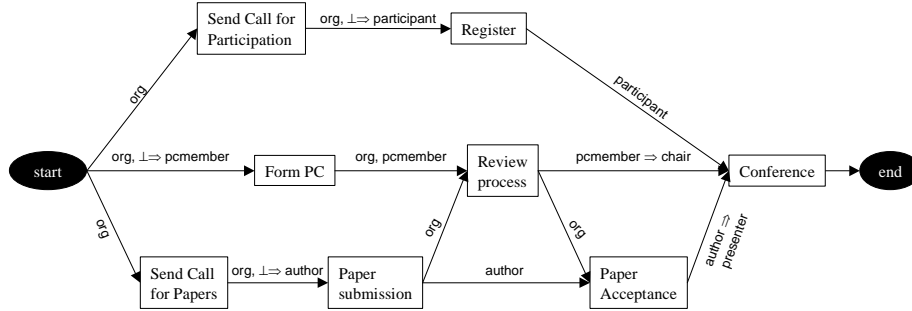


Fig. 2. Interaction Structure for 'Conference' Society

An interaction script, $IS \in S$, is defined by $IS = (P, CF', L)$ where:

- P is the set of participating roles, $P \subseteq roles(OM)$,
- CF' is the communicative framework used in the scene, possibly a specialization of *comm-framework* (OM),
- L are the landmarks describing the interaction (propositions that hold for the scene, including the norms)

Review process	
Roles:	Organizer (1) PCmember (N)
Landmarks:	$\forall p \in \text{Papers}, \quad \text{get_reviews}(p, \text{review1}, \text{review2}).$ $\forall p \in \text{Papers}, \quad \text{paper_decision}(p, \text{decision}, \text{review1}, \text{review2}).$ $\text{DONE}(\text{Organizer}, \text{get_reviews}(P, \text{Rev1}, \text{Rev2})) \equiv$ $\text{DONE}(\text{Organizer}, \text{assign_paper}(P, \text{PCmemberA}, \text{DeadlineR}) \text{ BEFORE } \text{DeadlineA})$ $\text{AND } \text{DONE}(\text{Organizer}, \text{assign_paper}(P, \text{PCmemberB}, \text{DeadlineR}), \text{ BEFORE } \text{DeadlineA})$ $\text{AND } \text{DeadlineA} \text{ BEFORE } \text{DeadlineR}$ $\text{AND } \text{DONE}(\text{PCmemberA}, \text{receive_review}(\text{Org}, P, \text{Rev1}) \text{ BEFORE } \text{DeadlineR})$ $\text{AND } \text{DONE}(\text{PCmemberB}, \text{receive_review}(\text{Org}, P, \text{Rev2}) \text{ BEFORE } \text{DeadlineR})$
Norms:	$\text{PERMITTED}(\text{Organizer}, \text{assign_paper}(P, \text{Reviewer}, \text{Review_deadline})) .$ $\text{OBLIGED}(\text{Reviewer}, \text{receive_review}(\text{Org}, P, \text{Rev}) \text{ BEFORE } \text{Review_deadline}).$ $\text{OBLIGED}(\text{Organizer}, \text{paper_decision}(P, \text{Decision}, \text{Rev1}, \text{Rev2}) \text{ BEFORE } \text{Decision_deadline}).$

Fig. 3. Interaction scene ‘Review Process’

For example, in the conference society the interaction script for scene ‘Review process’ is partially described in figure 3. The arity (1 or N) of roles indicates the allowed number of different actors that can play the same role in the scene. Transformations indicates the possible consequences for actors of a role in future interactions (where \perp indicates the disappearance of a role). As shown, landmarks can be more specific than just goals; in this example organizational design makes explicit how reviews are supposed to be obtained. Actors are in this case still free to decide about the assignment of papers (for example, letting reviewers indicate their preferences), about how to deal with missing reviews or about the acceptance or not of papers (for instance, do reviewers get to vote about papers or does the organizer decide alone). Such decisions are to be fixed as contracts in the scene play agreements in the Interaction Model.

