Emergent Societies of Information Agents

Paul Davidsson

Department of Computer Science, University of Karlskrona/Ronneby
Soft Center, 372 25 Ronneby, Sweden
Paul.Davidsson@ipd.hk-r.se

Abstract. In the near future, billions of entities will be connected to each other through the Internet. The current trend is that an increasingly number of entities, from smart personal devices to legacy databases, are controlled by software agents. Such agents often also posses a large amount of information about both the entity and its owner. Thus, a likely scenario is that the Internet will be populated by millions of information agents, all potentially able to communicate with each other. Unfortunately, we cannot assume that these agents are benevolent and are willing to cooperate in an altruistic fashion. As the amount of money transferred via the Internet is rapidly increasing caused by the breakthrough of e-commerce, we should actually expect a similar increase in the number of malicious agents. Another aspect that contributes to the complexity of agent interaction on the Internet is a desired openness, making it difficult to engineer agent societies in a top-down manner. Rather, we will here investigate the prerequisites necessary to form stable and trustworthy societies of information agents, and discuss some open problems and methodologies for studying them. The general conclusion is that more research is needed that takes into account the presence of malicious agents.

1 Introduction

In the near future, billions of entities will be connected to each other through a global communication channel, i.e., the Internet. Although some of these entities will be personal computers, the major part will be different kinds of smart devices, such as, mobile phones, digital personal assistants, and even refrigerators. The current trend is that more and more of these devices are controlled by software agents. In addition to the control of the device, these agents will in many cases also have access to a large amount of information about not only the device, but also about its user(s). Another trend is the use of "wrapper" and "transducer" agents [11] in order to increase the availability and usefulness of legacy information systems. If we extrapolate from these trends, a likely scenario is that the Internet soon will be populated by millions of information agents, all potentially able to communicate with each other through the Internet. The openness of the Internet brings about some desirable properties, e.g., supporting interoperability and cooperation. However, as we will discuss later, this openness also has some drawbacks.

A collection of agents interacting with each other can be seen as an agent society. The view of *open* agent societies, where, in principle, anyone with Internet access may contribute one or more agents without any particular restrictions, should be contrasted to that of *closed* agent societies, where an agent-based approach is adopted by a team of software developers in order to implement a complex software system (cf. distributed problem solving). In a closed multi-agent systems (MAS), it is often possible to precisely engineer the society, e.g., specify with which agents each agent interacts, and why. From this perspective, agent-based computing may be seen purely as a software engineering paradigm (cf. [31]). In between these types of agent societies, we have *semi-open* societies, where there are *institutions* to which an agent may explicitly register its interest to enter the society. The institution then either accepts or rejects this request. An example of such institution is the *portal* concept as used in SOLACE [15, 16].

Ideally, all agents belonging to a society should always cooperate in order to find globally optimal, or at least acceptable, solutions to the problems related to the fulfillment of the goals of the individual agents. However, agents often have conflicting goals that may result in competitive behavior. While being self-interested, agents are typically assumed to be sincere when interacting with other agents. Although such benevolence may be assumed in closed societies, this is an unrealistic assumption for most open agent societies. We must accept that in open agent societies there may be malicious agents trying to exploit the norm- and law-abiding agents, e.g., by stealing secret and/or personal information. As the amount of money transferred via the Internet is rapidly increasing caused by the breakthrough of e-commerce, we should expect a similar increase in the number of such malicious agents. Consequently, this must be taken into account when developing techniques, standards, software etc. for open (and semi-open) agent societies.

Unfortunately, the presence of self-interested and even malicious agents significantly increases the complexity of achieving stable societies. Although these aspects have been studied to some extent in particular situations, e.g., computational auctions, not very much work has been carried out at the general society level. However, the presence of malicious agents should not be confused with the limited and controlled competitiveness that is sometimes used in closed agent societies. In this case, the agents are designed to be competitive in order to achieve desired behavior at the global system level. Another aspect of open agent societies is how to achieve robust behavior in the face of imperfect knowledge about the environment (see e.g., [17]). This has, on the other hand, been quite well studied. However, it should be contrasted to the scenario we will concentrate on here, where agents deliberately may provide other agents with false information in order to gain some advantage.

