Flexible Autonomy in Holonic Agent Systems

Christian Gerber, Jörg H. Siekmann, and Gero Vierke
German Research Center for Artificial Intelligence (DFKI GmbH)
Im Stadtwald, Bau 36
66123 Saarbrücken, Germany
{gerber, siekmann, vierke}@dfki.de

Abstract

For many DAI researchers, autonomy is a vital definitional prerequisite for an agent. We argue that a strict commitment to this view is too rigorous and insufficient in complex and large-scale systems and propose the *holonic agent* paradigm, where the degree of autonomy in an agent society is dynamically adjusted to the problem at hand.

A *holonic agent* is an agent which is composed of subagents, with a self-similar structure. The sub-agents themselves may again be decomposable into holonic agents. A holonic multi-agent system is then used to implement a system called TELETRUCK in a traffic and transportation scenario. Finally, we introduce the notion of abstract resources and provide a categorization of different resource allocation mechanisms for holonic systems.

Introduction

Multi-agent systems (MAS) have become more mature, leading to an increasing reliance on MAS technology in complex domains, such as manufacturing (Parunak 1987), information systems (Klusch 1999), or transportation scheduling and monitoring (Fischer, Müller, & Pischel 1996).

The original motivation for MAS, namely to attack a naturally distributed problem by distributed problem solvers, can be extended as follows: the overall problem can be divided into sub-problems; and the problem solvers for these sub-problems can be agentified. Now, many distributed problems exhibit a recursive structure: an agent system that models such a problem solver may have a similar structure as the agents for the sub-problems, thus they should be structured recursively as well: we call agents consisting of sub-agents with the same inherent structure ***holonic agents***.

According to Koestler (Koestler 1967), a ***holon*** is a natural or artificial structure that is stable and coherent and that consists of several holons as sub-structures. He gives biological examples, for instance a human being consists of organs consisting of cells that can be further decomposed etc. In this paper, we elaborate and discuss different approaches on how to specify and implement holonic multi-agent systems, using existing definitions of agency as a starting point.

The paper is structured as follows: In the next section, we set up our terminology by comparing several agent definitions from sociology and AI, identifying those elements of agents that are relevant for our purpose. Furthermore, we shall sketch Koestler's concept of holonic organization and its current realization in the field of holonic manufacturing systems. In the subsequent section, we use the agent definition from the previous section as a basis for a definition of holonic agents, ensuring that this agent definition remains valid for the holonic case as well. The fourth section is concerned with resource management in holonic agent systems and introduces the notion of abstract resources and a categorization of different resource allocation mechanisms for holonic systems. Furthermore, we demonstrate the application impact of this work by presenting a holonic fleet management system that is implemented according to the concepts presented above. Finally, we conclude and provide an outlook on future work.

Basic Principles

In the following paragraphs, we shall present a collection of well-known definitions for agenthood and use them as a source for our own characterization of holonic agents. Since autonomy is not only a vital prerequisite of agency, but also a critical issue for holonic systems, we point to well-accepted definitions which we shall refer to throughout the paper.

Models of Agency and Autonomy

There is a great variety of agent or actor definitions, ranging from philosophically and sociologically inspired concepts to definitions that focus on implementational aspects such as a software architecture, its efficiency, and tractability.

The sociologist Parsons (Parsons 1969) takes an *actor* to be an agent that has goals. In his definition, an *agent* is an individual that shows behavior. *Behavior* is the capability to change the state of the world. The *world* can be differentiated into the agent itself and its environment. The *environment*, as it is perceived by the actor, defines the situation the actor is in. A *goal* behave in such a way as to achieve a goal. In general, an actor can choose from a set of actions, about which it has specific expectations of how they will change the world. The actor selects a specific action from its options according to its goals, the means at its disposal and its situation. Additionally, actors can use a shared language in order to communicate with one another.

Bratman (Bratman 1987), who analyzed human rational behavior in logical terms, used the mental category *intention* as a third concept in addition to the concepts of beliefs and desires and connected these by postulating certain requirements for an intelligent agent's mental capabilities. Based on these concepts, Cohen and Levesque (Cohen & Levesque 1990) and Rao and Georgeff (Rao & Georgeff 1991) founded a logical theory of beliefs, desires and intentions (BDI Theory) to ascribe these mentalistic notions to artificial agents as well.

