

JIAC IV – An Open, Scalable Agent Architecture for Telecommunications Applications^{*}

Ralf Sessler, Sahin Albayrak

DAI-Lab, TU Berlin
Sekt. FR 6-7, Franklinstr. 28/29
D-10587 Berlin, Germany
sesseler@cs.tu-berlin.de, sahin@cs.tu-berlin.de

Abstract

JIAC IV is a generic agent toolkit with an emphasis on the realisation of applications in the areas of telecommunications and electronic commerce. We motivate the need for open, scalable, and flexible systems and hence the advantages of agent technology in these areas. Then, we describe the basic architectural principles of CASA, which builds the structural and functional base of agents in JIAC IV. The component framework of CASA makes agents highly scalable and allows reconfiguration at run-time. The control architecture of CASA combines reactive, deliberative, and interactive capabilities to control agent behaviour in a flexible manner.

1 Introduction

The future of the telecommunications world will be dominated by increasing customer demands for new services and the integration of electronic and mobile commerce. One of the favourite slogans today is “information anywhere, anytime” claiming the ubiquitous availability of information and services via electronic networks like the Internet. On the other hand, deregulation leads to open markets and increasing competition. Therefore, the telecommunications companies are evolving from simple network providers to suppliers of service infrastructures supporting flexible and dynamic provisioning and combination of services.

We distinguish basic telecommunications applications from telematics services. By the notion of a telecommunications application, we generally denote any kind of information system that is necessary to run, to manage, and to administrate computer networks. On the contrary, the notion of a telematics service emphasises the aspect of trading between vendor and customer based on such networks. Hence, the telecommunications applications have to build the foundation for telematics services.

From this distinction, three main roles emerge in the area of telecommunications, each of them having specific re-

quirements of their own. First, the suppliers of the telematics services need to develop and introduce new services rapidly and maintain existing services on a robust and secure service environment. The customers of these services demand convenient, personalised access and mobility support. Both benefit from the possibility to combine services on demand, leading from basic to high value services. Finally, the providers of the telecommunications infrastructure have to deal efficiently with increasingly complex and open networks with heterogeneous structures and technologies. Also, they have to take into account the requirements of the other two roles, because they provide the basic layer.

Agent technology [1, 2] seems to be a promising way to realise such telecommunications infrastructures because of its inherent openness and distribution as well as because of the flexible and interactive capabilities of agents. The agent toolkit JIAC IV (Java Intelligent Agent Component-ware, Version IV) has been developed as a means to build and deploy telecommunications applications and telematics services as multi-agent systems to benefit from these advantages.

In the following, we provide a short overview of the JIAC IV agent toolkit as a whole. Then we concentrate on CASA (Component Architecture for Service Agents), which provides the structure and control mechanisms of single agents in JIAC IV. We describe how the component framework enforces openness and scalability by modularisation and how this modularisation facilitates the reuse of components at design-time as well as the reconfiguration of agents at run-time. The needed autonomy and flexibility of agents is ensured by a control scheme, which integrates reactive, deliberative, and interactive behaviour. Finally, we place our work in the context of related work and draw some conclusions.

2 Overview of JIAC IV

JIAC IV is intended as a comprehensive toolkit for developing and deploying agent systems covering design methodology and tools, agent languages and architecture, a

^{*} This work was funded by T-Nova Deutsche Telekom Innovationsgesellschaft mbH.

JIAC IV Agent Toolkit					
Architecture (CASA)		Development		Run-Time	
<i>components</i>	<i>control architecture</i>	<i>methodology</i>	<i>tools</i>	<i>infrastructure</i>	<i>administration</i>
components roles messages	knowledge languages reactive deliberative interactive	analysis design implementation evaluation	configuration language compilers debugger	market places migration white pages yellow pages security	configuration fault management end user access

Figure 1. The JIAC IV Agent Toolkit

FIPA compliant infrastructure, management and security functionality, and a generic scheme for user access. For a more comprehensive overview, we refer to [3].