We will here discuss the impact of the possible presence of malicious agents and what ingredients are necessary to form stable and trustworthy societies of agents. In the next section we specify what requirements that must be met in order for a society to emerge and in the following section take a closer look at societies of information agents. We then discuss what the research carried out in the social sciences may contribute to the study of artificial agent societies. Finally, we briefly describe a project that adopts an approach that to a large extent shares the view expressed in this paper on how to study agent societies.

2 Prerequisites for the Emergence of Information Agent Societies

What do we mean by a society? A typical definition of a human society is that it is a structured group of persons associated together for some kind of purpose, or/and residing in a specific locality. If the latter is the dominant factor, the group is often called a *community*. In addition, a society (or community) often has a set common rules and norms that the members are supposed to adhere to. In ecological contexts, similar definitions are used. For instance, an ecological society has been defined as "a closely integrated group of social organisms of the same species held together by mutual dependence and exhibiting division of labor", and an ecological community as "an assemblage of plant and animal populations occupying a given area."[1]

We will here use these concepts with respect to groups of "social" artifacts in analogous ways. That is, a collection of agents interacting with each other in some way or another, possibly in accordance with common norms and rules (cf. social laws [28]), are here called an agent society. The role of the society is to allow agents to coexist in a shared environment and pursue their respective goals in the presence of other agents. Note that both cooperation and competition between the members of the society are possible.

As pointed out by Wooldridge and Jennings [30], it is a common misconception that agent based systems can be developed simply by throwing together a number of agents in a melting pot where there is no structure and all agents are peers communicating with each other. A similar argument can, of course, be made with respect to the agents on the Internet. An agent will not interact with every other agent just because it may have a possibility to do so; there has to be a reason for the interaction, e.g., that it will bring the agent closer to achieving one or more of its goals. Thus, we will have some kind of structure within the Internet based on societies of agents, defined by the interaction that takes place between the agents. In accordance to what has been said earlier with respect to semi-open societies, additional constraints regarding what agents are allowed to belong to a society may be imposed by institutions. Such constraints can be related to the expected probability that the agent will comply with the norms and rules of the society, e.g., based on the agent's (or its owner's) credibility.

It would also be possible to make a distinction between agent societies and agent communities, where communities are characterized by closeness in proximity, e.g., residing on the same computer or local network. However, in what follows we will not make this distinction.

Another aspect that contributes to the complexity of agent interaction on the Internet is an often desired openness, i.e., given that an agent has an acceptable reason to join a society, it should be given the opportunity to do so. This makes it difficult to engineer agent societies in a top-down manner. Rather, we should provide an infrastructure that makes bottom-up "emergence" of stable societies possible. The most basic requirement that must be fulfilled for a society to emerge is that the entities of the society (or at least a subset) are able to interact with each other. For this we need at least: (i) a communication medium, (ii) a way of finding other entities to interact

[1] These definitions are taken from Webster's Encyclopedic Unabridged Dictionary of the English Language, 1989 edition.

with, (iii) a reason for the entities to interact (iv) and a communication language, including syntax, semantics, and pragmatics.

(i) In the case of information agent societies, the Internet (and the IP protocol) provides the necessary communication backbone. Fortunately, as this already exists and is extremely well tested, we can make the assumption that it is there and is working.

(ii) There must be mechanisms for an agent to enter (and leave) a society. These could be either *manual*, i.e., the owner at run-time, or programmer at "design time", tells the agent with other agents to interact with, or *automatic*, e.g., by making use of supporting mediator agents, such as, "yellow pages" or portal agents. In the latter case, the supporting agents also need to have knowledge about the competence of the agents of the society.

(iii) The interaction with other agents should in some way or another help the individual entity to fulfil one or more of its goals. As this is completely dependent of the goals of the particular agent there is difficult to discuss this in general terms. However, it can be noted that this is the reason why the billions of agents that may be populating the Internet will not all communicate with each other; there simply has to be a reason for the interaction.