Russell and Norvig (Russell & Norvig 1995) define an agent as "anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors."

Wooldridge and Jennings (Wooldridge & Jennings 1995) characterize an agent by the following traits: *autonomy* (with respect to states and actions in the sense of Castelfranchi (Castelfranchi 1995)), *social ability* (agents communicate with other agents by using some communication language), *reactivity* (agents respond to changes in their environment which they can perceive), and *pro-activeness* (agents display a goal-directed behavior in addition to their mere reaction to the environment).

Lange (Lange 1998) provides a more pragmatic definition that is oriented towards industrial demands. He defines an agent as a software object that has the following properties: situatedness, reactivity, autonomy with respect to its actions, and pro-activity. Furthermore, an agent should be temporarily continuous, i.e., it is continuously executing. Optionally, an agent can be communicative, mobile, believable, or able to learn.

Selected Requirements for Agency

All these definitions put different emphasis and requirements on agency, for our purpose, we focus on the following agent characteristics:

- **Autonomy** According to Castelfranchi (Castelfranchi 1995), agent autonomy means that "agents control their actions and internal states to enable them to operate without the direct intervention of humans or others." Russell and Norwig (Russell & Norvig 1995) define an agent's behavior "autonomous to the extent that its behavior is determined by its own experience." We decompose autonomy into three aspects:
  - **State Autonomy** An agent's state is determined only by its previous state and its perception. This requirement does not take the agent's actions into account... tions.
  - **Action Autonomy** Like in Castelfranchi's definition, the action of an agent is determined solely by its current state. This requirement does not take the possible success of the action into account.
  - **Computational Autonomy** The agent is supplied with computation time in a fair manner.
- **Goal-Directed Behavior** An agent has either explicitly represented goals/desires or it exhibits behavior that allows an external observer to ascribe implicit goals or desires to this agent.
- **Action** This requirement subsumes the request of Wooldridge and Jennings for reactivity and pro-activeness.
- **Belief** Agents have implicit or explicit representations of their environment.
- **Bounded Rationality** In analogy to Russell and Wefald (Russell & Wefald 1991) we require a rational agent to behave optimally with respect to its limited resources and its goals.
- **Communication** Similar to Wooldridge and Jennings' requirement of social ability, the agents have a shared communication language.

The Holonic Principle

The term *holon*, a combination of the Greek "holos" (whole) and the suffix "-on" (part), was originally introduced in 1967 by the Hungarian philosopher Arthur Koestler (Koestler 1967) in order to name recursive and *self-similar* structures in biological and sociological entities. According to Koestler a holon is a biological or sociological structure that is stable and coherent and that consists of further holons that function according to similar principles. No natural structure is either "whole" or "part" in an absolute sense, instead every holon is a composition of subordinate parts as well as a part of a larger whole. For example, a human individual is on the one hand a composition of organs consisting of cells that can be further decomposed, and on the other hand he or she is part of a group which in turn is part of the human society.

The organizational structure of a holonic society, or *holarchy*, offers advantages that the monolithic design of most technical artifacts lack: They are robust in the face of external and internal disturbances, they are efficient in their use of resources, and they can adapt to environmental changes.

Koestler's ideas have been picked up and applied in Flexible Manufacturing Systems (see e.g. (Deen 1994)), where the positive features of Koestler's holarchies, namely stability, adaptability, flexibility, and efficiency motivated a similar design for the management of manufacturing processes.

Holons in Holonic Manufacturing Systems (HMS) autonomy and cooperativeness. Here, autonomy has roughly the same semantics as for multi-agent systems; namely, the capability to create and control the execution of plans. Cooperativeness stands for joint planning and coordination of joint plan execution, and, therefore, is subsumed by our agent's attribute of social ability. In a holonic manufacturing system, holons consist of an information processing part and sometimes a physical part which is responsible for transforming, transporting, storing and validating information as well as physical objects. Manufacturing holons can be build recursively out of other holons.

The idea of agents consisting of agents is not new to the AI community either: for example Minsky's *Society of Mind* (Minsky 1986) from 1986 proposed that the human mind is structured as a recursively organized society of actors/agents.

Holonic Agents

We shall now extend a multi-agent system by these ideas, as many application domains exhibit a holonic structure and it is often natural to map these structures directly onto the multi-agent society.