5. Matching OM and agent autonomy: contracts

The Organizational Model as described above is not enough to describe the activity of an agent society. The OM provides the specification of the society behavior as desired by the society designers. However, the actual behavior of the society emerges from the goal-pursuing behavior from the individual agents within the constraints set by the Organizational Model. From an organizational point of view this creates a need to check conformance of the actual behavior to the desired behavior. In the **Social Model** (SM) the social contracts that hold at a given moment in a society are described. These contracts make explicit the commitments between actors and the society, that is, specify all the aspects that govern and determine the interaction of a member with the rest of the society. Depending on the capabilities of the agents and on the organizational constraints, such contracts can be negotiated by the agents

seeking admission. The social model contains the instantiation of roles to actual agents (this combination of agent and role is called an **actor**).

Furthermore, actual interactions between actors must also be made explicit. We call this the **Interaction Model (IM)**. In the agent model interaction scenes are instantiated or created as a result of previous scenes for individual agents. Interaction scenes are described for generic agent roles and can be instantiated for enacting agents. A scene instantiation is described in a **contract** between the interacting agents. In the following, we collect the most important requirements for a contract language – a formal logic that meets these requirements is described in [10].

5.1. Contract requirements

The interaction structure as specified in the OM, gives a partial ordering of interaction scenes and the consequences of transitions from one to the other scene to the roles (creation, destruction, etc). That is, the interaction structure indicates the possibilities of an agent at each stage and the consequences of its choices. Explicit representation of commitments is necessary in order to know how to proceed if the norm is violated but first of all to inform the agents about the behavior they can expect from the other agents. In this way, coordination can become possible. A formal language for contract specification is needed in order to be able to evaluate and verify agent interaction.

Contract rules usually do not express an immediate obligation, but the obligation to achieve a certain result before something else has happened. The fact that usually contract obligations refer to a deadline indicates that temporal aspects are of paramount importance for the specification of contracts. Furthermore, because agents are autonomous entities that can decide on the course of action they will take, different futures must be considered at any moment. Deontic logic provides possibilities to reason about violability of norms, that is, about how to proceed when norms are violated. In systems where agents are assumed to be autonomous and intelligent, norm violation must be considered (agents can deliberately choose to violate a norm if they think that their goals are better achieved in this way).

Furthermore, verification of the behavior of an open society, where the design of participating agents cannot be controllable, must be based on the externally observable effects of agent actions. That is, from the society perspective, different actions that bring about the same state of affairs in the world cannot be distinguished.

6. Society types

The way organizational coordination models achieve coordination is determinant to the motivation of coordination in agent societies. Table 2 derived from [11] gives an overview of the characteristics of agent societies with different coordination models.

	Market	Network	Hierarchy
Type of society	Open	Trust	Closed
Agent 'values'	Self interest	Mutual interest/ Collaboration	Dependency
Facilitation roles	Matchmaking Banking	Gatekeeping Registry Matchmaking	Interface Control

Table 1. Coordination in agent societies

In **markets**, agents are self-interested (determine and follow their own goals) and value their freedom of association and own judgement above security and trust issues. Openness is thus per definition a feature of markets. Facilitation is, in the most extreme case, limited to identification and matchmaking activities. Interaction in markets occurs through communication and negotiation.

Network organizations are built around general patterns of interaction or contracts. Relationships are dependent on clear communication patterns and social norms. Agents in a network society are still self-interested but are willing to trade some of their freedom to obtain secure relations and trust. Therefore, agents need to enter a social contract with the network society in which they commit themselves to act within and according to the norms and rules of the society. The society is responsible to make its rules and norms known to potential members. Coordination is achieved by mutual interest, possibly using trusted third parties, and according to well-defined rules and sanctions.

Finally, in a **hierarchy** interaction lines are well defined and the facilitation level assumes the function of global control of the society and coordination of interaction with the outside world. In a hierarchy, agents are cooperative, not motivated by self-interest and all contribute to a common global goal. Coordination is achieved through command and control lines. Agents are cooperative and not motivated by self-interest. Such agents are said to be benevolent, that is, agents are assumed to always attempt to do what is requested from them.