(iv) The linguistic interaction between agents within a society can be studied at different levels, e.g.:

- *Syntactic level*: What language (syntax) is used for the communication?
- *Semantic level*: What ontologies are used for the communication?
- *Coordination/conversational level*: What is the structure of the communication?
- *Social level*: What norms and rules are used in the communication? With which agents does an agent want to, or need to, communicate?

In order for a stable society to emerge, a common language, a common ontology, and common norms and rules are needed. Coordination and conversation models may not be strictly needed, but are from a practical perspective very useful since they make the interaction between agents more structured and efficient. However, all of these levels have recently been the topic for much attention in agent research. In fact, they are now almost research fields in their own right, arranging separate conferences or workshops, e.g., ACL (workshop on Agent Communication Languages) [7], FOIS (conference on Formal Ontology in Information Systems) [14], SICP (workshop on Specifying and Implementing Conversation Policies) [13] and the more general conference on Coordination Models and Languages [2], and the workshop on Norms and Institutions in Multi-Agent Systems [6].

However, most of this research assumes that agents are benevolent and, as pointed out earlier, this assumption is not realistic in most open agent societies. For instance, consider the two most commonly used and well-defined agent communication languages, KQML [8] and FIPA ACL [9], which both include a primitive performative for communicating information by a declarative speech act (*tell* and *inform* respectively). Attempts have been made to define the semantics of this speech act in terms of the intentional stance, i.e. mental attitudes in terms of beliefs, desires and intentions. In both cases, the meaning is specified using a semantic condition, sometimes called the *sincerity condition*, saying that an agent actually believes what it communicates [21, 9]. As pointed out by Pitt [24], this condition is not valid in many open (and semi-open) agent societies. Similar objections can be made for the other levels of linguistic interaction. In fact, we argue that the consequences of the non-benevolence assumption and the openness of the Internet are understudied topics, and need to be taken into account in nearly all aspects of agent research. The awareness of these matters are probably highest in the research on the social level, where concepts such as *trust* and *deception* are given more and more attention [1].

3 Information Agent Societies

The software agents (and other software entities, but we will here concentrate on the agents) populating the Internet forms a "Universal Information Ecosystem" (UIE) that can be seen as a collection of societies. These societies may be either open, closed, or semi-open. An agent may belong to several societies, but agents belonging to different societies typically do not interact with each other. Moreover, each society will have its own (possibly empty) set of norms and rules. Within a society there may be a number of coalitions, each consisting of a number of individual agents. Whereas a society is neutral with respect to cooperation and competition, coalitions are formed with the intention of cooperation. This view of the UIE is illustrated in Fig. 1. Note that also a closed society may include coalitions.

Each agent on the Internet is "owned" by a real person (or organization) on whose behalf it acts and who caused its creation. A person may have several agents, even in the same society. Thus, we can see interaction between agents as social interaction between humans mediated by agents. It is important to note that such interaction may have both social and legal consequences, a fact that often is ignored in agent research. It is the owner of the agent that should be held responsible for any illegal or immoral actions performed by the agent, e.g., breaking a contract by not fulfilling a commitment, or spreading secret information. One consequence of this ownership relation is that an agent may be fully autonomous with respect to deciding what plans and actions to adopt in order to achieve a particular goal, but only limitedly autonomous with respect to deciding what goals to adopt and what rules and norms to obey, which then are decided by its owner (cf. the distinctions between plan, goal, and norm autonomous agents [29]).

It is sometimes easier for an agent to achieve its goals if it cooperates with some of the other agents in the society. In an open (or semi-open) agent society, it is typically not possible to determine with which agents to cooperate when the agent is created. That is, the agent itself must often decide at run time whether it is fruitful or not to cooperate with a particular agent, or set of agents. Thus, *coalition formation* is a very important activity in agent societies. Unfortunately, most research in this area assumes societies that are closed and where agents have unlimited computational resources. Some attempts to ease up these restrictions have been made [27, 20] and a promising general framework has been suggested by Johansson [19]. It should be noted that there are some similarities between society and coalition formation, e.g., similar reasoning is required by the agents when deciding whether to join a coalition or a society, i.e., weighing the pros and cons of being a member of the coalition / society. Compared to