In the following sections we use the terms *holon* and *holonic agent* synonymously. By *super-holon* we denote a composition of subordinate agents, which we call *sub-holons* or *sub-agents*. As these sub-holons may be further decomposable into sub-sub-holons we shall use the term *immediate sub-holons* to distinguish them from its transitive closure.

According to Koestler's framework, arbitrary structures could be viewed as holons, instead we restrict this definition to our notion of an agent, and, furthermore, we require that sub-holons always have the same structure as the super-holon. Now, the essential idea is as follows: A holonic agent has a well-defined software architecture may join several other holonic agents to form a super-holon; and this group of agents now acts *as if it were a single holonic agent* with the same software architecture. How this merging of several separate entities into one entity can be achieved is the subject of the next section. However, let us first recall our basic criteria for agenthood again and extend these to holonic agents:

Autonomy An agent group forming a holon acts as a single entity in its environment. It interacts with the environment as an autonomous agent in the sense of the above presented criteria (state autonomy—action autonomy—computational autonomy).

By joining a holon, agents accept restrictions of their autonomy: they commit themselves to accept the goals from the holon and they accept restrictions of their capabilities to act or to communicate according to the capabilities of the super-holon. Nevertheless, they keep free to leave the holon.

Common Goal-Directed Behavior Sub-agents of a holon still pursue their private goals. In addition, they have to pursue at least one common goal of the holon which may be represented explicitly or implicitly.[2]

Hence, the super-holon's overall goals emerge from the common goals of the sub-agents. We do not require that the holon's goals are also goals of the sub-agents, but they must not contradict any goal of a sub-agent. Consequently, an agent can only be a member of several holons with conflicting goals if the agent is indifferent to these goals, i.e. the overall goals must not contradict the individual agent's goals. This requirement corresponds well to the cooperation feature of an HMS holon and extends the definition for goal-directed behavior.

Increased Group Capabilities A single agent's capabilities to act are extended at the group level to *macro* actions: a super-holon may have actions at its disposal that none of its sub-agents could perform alone.

Belief The requirements for an agent's belief remain unchanged: holons have some representation of their environment, i.e. they hold beliefs about their surroundings. This knowledge might be represented explicitly within the super-holon, or it is implicitly provided by the local knowledge of the individual sub-holons. Inconsistencies between the holon's and some of its members' beliefs may be tolerated since the members are autonomous to withhold their beliefs.

Bounded Rationality A holon has to control its resources in order to exhibit a bounded rational behavior. Resource management of the sub-holons is monitored at the level of the super-holon and guidelines on local resource distribution are propagated to its sub-holons. This is an essential issue and will be discussed further in the subsequent section.

Communication The ability to communicate is part of an agent's autonomy and is solely the right of the holon and not its sub-holons: the right to communicate with some other agent is a resource that corresponds to a communication channel between the two parties. These channels are managed solely by the super-holon. However, we have to distinguish between communication inside the holon and communication between holons: the internal communicational needs for efficient holonic resource management are discussed in more detail in the next section.

[1] For more details, see for instance the HMS web page `http://hms.ifw.uni-hannover.de/public/hms_tech.html`

[2] For example, BDI architectures provide an explicit representation of goals. Implicit goals can be ascribed to any agent that exhibits some kind of pro-activeness.

internal communication load can be high, hence, an efficient data structure for internal communication should be provided.

Flexible Autonomy of Holonic Agents

As mentioned before, the autonomy of a holonic agent can vary on a broad scale. On the one hand an agent that is not incorporated into any holonic structure remains completely autonomous within the framework of its architecture. On the other hand, agents who give up their autonomy as they become part of a holon, sometimes resign from their autonomy completely.

The agents in a holonic society can form, disband, and reconfigure holonic structures dynamically. Hence, they adjust their degree of autonomy in a flexible fashion. Furthermore, even a in stable holon the agents can adjust the degree of autonomy within the holon by modifying the allocation of resources.

The area of *computer supported cooperative work* integrates human and artificial agents. The holonic paradigm can support this integration by enabling humans to take part in holonic structures. A suitable interface agent allows the user to interact with other agents in the society. Whenever humans enter a holonic society, the degree of autonomy of the involved agents is modified by adjusting the resource allocation to the needs of the user.