The development process is guided by an agent-oriented software engineering model, which is tailored to the specifics of JIAC IV. It comprises tools for agent specification including compilers for the different agent languages and tools to analyse and debug the running system.

On the single-agent level, CASA provides a scalable component framework and a flexible, knowledge-based behaviour control scheme, which we will describe further in the following sections.

The agent infrastructure has to facilitate dynamic interactions between agents. Based on the specifications of FIPA [4], it comprises the communication infrastructure as well as services to administer the agents of a society (Agent Management Service, AMS) and the services they supply (Directory Facilitator, DF).

Commercial applications have special requirements on the reliability and trustiness of the system. Therefore, the JIAC IV toolkit provides several management and security functionalities, which can be easily integrated and adapted as needed thanks to the overall scalability. Management services include configuration, fault management, and logging of system processes, which may be used for analysis as well as for accounting. Security issues are addressed by authorisation, authentication, and privacy mechanisms.

To provide a convenient way for the user to interact with the agent system, JIAC IV contains a generic scheme to translate agent functionality into human accessible services. Thereby, dedicated agents provide graphical user interfaces, which can be used by a single access point.

3 Component Framework

Agents have to be adaptable to different purposes, tasks, and domains not only by varying knowledge but also by specific capabilities to process this knowledge accordingly. Therefore, CASA agents have a modular internal structure consisting of an open set of components that can be adjusted to different requirements at design-time as well as at run-time. Since components are reusable, new

agents can be created using existing components for generic functionalities, which reduces development effort to application-specific implementations.

CASA provides a framework to manage the components of an agent and their interactions. Components are integrated into an agent by a common interface to the core agent, which realises the internal infrastructure. By this interface, the core agent controls the component and its configuration and gives it access to the message passing mechanism for component interactions. Among one another, components are identified by the roles they take within the agent describing their interactive capabilities to abstract from different implementations of the same functionality.

3.1 Component Interactions

Components interact by passing messages relative to their roles. Received messages are stored in a buffer to allow asynchronous processing. Thus, interactions are decoupled both from time and from the component realising a role, which is needed to allow changes of the component set at run-time without affecting the functioning of the agent.

The interacting components know only the roles of their counterparts. Thus, dependencies only exist between roles, but not between specific components. The roles serve as an interface description for the interactions of components and for their integration into the agent. A single component can have several roles, but each role must be unique to an agent to identify exactly one component.

Each role belongs to one or more groups of roles. Such a group subsumes roles with a common functionality. A group has the same kind of interface description as a role, which their roles inherit. On top of the hierarchy is a most common role subsuming all groups.

The specification of a role covers two parts. First, it declares properties, by which a component implementing the role can be configured and its state retrieved.

For interactions, a role declares the message types it can process and which roles and groups of roles have permission to send them. Thus, control structures consisting of several dependent roles can be defined.

3.2 Components of the Core Agent

The core agent itself consists of three components. The Agent Kernel manages the components and the agent as a whole. The delivery of messages between components is done by the Message Server. The Control Cycle organises all processing of the agent including that of messages.

Agent Kernel

The main task of the Agent Kernel is to manage the component structure of an agent. Also, it represents the agent as a whole and its properties and life-cycle state.

The common interface of the components allows the Agent Kernel to configure them by changing their properties. Properties can be declared for roles as well as for components. The properties of the agent are realised as properties of the Agent Kernel. A special property is the life-cycle state that determines its status of activity. All components have the same life-cycle state as the agent, except while they are added or removed or when they are defective.

Via the Agent Kernel, components of an agent can be added, removed, and exchanged. The Agent Kernel only allows adding a component, if no role of it is already occupied by an existing component. To add a component, its interface is connected to the components of the core agent for mutual access until it is removed again. Exchanging a component means adding a component while removing all previously existing components with corresponding roles.

The Agent Kernel creates an agent out of a specification containing a list of components and their initial properties. This is done by creating instances of the components, configuring them by the properties, and adding them to the agent.

Message Server

The Message Server provides the infrastructure for component interaction. It manages an address list that associates roles and existing components of the agent (Figure 2).