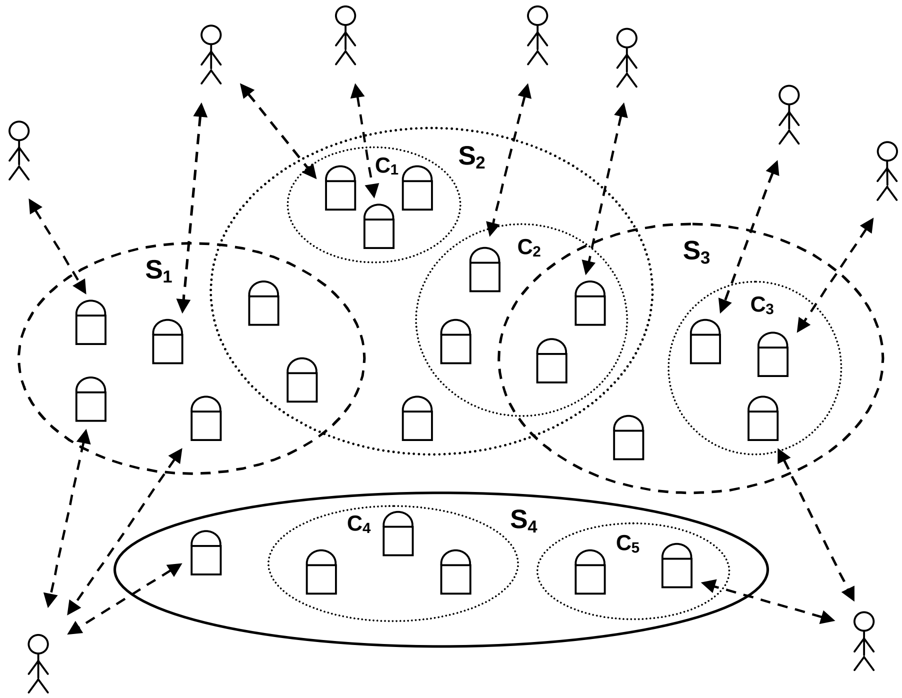

Fig. 1. A schematic illustration of the structure of the UIE. In this example there are four societies; S_1 and S_3 are semi-open (marked by dashed ellipses), S_2 is open (dotted ellipsis), and S_4 is closed. Within these societies there are five coalitions in total ($C_{1\text{-}5}$, thin dotted ellipses). Also, some of the persons (or organizations) that "owns" the agents are illustrated.

the society, a coalition within the society typically has a set of additional rules and/or norms that the agents of the coalition should adhere to.

If we study current information agent societies, we notice that they may include a number of different types of agents, e.g.:

– *Interface agents*: Provides an interface between human users and the Internet.
– *Information agents*: such as, information retrieval agents and database agents.
– *Middle agents*: Different kinds of mediating agents that glues the society together, e.g., portals, brokers, and match makers.
– *"Smart equipment agents"*: Provides an interface between hardware entities and the Internet.

However, open and semi-open societies may call for new types of agents. For instance, it would be useful to have particular agents that cope with malicious agents. One approach would be to introduce "police agents" that monitor the society looking for agents that break the norms and rules of the society. If such an agent is detected, it takes appropriate measures, e.g., excluding the agent from the society or punishes its owner in some way. In semi-open agent societies this may be taken cared of by institutions accepted by the members of the society. Similarly, institutions in the form of "information banks" may be introduced to provide safe storage of personal and other secret information. Another possible type of agent would be "meta-mediators", which have knowledge about different societies, and are able to guide (new) agents to the right society. Similarly, coalition brokers may be introduced to facilitate coalition formation.

4 Human Societies and Social Theories

As should be apparent from the discussions above, societal issues assuming the presence of self-interested and possibly malicious agents are very relevant to research on both open and semi-open information agent societies. However, instead of developing theories and models from scratch, we may learn from the several hundreds of years of research that has been carried out in the social sciences on these topics. We will here use a very broad definition of social sciences that includes, e.g., economics, psychology, sociology, anthropology, and philosophy, as well as ecology and even some areas of biology.