Implementation of Holons

We shall now address issues of implementation of our model of agency and holonic structures.

Holonic Structures in Agent Societies

Let us first look at some possibilities for modeling holonic structures within a multiagent system and evaluate whether they are suitable for the design and implementation of *holonic multi-agent systems*. The following notions differ in the degree of autonomy the subholons have.

A Holon as a Federation of Autonomous Agents A straightforward way to model holonic structures is to assume that the sub-holons are atomic agents with their predefined architecture and the super-holon is just a new conceptual instance of the generic agent definition, whose slots are provided by the sub-holons and filled dynamically. In this case no agent has to give up its autonomy, and the super-holon is realized exclusively through communication and cooperation between the immediate sub-holons.

The most transparent way of cooperation is an *explicit coordination* by *commitments* via communication. Agents negotiate over joint plans, task distribution, or resource allocation. If commitment structures between agents can not be established through communication, *implicit coordination* may be established in two ways: firstly, the holons are designed such that a goal directed agents. Secondly, sub-holons that are able to reason about the goals and intentions of others could coordinate their actions without communication.

Several Agents Merging into One Another way of building a holon is to terminate the participating subagents and to create a new agent with capabilities that subsume the functionalities of these agents. In this case the merging agents give up their autonomy completely.

This approach violates the paradigm of autonomy since the merging agents loose their autonomy completely and delegate it to the super-holon. The realization of this approach requires procedures for merging and splitting holons that lead to the creation of a new agent. Especially for a heterogeneous group of agents this can be intractable. In the case of agents of the same kind with explicit representations of goals and beliefs (e.g., BDI agents) merging can be performed by creating an agent with the union of the sub-agents' beliefs and goals provided consistency is guaranteed. According to this model, agents cannot participate in more than one holon, unless they are copied.

A Holon as a Moderated Group As an alternative to the two extreme solutions, we consider a hybrid way of forming a holon where agents give up only part of their autonomy to the super-holon. This is achieved by one sub-holon acting as a representative or *head*. This head represents the holon to the outside world, i.e. to the rest of the agent society by its ability for a consistent interaction with its fellow agents. Its competence ranges from pure administrative tasks to the authority to issue directives to the other sub-holons. Furthermore, the head may have the authority to allocate resources to the other agents in the holon, to plan and negotiate for the holon on the basis of its sub-agents' plans and goals, and even to remove some sub-holons or to incorporate new sub-holons.

We can distinguish two methods to determine the head. Either, a new agent may be created for the lifetime of the holon, or one of the members of the holon takes the role of the head and gains the additional functionality. In the second case either one member of the holon is a priori predestinated for the leadership, or an election procedure is needed to promote one of the agents to leadership. Among other things, the representative serves as the interface and contact of the holon to the rest of the agent society. In this model, the autonomy of member agents is partially traded for control. Depending on the application domain, the competence of the representative may vary. The resulting super-holon's structure can range from a loosely moderated group up to a hierarchical structure. However, in any case, the members of the super-holon are still represented as agents, and, hence, we do not lose the capability to solve problems in a distributed fashion.

tions of these three approaches, we prefer the hybrid one: It allows for an explicit modeling of holons, a flexible formation of holonic groups, and the degree of autonomy of the participating agents is scalable. In this approach, the challenging problem is how resources are allocated within the holonic structures and this will be addressed in the following section.

Resource Allocation in Holonic Systems

How are the resources of the super-holon distributed to its immediate sub-holons? As in the previous section, we shall first present several possible approaches and discuss their strength and shortcomings. All approaches are based on the distribution of *abstract resources*[3] and an allocation of tasks to the sub-holons. The spectrum of mechanisms ranges from totally decentral approaches over moderated ones to more centralistic mechanisms. But first let us introduce the concept of an abstract resource.

Characterization of Abstract Resources All mechanisms are based on the distribution of *abstract resources* (Gerber & Jung 1998): this concept is a generalization of the familiar notion in computer science, where resources are mainly *computational time* and *memory space.* We use the term abstract resource to denote *any environmental device or tool that enhances the behavior of an agent:* A resource can enhance, or by its absence, constrain the agent's choice of action (e.g., information, perception, capabilities) or it can constrain the effects of the agent's execution of an action in a quantitative or qualitative fashion. This definition includes not only classical resources such as time and space, but also pieces of information, external tools, or other agent's capabilities. For instance, the construct of a *semaphore* can be viewed as an abstract resource: only one of the agents is able to get a hold on it and therefore is allowed to compute or act.