To send a message, a component creates it and passes it to the Message Server. The Message Server validates the message by verifying that the sending component takes the role declared as sender and that this role has the permission to send messages of the given type to the receiver as specified in the role definition for the receiver address.

If the message is not valid or no component exists in the address list for the receiving role, the sender is informed accordingly. Otherwise, the Message Server delivers the message by adding it to the end of the queue of the receiver, which is part of the common interface of the components.

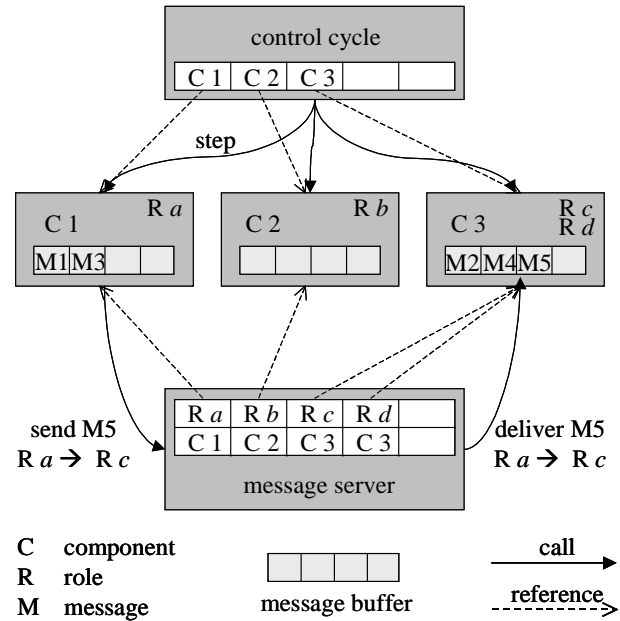


Figure 2. Interaction by Message Passing / Execution

Control Cycle

The Control Cycle provides and manages the processing resources for the components. For controlled interruptions, all processes of the components have to work step by step. Thus, components are ensured to be in a stable state when they are removed.

Components declare by the common interface if they have steps to execute. Then, the Control Cycle assigns processing resources as available. Steps can be executed in a sequential, parallel, or mixed mode, but only one step per component is executed at one time.

For message processing, the Control Cycle removes the first message from the buffer queue and passes it to the component for processing (Figure 2).

3.3 Reconfiguration at Run-Time

A main design motivation for the component framework is the reconfiguration of agents at run-time. Especially for agents that provide services, tasks have to be changed, updated, or adapted without the agent or some of its services being not available in the meantime. Thereby, not only the knowledge like facts and operators, but also the component functionality may need to be revised.

Changes of knowledge are a special case of component reconfiguration of the components containing that type of knowledge. The configuration of any component at run-time can be done via the Agent Kernel using tools and components for configuration by changing the declared properties of roles and components.

To change the functional part of an agent, components are added, removed, or exchanged. Like creating, adding a component simply means to configure it and to connect it to the core agent. On removal, the component gains no more new processing resources from the Control Cycle, but it can finish the current step. Then it is disconnected from the components of the core agent. Components can register at the Agent Kernel to be informed about changes of presence of components for some or all roles.

Exchanging a component means to replace a role without affecting the current working of the agent, especially of components interacting with that role by messages. The new component replaces the old one directly at the Message Server, so that at every time during the exchange messages can be sent to that role. In addition, the messages addressed to the shared role waiting in the buffer of the old component are moved to that of the new one ensuring no message can get lost.

The exchange of a component takes place by simultaneously adding the new and removing the old component. During the exchange, neither component has access to processing resources. The new component takes over the configuration and run-time state of the old one given by their properties. Thus, it can continue the work of the former without interruption or loss of information.

Since any component can take several roles that have to be replaced at the same time and possibly by different components, component exchange is done for a set of new components. All existing components taking at least one of the roles of the new components are removed to ensure uniqueness of roles.

By the described mechanism, component exchange is completely transparent to the remaining components even with respect to ongoing interactions. All dependencies between components are restricted to their roles ensuring an open and scalable structure of the agent.