Let us consider now the way agent interaction is defined in each of these society types using the agent society model. In a **market**, the OM is relatively simple. It contains the roles of Buyer and Seller, as well as some facilitating roles such as Market Owner, Matchmaker and Bank. The Interaction Scenes fix the interaction based on standard protocols, such as auctioning protocols. In a typical market, the agents cannot refine these protocols, in other words, the contracts at IM level are identical to the scenes (of course, the parameters are different for each instance, but the contract structure is fixed). However, a market may also allow for refinement (in which case the society will start to resemble the network type). The Social Contract (SM level) specifies for an agent if it can play the role of Buyer, Seller, or both. The social contract also specifies what can happen in the case of misbehavior. However, the parties themselves are responsible for monitoring the actual performance. If the counterparty does not behave according to the scene/contract, and the issue is not solved bilaterally, one party can raise a claim with the Market Owner. Depending on

the Social Contract, the Market Owner may then decide to sanction the non-behaving party.

In a typical **hierarchy**, the OM has a rich differentiation. It contains various functional role descriptions and the Interaction Scene binds these roles together into process models. One or more Controller roles must be defined that have the task of monitoring the performance of a functional role, comparing it with the organizational goals or some targets derived from them, and possibly initiating some remedial action. The extreme case is that the Controller role is performed by one agent and the OM does not allow for refinement. In such a situation, the advantages of an agent architecture such as local autonomy and flexibility become problematic. The other extreme is that the Controller role is distributed over many agents, and that refinement of contracts is possible. In this situation, the hierarchy starts to resemble a network. Normally, the structure of the hierarchy will be somewhere between these extremes, depending on the level of standardization that is desirable in the given environment. The assumption that agents are cooperative, that is, follow-up requests, is preferably not taken for granted, but put explicitly in the Social Contract. Agents are supposed to follow-up all requests authorized directly or indirectly by the Social Contract. On this point, the hierarchy is fundamentally different from both the market and the network.

The **network** type is the type that maximizes collaboration autonomy. The Organizational Model can be simple, including not more than a role Participant and some facilitating roles such as Notary and Monitor, or it can be more elaborate, such as the Conference Organization example described above. The function of the Notary is to verify the correctness of contracts (whether they conform to the scene descriptions), and to register them. The role of the Monitor is to validate the contract, that is, to monitor actual performance. This role can also be assigned to the parties themselves. If actual performance deviates from the contract, typically the contract will indicate remedial action. If there is breach of contract, the parties can report this to the Network Owner or to some Reputation role. This needs to be defined in the Social Contract to which each agent in the network has to commit.

It can be observed in this discussion is that both markets and hierarchies, however different in nature, can make use of the same two basic coordination mechanisms. These basic mechanisms are *standardization* and *collaboration autonomy*. Collaboration autonomy corresponds with what Mintzberg [15] has called *mutual adjustment*, and with respect to standardization he has distinguished several types of standardization (of work processes, outputs, skills, or norms). The sixth organizational coordination mechanism distinguished by Mintzberg is direct supervision, which we would like to generalize to *authorization* and which is the basic structure of hierarchies.

7. Conclusion

In this paper, we have described an approach to the specification of interactions that respects both the internal autonomy of agents as well as the collaboration autonomy. In order to profit from the flexibility of autonomous agents while at the same time ensuring the robustness of the complete system, we have argued for the

need to define an organizational structure for the multi-agent system (similar to real-life organizations). This organizational structure is based on the goals of the overall system and indicates the behavioral boundaries for the individual agents within the system. We have described an Organizational Model that allows the designer of the Agent Society to specify some very basic structures (landmarks) to be respected by all agents within the society. These basic structures include role definitions as well as interaction scenarios.

Because agents come in with their own goals and capabilities, we need to specify how the agent fits into a certain role it takes in the multi-agent organization. To this end, a social component (the Social Model) has been added to the agent architecture. This component takes care that the goals of the agent fit within those for the role it plays and that the agent fulfills all norms belonging to the particular role. We have shown how the Social Model and the Organizational Model differ for the various society types (hierarchies, networks, markets). One conclusion that can be drawn from that discussion is that refinement is most typical in a network type of society, but it can be combined with the other types as well. The extreme market and hierarchy types do not support collaboration autonomy, but they can introduce it to some degree, thereby evolving into a network.

The combination of “refinement” and “agreement” in our proposal is rather unique. Refinement in itself is not new; for example, it is also used under the name of “levelling” in AUML [16]. However, in our proposal refinement is more than an abstraction device: it is a matter of the proper balancing of responsibilities.