One idea would be to transfer social theories about human (and animal) societies to the artificial societies of the UIE. However, there are differences between natural and artificial societies, e.g., the consequences of terminating an entity of the society, the importance of physical distance between entities, and the reasoning abilities and predictability of the entities themselves. Thus, we must ask ourselves what theories are relevant, e.g., sometimes the theory needs to be modified to better suit artificial societies. Also, if the social theory is only loosely described, e.g., in a natural language, it may need to be formalized.

We will by social theory mean every theory, informal or formal, that describes situations in which two or more humans or animals cooperate and/or compete. Now, if the designer of an agent (or an institution) predicts that the agent (or institution) will find itself in a specific situation where it needs to interact with other agents (or institutions) in some organized way, we may try to identify a social situation analogous to this situation and then look for social theories and models describing the situation. If such a theory is found, we try to transfer to it agent societies, if possible. A more detailed description of these ideas can be found in [5].

This idea has already been applied successfully in several instances. Take for instance the idea of *market-oriented programming* [3], which exploits the concept of economical markets and uses them to build computational economies to solve particular problems, such as, distributed resource allocation. This approach is inspired in part by economists' metaphors of market systems "computing" the activities of the agents involved, but also by the agent researchers' view of autonomous software entities forming a society of agents. In market-oriented programming these metaphors are taken literally, and are directly used to implement the distributed computation as a market price system. It is sometimes argued that economic models of agency reduce agents to be selfish entities and are therefore of limited value in practical application

[18]. However, it is exactly this aspect that makes them interesting in the open society context we are discussing here. Other examples of social theory transformations are the *game theory* [22], which has been applied to agent societies, e.g., by Rosenschein and Zlotkin [26], and the *genetic theory of natural selection* [10], which has resulted in *genetic algorithms* [12].

When applying social theories and models to artificial agent societies there are some things that one should bear in mind. For instance, great care should be taken when deciding whether to make a complete or only a partial transformation of the social theory. Some parts of the social theory may be irrelevant in the computational context, e.g., the exclusion of sex in genetic algorithms. Another thing to keep in mind is to make sure that all parts of the theory to transfer really are describable. For example, problems typically occur when the social theory depends heavily on mental states rather than just the observable interaction. Consequently (and quite interestingly), the simpler view of the individual assumed in the social theory, the better it is suited for transformation to the computational domain. If the theory nevertheless is based on non-observable mental states, it is essential to make explicit what assumptions are made about the mental structure and capability of the agents: Is it assumed that the agents are based on a BDI (Belief, Desire, and Intention) [25] architecture? How sophisticated are their world models, e.g., do they include the mental states of other agents, or just observable features? What are the reasoning and planning capabilities of the agents?

As the fielded application of theories and models of agent societies on the Internet typically has both legal and economical consequences, it is important to validate them thoroughly. Because of the inherent complexity, the most (perhaps only) viable way of doing this is by means of simulations. To make such simulations meaningful, the whole system should be simulated, including the human owners of the agents. As been suggested elsewhere [4], Multi Agent Based Simulation (MABS) is particular useful for this kind of complex systems that involves interaction between humans and computer controlled devices.

5 A Case Study

We will now describe a project that investigates the concept of “ethical” behavior in agent societies. The project is called ALFEBIITE (A Logical Framework for Ethical Behaviour between Infohabitants in the Information Trading Economy of the universal information ecosystem)[2] and is a part of the Universal Information Ecosystem (UIE) initiative made by FET (Future and Emerging Technologies) within the IST (Information Society Technologies) fifth framework programme issued by the European Commission. UIE stems from the vision of having billions of infohabitants (e.g., agents) populating the universal information ecosystem.

---

[2] ALFEBIITE is a joint project between Imperial College, London UK, National Research Council, Rome Italy, University of Oslo, Norway, Queen's University, Belfast UK, and the University of Karlskrona/Ronneby, Sweden. It is coordinated by Jeremy Pitt, Imperial College. For more information, see the project home page: www.iis.ee.ic.ac.uk/alfebiite.