4 Control Architecture

The components of the core agent do not determine the mechanisms to control the behaviour of an agent, but they provide an open framework to design functional architectures. Using the CASA component framework, such a control structure is defined by a set of roles. The interdependencies between these roles are stated by the message types, a role accepts to receive from a selected set of other roles, or at a more abstract level as the functionality, a role has to provide to other roles.

Before proposing a default architecture for service agents in CASA, we present a basic scheme for reactive, deliberative, and interactive behaviour. This scheme uses a knowledge-based approach in the tradition of artificial intelligence [5] and the BDI-theory of agency [6, 7] in a

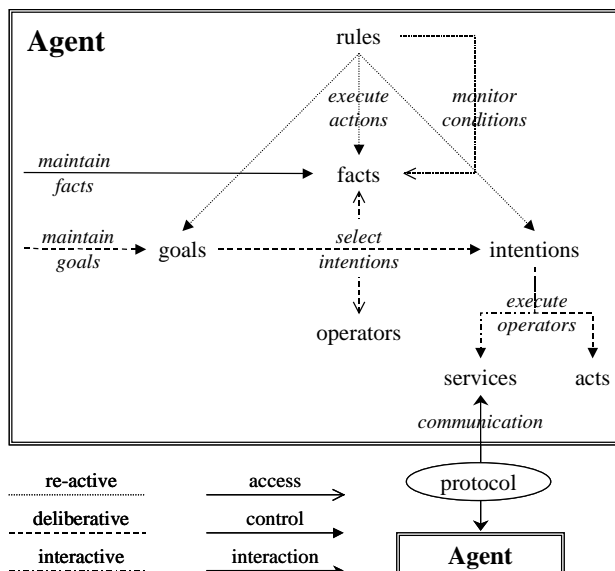


Figure 3. Behaviour Control Scheme
pragmatic way without claiming to provide a base for rationality or intelligence in a human sense. Instead, the formal representation of declarative knowledge serves mainly as a common ground for interactions between agents providing a flexible and open scheme to express meaningful contents in a standardised and expressive way. On the other hand, the formal representation of procedural knowledge enables more autonomous, flexible, and explicit control of the behaviour of an agent. Combining both aspects, the formal descriptions of services enable reliable and dynamic interactions between agents.

4.1 Control Scheme

Figure 3 gives an overview of the control scheme for the default architecture of CASA. An agent represents assumptions about the state of its environment as factual knowledge. Its goals are states to reach and its intentions actions to take. Rules describe the reactive dispositions, while operators describe the options for deliberative actions and interactions. All knowledge is expressed and structured by terminologies and representation schemes contained in ontologies for different domains.

An agent updates its facts and goals permanently, in order to reflect the current environmental and motivational state, and reacts to new situations. To reach its goals, it selects actions by deriving intentions using appropriate operators. The intentions are coordinated and the contained operators executed. To interact via services, agents communicate guided by protocols.

4.2 Default Architecture

The CASA default architecture defines a set of component roles mapping this control scheme into a functional con-