From the point of view of the super-holon, the capabilities of its sub-holons can be treated as abstract resources. So, in a holonic society, abstract resources can be, for instance, the *agents* building a holon, the *number of specialists* (i.e. the sub-sub-holons) for certain tasks, *different communication protocols,* etc.

Due to the recursive structure of a holon, a resource allocated to a super-holon can be redistributed at a finer granularity to its sub-holons. Hence, the coarse allocation of an abstract resource on a higher stage can be viewed as a "guideline" or constraint on the allocation to the lower ones. Of course some "typing" mechanism for abstract resources is required, which ensures that only meaningful resources are passed on.

[3]There is a "Collaborative Research Center" SFB-378 on resource adaptive cognitive processes in Saarbrücken, funded by the DFG, that researches such problems of resource bounded computation, abstract resources, and anytime algorithms. We refer the reader to their results (`http://www.coli.uni-sb.de/sfb378/index-en.html`).

be allocated can be either controlled a priori by some central device or they can be controlled by the individual holon members through a request. We discuss various options to realize resource allocation and comment on their suitability for different settings. In particular, we distinguish between *cooperative* and *non-cooperative* settings: In a *cooperative setting,* the participants have no local utility valuation. They are eager to maximize the utility of the group, and, hence, a utility measure is needed that enables the agents to decide locally whether a trade is globally beneficial. In a *non-cooperative setting,* each agent tries to maximize its local utility. So the possibility of utility transfer via side-paying must be given[4].

Market Mechanisms Market-based mechanisms serve to distribute and re-distribute tasks or resources among a holonic group of agents having roughly predefined schedules of tasks and resources. The purpose of a market mechanism is the iterated exchange of tasks and resources between the agents, so that the quality of the overall distribution increases (see e.g. (Fischer, Müller, & Pischel 1996)).

We can distinguish between coordinated and uncoordinated market mechanisms. In an uncoordinated market, agents negotiate and decide locally whether or not to agree on a deal. In a coordinated market a central instance, in our case the holon's head, moderates the trading process.

Game Theoretic Allocation Mechanisms In contrast to market mechanisms, a central instance has not only the job to mediate between agents, but to allocate one or more resources or tasks to a group of agents on the basis of reported valuations. For our purpose, this central instance is represented by the head of the holon. Again, we can distinguish between the cooperative case (i.e., truthful behavior is guaranteed) and the non-cooperative case where the agents may try to increase their own benefits at the expense of others. In the latter case, for the sake of global performance, it might be useful to apply *truth revealing mechanisms.* Furthermore, in this case, some sort of currency is needed that allows an explicit utility transfer. For a survey see (Fischer, Ruß, & Vierke 1998) where also multiple-item auctions are discussed.

Coordination Approaches based on Heuristics Decentral approaches may face sub-optimality if local knowledge is inadequate for decision making. In such cases, more optimal solutions can be achieved by shifting decision power to the holon's head.

[4]It might seem contradictionary to consider self-interested members of a holon since we require a common overall goal. However, such a common goal does not prevent conflicting goals of minor priority.

the competence to allocate resources within a holon. (Of course, members representing sub-holons can do so for their sub-domain.) The representative requests the member agents' relevant local information. By collecting this information, the head obtains a rather global picture for resource/task allocation.

Such an approach is presented in (Gerber 1998).

Discussion All these techniques could be used for resource and task assignment. However, they differ with respect to the role of the central control: Market mechanisms are based on very little control; the other approaches require a central control whose competences, however, vary. In the case of the non-cooperative setting, the head has only administrative competences and therefore mechanisms that enforce cooperative behavior have to be applied. In a cooperative case, decision power is split: the head decides on the basis of the local calculations of its sub-holons. Finally, in the central heuristic-based approach the holon's head has the complete resource allocation competence. Local information is only used as a heuristic.

More decentral approaches are well suited to deal with high complexity allocation problems, as they can often be reduced to a set of problems with less complexity. On the other hand, the use of these methods may be less than optimal. Hence, the choice of the appropriate mechanism depends on the nature of the application and there is a tradeoff between optimality and complexity.