For the society, it is important that these contracts adhere to the landmarks given by the Organization Model. Therefore, the specification languages must be formalized syntactically and semantically so that formal verification is possible. We have developed a branching time deontic logic called LCR [10] that can provide the logical semantics of both the Organization Model and the contracts.

Using the models presented in this paper, and the LCR logic, it becomes possible to give a precise and implementable specification of agent societies. One important topic for future research is to demonstrate precisely that, that is, the development of architectures and design patterns for various types of agent societies.

8. References

- [1] Andrade, L., Fiadeiro, L., Gouveia, J., Koutsoukos, G. Lopes, A. and Wermelinger, M.: Coordination Patterns for Component-Based Systems. Musicante M., Haeusler H. (Eds.): Proc. of Brazilian Symposium on Programming Languages, UFPR 2001, (2001), B: 29-39.
- [2] Artikis, A., Pitt, J.: A Formal Model of Open Agent Societies. Proc. Autonomous Agents 2001, (2001) 192-193
- [3] Bons, R., Lee, R., Wagenaar, R., Wrigley, C.: Modeling Inter-organizational Trade Procedures Using Documentary Petri-nets. Proc. Hawaii International Conference (1995).
- [4] Castelfranchi, C.: Engineering Social Order, Omicini, A., Tolksdorf, R., Zambonelli, F., (Eds.) Engineering Societies in the Agents World, First International Workshop, ESAW 2000, Berlin, Germany, LNAI 1972, Springer-Verlag (2000), 1 – 19
- [5] Dastani, M., Dignum, V., Dignum, F.: Organizations and Normative Agents. Submitted to the First Eurasian Conference on ICT, Teheran, Iran, October, (2002).

- [6] Dellarocas, C.: Contractual Agent Societies: Negotiated shared context and social control in open multi-agent systems. Proc. WS on Norms and Institutions in Multi-Agent Systems, Autonomous Agents-2000, Barcelona (2000).
- [7] Dignum, F.: Autonomous Agents with Norms. In *AI and Law*, (7), (1999) 69 – 79.
- [8] Dignum, V., Dignum, F.: Modeling agent societies: co-ordination frameworks and institutions. In: P. Brazdil, A. Jorge (Eds.): *Progress in Artificial Intelligence*. LNAI 2258, Springer-Verlag, 2001
- [9] Dignum, V., Meyer, J.-J., Weigand, H., Dignum, F.: An Organizational-oriented Model for Agent Societies. In: Proc. Int. Workshop on Regulated Agent-Based Social Systems: Theories and Applications (RASTA'02), at AAMAS, Bologna, Italy, July, (2002).
- [10] Dignum, V., Meyer, J.-J., Dignum, F., Weigand, H.: Formal Specification of Interaction in Agent Societies. 2nd Goddard Workshop on Formal Approaches to Agent-Based Systems (FAABS), Maryland, Oct, (2002)
- [11] Dignum, V., Weigand, H., Xu L.: Agent Societies: Towards framework-based design. Wooldridge, M., Weiss, G., Ciancarini P. (Eds.): *Agent-Oriented Software Engineering II*, LNCS 2222, Springer-Verlag, (2002) 33-49.
- [12] Esteva, M., Padget, J., Sierra, C.: Formalizing a language for Institutions and Norms. Proceedings of the 8th International Workshop on Agent Theories, Architectures and Languages, ATAL-2001, Seattle, (2001)
- [13] Hindriks, K., De Boer, F., Van der Hoek, W., Meyer, J.-J.: Agent Programming in 3APL. *Autonomous Agents and Multi-Agent Systems*, 2(4):357–401, (1999).
- [14] Kumar, S., Huber, M., Cohen, P., McGee, D.: Towards a Formalism for Conversation Protocols Using Joint Intention Theory. Chaib-draa, B., Dignum, F. (Eds.): *Computational Intelligence Journal*, Special Issue: Agent Communication Language, (2002). *To appear*.
- [15] Mintzberg, H. *Mintzberg on Management*. The Free Press, New York (1989).
- [16] Odell, J., Van Dyke Paranak, H, and Bauer, B., Extending UML for Agents. Proc. AOIS Workshop, Austin, July (2000), pp.3-17.
- [17] Verharen, E.: A Language/Action Perspective on the Design of Cooperative Information Agents. Ph.D. Thesis, Tilburg University, (1997).