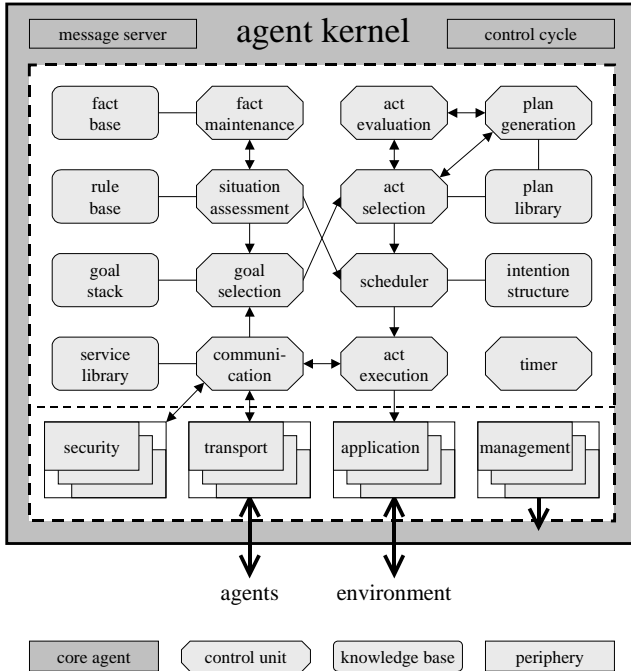


Figure 4. Roles of the Default Architecture. This architecture is still open for different implementations of the defined roles, but it is reasonable to implement default components for application-independent roles. Thus, an agent designer can compose a new agent reusing generic components and only needs to develop new components for application-specific tasks. Also, he can use only particular parts of the default architecture or even another control structure, as long as its interactive behaviour conforms to the requirements for service interactions (see Section 4.3).

The default architecture consists of four parts: the core agent, the control unit, the knowledge base, and the periphery (Figure 4). The components of the core agent build the component system as described in Section 3.2. Control unit and knowledge base together realise the control scheme. The knowledge base is a storage divided into roles for the different types of knowledge, while control functions are assigned to the roles of the control unit. The periphery contains additional types of roles for auxiliary and application-specific tasks.

Knowledge Base

The six roles of the knowledge base are the fact base, the goal stack, the intention structure, the rule base, the plan library, and the service library. Each of them is responsible to store, maintain, and provide knowledge of the according type.

The fact base contains the assumptions of an agent about the current state of the world. It has to be consistent at the ground level, i.e. no two facts contradict each other. The goal stack is an ordered list of the current goals. Its order

reflects the priority of the goals and hence the probability to be pursued next. The intention structure reflects the intentions coordinated according to conflicts, redundancies, and execution order.

Rule base, plan library, and service library are mere containers for the procedural knowledge containing rules, operators, and services.

All roles of the knowledge base provide the same functionality to manage their contents by accessing and changing knowledge and monitoring changes of knowledge.

Control Unit

The control unit operates on the knowledge stored in the knowledge base. Its roles can be grouped by the following four tasks: representation, reaction, deliberation, and interaction.

The representational tasks concern the environment and as a special aspect time. The fact maintenance manages the factual knowledge to reflect the current state of the world. It has to keep the facts up-to-date and coherent, but thereby it depends on the information other components gain about the environment. The fact maintenance also realises the reasoning capabilities of the agent by handling not only isolated facts, but also complex propositions that express changes of the world or whose truth has to be evaluated relative to the current facts. Since time changes continuously, the current time is not handled simply as a frequently changing fact. Instead, the timer role serves as an internal clock offering notifications for absolute, relative, and periodical time events.

For reactions, the situation assessment monitors changes of factual knowledge for the situations described by the rules contained in the rule base. If such a situation occurs, the situation assessment changes the knowledge of the agent according to the action part of the rule resulting in a behaviour adapted to the new situation.

Deliberative behaviour emerges from six roles leading from goals to acts: goal selection, act selection, plan generation, act evaluation, scheduler, and act execution. The goal selection organises the motivational state of an agent. All roles can add new goals. The goal selection updates, evaluates, and coordinates the goals in the goal stack and selects those to pursue next. These goals are passed to the act selection to find actions to reach them. Therefore, suitable operators have to be found and the best alternative has to be selected as a new intention. Since both tasks may depend on the application domain, they can be delegated to own roles for plan generation and act evaluation. These roles are separated from the act selection in favour of modularisation to allow different domain-specific implementations. The act selection thereby only realises the basic capabilities. The plan generation combines operators to plans in a goal driven manner. The act

evaluation compares different acts to decide between alternatives.

The scheduler coordinates the intentions of an agent resolving conflicts and utilising redundancies. It determines the order of execution and selects the next intentions to execute, which are passed to the act execution. The act execution interprets the operators contained in the intentions. Depending on the operator type, it is either executed by the act execution itself or passed to the appropriate role. The results from executing operators are reported back.