In addition to the very specific goals of ALFEBIITE, which we will describe below, the project has some more general objectives: (i) develop a new paradigm for designing, deploying, and managing open distributed intelligent systems, in which social relationships between agents are just as important as interface definitions in providing interoperability, (ii) provide users with greater trust and confidence in agent societies by ensuring that they are not exploited, their rights protected, and their privacy is not violated, and (iii) bridge the gap between human sciences (social psychology, sociology, and philosophy) and more technical sciences (AI, agent-based computing, and information systems). These objectives are very ambitious and we do not believe that ALFEBIITE will resolve them completely, but that we will at least provide guidelines for future research and offer partial solutions. Some of the more concrete goals of the project are to:

- develop a logical framework to characterize norm-governed behavior and social relations based on the results of psychological studies,
- use this framework to specify axiomatic descriptions of communicative acts regarding social commitments and to specify the mechanics of society formation,
- implement a "society simulator" to experiment with and validate these formal models in a number of realistic scenarios, and
- contribute in the development of a software (agent) ownership model and the legal mechanisms necessary for enforcing the consequences of such a model.

To conclude, ALFEBIITE is an ambitious project that tackles one of the core problems of open agent societies, and when doing this, it fully adopts the open (semi-open) society view and takes into account its consequences.

6 Conclusions

A distinction was made between open, semi-open, and closed agent societies. If we want open or semi-open societies, it is often not possible to engineer such societies in a strictly top-down manner as is typically done with closed agent societies. Thus, we need to investigate what are the necessary conditions to form stable and trustworthy societies of information agents in a more bottom-up fashion.

It can be argued that semi-open societies have a greater potential than totally open societies to become stable and trustworthy because of their use of explicit norms and rules and their control of which agents are entering the society. In open societies, on the other hand, norms and rules must emerge in some way and then spread to new members of the society. It is actually not clear how this should be implemented, or whether it is possible at all.[3] However, any regulation that is introduced, e.g., by in-

---

[3] Shoham and Tennenholtz [28] have shown that given the capabilities and goals of all agents in the society, the general problem of automatically deriving useful social laws is NP-complete. Note that this assumes that all agents are cooperative and are willing to give away all information about their goals and capabilities.

stitutions, must be carefully balanced between censorship and openness in order not to result in totalitarian or anarchistic societies respectively.

The general conclusion is that more research is needed that takes into account the presence of non-cooperating and maybe even malicious agents. When doing this, we may be inspired by social science research and borrow concepts and theories. However, great care should be taken, e.g., by explicitly stating what assumptions are made regarding the mental capabilities of the agents.

Finally, this open (and semi-open) society approach to information agent research may in fact result in advances also in the social sciences. Some obvious ways of achieving this are corroboration of existing theories, suggestions to modify existing theories, and the development of new theories.

Acknowledgements

The author wishes to thank Bengt Carlsson, Martin Fredriksson, Rune Gustavsson, Stefan Johansson, Christer Rindebäck (all members of the SoC research group), Johan Kummeneje, and Harko Verhagen for stimulating discussions and useful comments on earlier drafts. He also acknowledges the valuable contribution from the colleagues in the ALFEBIITE project.

References

1. Castelfranchi, C. and Tan, Y.H. (eds.): *Trust and Deception in Virtual Societies*. Kluwer Academic Publishers, 2000.
2. Ciancarini, P. and Wolf, A.L. (eds.): *Coordination Languages and Models*, Springer, 1999.
3. Clearwater, S. (ed.): *Market-Oriented Programming: Some early lessons*, World Scientific, 1996.
4. Davidsson, P.: Multi Agent Based Simulation of "Socio-Technical" Systems. In: *Multi Agent Based Simulation (MABS'2000)*, 2000. (In press.)
5. Davidsson, P. and Ekdahl, B.: Towards Situation-Specific Agent Theories. In: *Intelligent Agent Technologies*, 227-231, World Scientific, 1999.
6. Dellarocas, C. and Conte, R. (eds.): *Agents'2000 workshop on Norms and Institutions in Multiagent Systems*, 2000
7. Dignum, F. (ed.): *IJCAI-99 Workshop on Agent Communication Languages*, 1999.
8. Finin, T., Labrou Y., and Mayfield J.: KQML as an Agent Communication Language. In: J. Bradshaw (ed.), *Software Agents*, MIT Press, 1995.
9. FIPA, FIPA 97 Specification Part 2: Agent Communication Language. FIPA (Foundation for Intelligent Physical Agents), http://drogo.cselt.stet.it/fipa/, 1997.
10. Fisher, R.A.: *The Genetic Theory of Natural Selection*, Dover, 1958.
11. Genesereth, M.R. and Ketchpel, S.P.: Software Agents. In: *Communications of the ACM*, Vol. 37(7): 48–53, 1994.
12. Goldberg, D.E.: *Genetic Algorithms*, Addison Wesley, 1989.
13. Greaves, M. and Bradshaw, J. (eds.): *Agents'99 workshop on Specifying and Implementing Conversation Policies*, 1999.
14. Guarino, N. (ed.): *Formal Ontology in Information Systems*, IOS Press, 1998.