Applications

We have applied the holonic agent methodology in a multi-agent fleet scheduling system, called TELETRUCK that has been developed at DFKI (Bürckert, Fischer, & Vierke 1998a; 1998b).

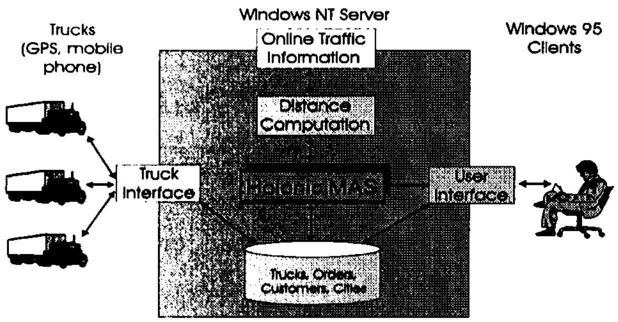

Figure 1: The TELETRUCK Architecture

The aim of the TELETRUCK system is to model the business processes, in particular the allocation of incoming transportation requests, of a transportation company. The company has a fixed number of transportation units like drivers, trucks, or trailers, where the units in turn may differ in many ways: trucks can be and those without, etc. The type and size of the loading space of the containers constrains the type of cargo that can be transported. Also human drivers differ in their supplied working time and the type of cargo he or she may transport depending on issues such as special training or certain licenses, e.g. for dangerous goods. Resources have to be scheduled in such a way that the transportation tasks at hand can be executed with minimal cost.

In our implementation, the multiagent system consists of holons of several types. For each transportation device of the forwarding company as well as for each of its drivers there is an agent, which administrates the resources of the device or the driver. These agents have their own plans, goals, and communication facilities in order to provide their resources for the transportation plans according to their role in the society. These agents merge with a *Plan'n'Execute Unit* (PnEU) and form a holon that represents a complete vehicle that is actually capable of executing transportation tasks. For example such a vehicle holon may consist of a PnEU, a driver, a truck, a trailer, and two containers, each being modeled as a sub-holon that is merged into the super-holon. (see Figure 2)

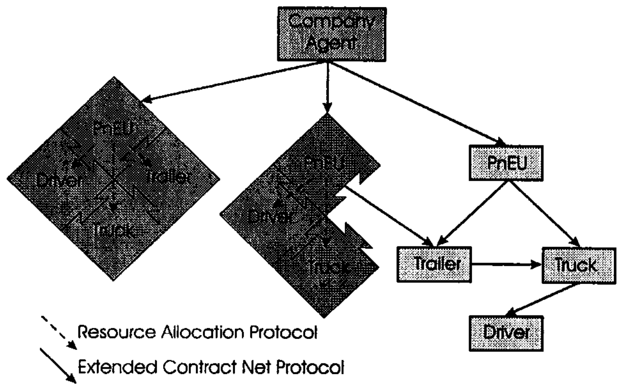

Figure 2: Holonic Agents in TELETRUCK

The PnEUs are special sub-holons, which coordinate the formation of the super-holons representing the transportation entities and plan the vehicles' routes, loading stops, and driving times. The PnEU is the head of the vehicle holon, represents it to the outside world, and is authorized to reconfigure it. A PnEU is equipped with planning, coordination, and communication abilities, but does not have its own resources. Each transportation holon that has at least one task to do is headed by such a PnEU. Additionally, there is always exactly one idle PnEU with an empty plan that coordinates the formation of a new holon from idle components if needed.

The vehicle holons are sub-holons of the super-holon that represents the entire transportation company. The holon that subsumes the complete agent society is nounces and distributes the incoming orders, gives the acceptance of the tenders, controls global optimization, coordinates the execution, and channels all communication of the system with the user, i.e. the dispatch officer. The company agent represents the society to the user—and according to future extensions, to partner companies to be represented also by such company agents—, and coordinates the internal cooperation and interaction between the PnEUs.

The holonic structures for the TELETRUCK project were implemented through moderated groups of agents. For the formation and coordination of a holon we have chosen an extension of the contract net protocol and a coordinated market mechanism (Fischer, Müller, & Pischel 1996). This decentralized approach is well suited for this complex setting since local information is sufficient for globally efficient resource and task distribution. The model has been implemented and tested in cooperation with a haulage company (Bürckert, Fischer, & Vierke 1998b).