Operators for services are executed by the communication. This role organises the interactive behaviour of an agent. It initiates the usage of services and handles service requests. In addition, the communication creates speech acts to send and processes received speech acts.

Periphery

Since the structure of the periphery is not strictly determined by the CASA control scheme, it only defines groups of roles each with a common functionality: the application group, the transport group, the security group, and the management group.

Application-specific tasks of an agent are realised by the roles of the application group. This includes interactions with the environment like sensory input and behavioural output, but not communications. These roles can affect the agent behaviour by stating new facts and raising new goals. For deliberative acts, each primitive operator denotes an application role that implements its execution.

The communication role is just responsible for organising communication at a higher level, while the transport of speech acts between agents is left to the roles of the transportation group. Each role of this group enables communication via a single communication channel like TCP/IP, SSL, or IIOP. At run-time, it constitutes an address by which the agent can send and receive speech acts. Received speech acts are passed to the communication role, which also is the source for speech acts to send.

To address security issues in communications between agents, the communication role has to rely on the roles of the security group. Security requirements for services may include among other authentication of the communication partner by certificates, privacy of communication via encryption, authorisation for service usage, and trust relationships between agents [8, 9].

The roles of the management group control the agent and its components at run-time. Their instances gain direct access to the Agent Kernel and its functionality. There are two directions of management: introspection and manipulation. The introspection collects and analyses run-time information about the agent. The manipulation can

modify properties and the component set. Management tasks include among other reconfiguration, fault detection and correction, and performance measurement and improvement.

4.3 Agent Interaction

The control scheme and the default architecture of CASA integrate the interactive capabilities of agents into the overall behaviour control. All interactions between agents are based on services, which are actions one agent performs in behalf of another. A service is described as an operator for planning from the point of view of the customer of the service. Thus, the customer can use the services of other agents to decide about its behaviour as any other action, while only the execution is delegated to another agent, which is called the provider of the service. This allows for flexible and dynamic selection and combination of services for an agent to pursue its goals.

Communication between agents is based on shared ontologies, an agent communication language, and protocols to ensure interoperability.

Since all interactions are according to a general scheme of service usage, CASA defines a generic meta-protocol, which handles the common aspects of service interactions. It is initiated by the customer when requesting the service. The protocol comprises three steps. First, a generic negotiation phase is used to establish the possibility for interaction. The provider ends this step by accepting or refusing the initial request. In the second phase, the customer may select one of several providers by negotiating service-specific parameters using an embedded protocol. Then, another embedded protocol may be used for service provision. Finally, the provider ends the meta-protocol by communicating the result of the service usage. While the generic meta-protocol is an integral part of CASA, the embedded protocols for negotiation and service provision are open for specific communications.

5 Related Work

Only few agent systems attempt to provide comprehensive toolkits covering all aspects of agent development and deployment. ZEUS [10] and AgentBuilder [11] both emphasise the relevance of supporting the development process by a methodology and a set of tools for agent specification and system analysis. Also, each offers a sound control architecture for knowledge-based agent behaviour including appropriate specification languages. ZEUS utilises a traditional planning approach and integrates interactive capabilities. AgentBuilder uses a more reactive control structure based on Agent-0 [12], but supports only simple message passing for agent interaction. An infrastructure for multi-agent systems is provided by ZEUS via dedicated Utility Agents, but is missed in

AgentBuilder. Both systems do not address security issues or a generic user access scheme.

Many agent architectures concentrate on the internal mechanisms for knowledge-based behaviour control. Well-known examples are PRS/dMARS [13, 14] or InTeRRaP [15]. These systems are usually based on theories of action like BDI [6] and are inspired by the tradition of artificial intelligence. To put the focus on the control architecture leads to powerful systems, but most of these systems lack to support the development process and to provide a reliable infrastructure.

Other systems mainly provide an agent infrastructure, but rely on only minimal agent models. Good examples are JATLite [16] and the FIPA-compliant JADE [17]. These systems support only the multi-agent aspect, but do not integrate the single-agent view. Instead, agents are implemented by conventional programming techniques and languages like Java or by using a control architecture, which is not part of the system itself. Either way, the multi-agent functionality is not an integral part of the single-agent behaviour mechanisms and vice versa, which means that these will have to be adapted for each application anew.