15. Gustavsson, R. and Fredriksson, M.: Coordination and Control of Computational Ecosystems: A Vision of the Future. In: Omicini, A., Klusch, M., Zambonelli, F., and Tolksdorf, R. (eds.): *Coordination of Internet Agents: Models, Technologies, and Applications*. Springer, 2000. (In press.)
16. Gustavsson, R., Fredriksson, M., and Rindebäck, C.: Computational Ecosystems in Home Health Care. In: Dellarocas, C. and Conte, R. (eds.): *Agents'2000 workshop on Norms and Institutions in Multiagent Systems*, 2000. (In press.)
17. Huberman, B.A. and Hogg, T.: The Emergence of Computational Ecologies. In: Nadel, L., Stein, D. (eds.): *Lectures in Complex Systems*, 185–205, Addison-Wesley, 1993.
18. Huhns, M.N. and Stephens L.M.: Multiagent Systems and Societies of Agents. In: Weiss G. (ed.): *Multiagent Systems*, MIT Press, 1999.
19. Johansson, S.J.: Mutual Valuations between Agents and their Coalitions. In: *Intelligent Agent Technology*, World Scientific, 1999.
20. Klusch, M. and Shehory, O.: A Polynomial Kernel-Oriented Coalition Algorithm for Rational Information Agents, *ICMAS'96*, 1996.
21. Labrou, Y. and Finin, T.: Semantics for an Agent Communication Language. In: Singh, M., Rao, A., and Wooldridge M. (eds.): *Intelligent Agents IV*, Springer, 1998.
22. von Neumann, J. and Morgenstern, O.: *Theory of Games and Economic Behavior*, Princeton University Press, 1944.
23. Ossowski, S.: *Co-ordination in Artificial Agent Societies*, Springer, 1999.
24. Pitt, J. and Mamdani, A.: Inter-Agent Communication and the Sincerity Condition. In: Dignum, F. (ed.): *IJCAI'99 Workshop on Agent Communication Languages*, 1999.
25. Rao, A. and Georgeff, M.: BDI Agents — From Theory to Practice. In: *ICMAS'95*, AAAI/MIT press, 1995
26. Rosenschein J. and Zlotkin, G.: *Rules of Encounter*, The MIT Press, 1994.
27. Sandholm, T.W. and Lesser, V.R.: Coalition Formation among Bounded Rational Agents. In*: IJCAI'95*, 1995.
28. Shoham, Y. and Tennenholtz, M.: On Social Laws for Artificial Agent Societies: Off-line design. In: *Artificial Intelligence*, (73) 1-2, 1995.
29. Verhagen, H.J.E.: *Norm Autonomous Agents*, Ph.D. Thesis, Stockholm University / Royal Institute of Technology, Sweden, 2000.
30. Wooldridge, M.J. and Jennings, N.R.: Pitfalls of Agent-Oriented Development. In: *Autonomous Agents'98*, ACM Press, 1998.
31. Wooldridge, M.J., Jennings, N.R., and Kinny D.: A Methodology for Agent-Oriented Analysis and Design. In*: Autonomous Agents'99*, ACM Press, 1999.