Conclusion and Outlook

We have presented criteria for identifying holonic structures in a multi-agent system, that were inspired by biological systems in the sense of Koestler (Koestler 1967) and Minsky (Minsky 1986).

These criteria have been used as a design aid for modeling and implementing holonic agents in an industrial setting. The advantage of the holonic approach is the possibility to map an application domain directly onto a multi-agent system by agentifying entities at any granularity without losing higher level abstractions.

Our ongoing research deals with the usage of our concepts for further realistic applications in intermodal transportation, in telematics, and manufacturing systems. Another goal is to look at the scalability problem of multi-agent systems through resource management in holonic structures. In addition, we are working on a theoretical framework for an algebraic description of holonic agent systems, in particular an algebraic description of the merging operation that forms a superholon out of several sub-holons.

Acknowledgments

The authors are indebted to Hans-Jürgen Bürckert and Klaus Fischer for interesting discussions. This work was supported by the Siemens AG and the German Federal Ministry of Education, Science, Research, and Technology (BMBF).

References

Bratman, M. E. 1987. *Intention, Plans, and Practical Reason.* Cambridge, Mass.: Harvard University Press.

Bürckert, H.-J.; Fischer, K.; and Vierke, G. 1998a. Teletruck: A holonic fleet management system. In *Proceedings of the 14th European Meeting on Cybernetics and Systems Research.*

Transportation Scheduling with Holonic MAS – The TeleTruck Approach. In *Proceedings of the Third International Conference on Practical Applications of Intelligent Agents and Multiagents (PAAM'98).*

Castelfranchi, C. 1995. Guarantees for autonomy in cognitive agent architecure. In Wooldridge, M., and Jennings, N., eds., *Intelligent Agents: Theories, Architectures and Languages*, volume LNAI Volume 890. Heidelberg, Germany: Springer-Verlag. 45–70.

Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. In *Artificial Intelligence*, volume 42. 213–261.

Deen, S. M. 1994. A cooperation framework for holonic interactions in manufacturing. DAKE Centre, University of Keele.

Fischer, K.; Müller, J. P.; and Pischel, M. 1996. Cooperative transportation scheduling: an application domain for DAI. *Journal of Applied Artificial Intelligence. Special issue on Intelligent Agents* 10(1).

Fischer, K.; Ruß, C.; and Vierke, G. 1998. Decision Theory and Coordination in Multiagent Systems. Research Report RR-98-02, DFKI.

Gerber, C., and Jung, C. 1998. Resource Management for Boundedly Optimal Agent Societies. In *Proceedings of the ECAI Workshop on Monitoring and Control of Real-Time Intelligent Systems.*

Gerber, C. 1998. Bottleneck Analysis for Self-Adapting Multi-Agent Societies. In *Proceedings of the IEEE Joint Conference on the Science and Technology of Intelligent Systems.*

Klusch, M. E. 1999. *Intelligent Information Agents.* Springer Verlag.

Koestler, A. 1967. *The Ghost in the Machine.* Hutchinson.

Lange, D. 1998. Mobile agents: Environments, technologies, and applications. In *Proceedings of the Third International Conference on Practical Applications of Intelligent Agents and Multiagents (PAAM'98).*

Minsky, M. 1986. *The Society of Mind.* Simon and Schuster (Touchstone).

Parsons, T. 1969. *Politics and Social Structures.* New York: Free Press.

Parunak, H. 1987. Manufacturing experience with the contract net. In Huhns, M., ed., *Distributed Artificial Intelligence.* San Mateo, CA: Morgan Kaufmann Publishers. 285 – 310.

Rao, A. S., and Georgeff, M. P. 1991. Modeling Agents Within a BDI-Architecture. In Fikes, R., and Sandewall, E., eds., *Proc. of the 2rd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, 473–484. Cambridge, Mass.: Morgan Kaufmann.

Russell, S., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach.* Prentice Hall.

Russell, S. J., and Wefald, E. H. 1991. *Do the right thing: studies in limited rationality.* MIT Press.

Wooldridge, M. J., and Jennings, N. R. 1995. Intelligent agents: Theory and practice. *Knowledge Engineering Review* 10(2):115–152.