6 Conclusion and Outlook

JAC IV is a comprehensive agent toolkit that is already capable of providing the infrastructure for real-world telecommunications and electronic commerce applications.

By its openness and scalability, the toolkit makes it easier to develop new applications and to adapt running applications according to dynamic changes of customer demands and by the integration of new services. This is achieved at the single-agent level by the component framework of CASA and at the multi-agent level by the agent infrastructure. Furthermore, application development and deployment is supported by a process model and a set of tools.

The control architecture allows to create agents that are flexible enough to fulfil the tasks delegated to them by the user reliable and autonomous. The interactive capabilities enable agent societies with dynamic selection and combination of services.

The end user of applications realised with JAC IV is provided with a uniform and convenient way of accessing the services, which are realised by agents. The supply and integration of security and accounting functionalities are prerequisites for commercial applications and are needed for user acceptance.

Since agent systems tend to be very complex systems, there are still some improvements to make on the concepts and the implementation of JAC IV. A critical point is the competence needed by agents to make decisions accord-

ing to the intentions of the user it is acting for. The mechanisms of the control architecture of CASA are still generic. We aim to integrate aspects like personalisation and learning that allow to adapt the decision making to the demands and preferences of the individual user.

References

- [1] Michael Wooldridge, Nicholas R. Jennings. *Intelligent Agents: Theory and Practice*. Knowledge Engineering Review, October 1995.
- [2] Nicholas R. Jennings, Michael J. Wooldridge (Eds). *Agent Technology: Foundations, Applications, and Markets*. Springer-Verlag, 1998
- [3] Stefan Fricke, Karsten Bsufka, Jan Keiser, Torge Schmidt, Ralf Sessler, Sahin Albayrak. *Agent-based Telematic Services and Telecom Applications*. in: *Communications of the ACM*, April 2001
- [4] FIPA 2000 Specification. www.fipa.org.
- [5] Stuart Russel, Peter Norwig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [6] Philip R. Cohen, Hector J. Levesque. *Intention is Choice with Commitment*. *Artificial Intelligence*, 42(3), 1990.
- [7] Anand S. Rao, Michael P. Georgeff. *BDI Agents: From Theory to Practice*. *Proceedings of the First International Conference on Multi-Agent Systems, ICMAS 1995*.
- [8] FIPA '98 Specification. www.fipa.org.
- [9] Chelliah Thirunavukkarasu, Tim Finin, James Mayfield. *Secret Agents – A Security Architecture for the KQML Agent Communication Language*. *CIKM'95 Intelligent Information Agents Workshop*, 1995.
- [10] N. Azarmi, N. Thompson. *ZEUS: A Toolkit for Building Multi-Agent Systems*. *Proceedings of fifth annual Embracing Complexity conference*, Paris, 2000.
- [11] Reticular Systems. *Agent Builder - An Integrated Toolkit for Constructing Intelligent Software Agents*. www.agentbuilder.com
- [12] Yoav Shoham. *Agent-Oriented Programming*. *Artificial Intelligence*, Vol 60, pages 51-92, 1993
- [13] Michael P. Georgeff, Francois F. Ingrand; *Decision-Making in an Embedded Reasoning System*. in: *Proceedings of the Eleventh International Jopint Conference on Artificial Intelligence (IJCAI-89)*
- [14] Mark d'Inverno, David Kinny, Michael Luck, Michael Wooldridge; *A Formal Specification of dMARS*. *Proceedings of the Fourth International Workshop on Agent Theories, Architectures, and Languages, ATAL'97*
- [15] Christoph G. Jung & Klaus Fischer, *Methodological Comparison of Agent Models*; DFKI Research Report RR-98-01
- [16] JATLite, java.stanford.edu
- [17] F. Bellifemine, A. Poggi, G. Rimassa. *Developing multi-agent systems with JADE*. *Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL-2000)*, Boston, MA, 2000.