Agent Technologies in Virtual Enterprises

Klaus Fischer*and Christian Ruß†
German AI Research Centre — DFKI GmbH
Stuhlsatzenhausweg 3, D-66123 Saarbrücken

Abstract

Decreasing innovation cycles, changing market situations as well as growing specialisation in individual market segments demand new ways of economic thinking, increasingly forcing enterprises into cooperations, sometimes even with direct competitors. Presently discussed and designated as the corporate and cooperation model of the future is the so-called *virtual enterprise*. In this paper, we advocate the use of *intelligent agents* as a useful metaphor and as a software engineering methodology for the design and the operation of virtual enterprises. We focus on how agents can support the management of virtual enterprises through the Internet by performing tasks such as negotiation, information retrieval and extraction, and the participation in in electronic markets. This paper does not describe completed research. It rather offers a perspective of the high potential of agent-based technology for one of tomorrow's key industrial areas.

1 Introduction

At present, the time interval within which a product can be marketed successfully is dramatically decreasing. Due to changing market situations as well as growing specialisation in individual market segments, enterprises are increasingly faced with hard terms of competition. This forces them into cooperations, sometimes even with direct competitors.

Currently discussed and designated as the corporate and cooperation model of the future [Byrne et al. 93] is the so-called *virtual enterprise (VE)*. Unfortunately, companies setting up or constituting a virtual enterprise are faced with new *interorganizational* coordination problems because a virtual enterprise *"refuses an institutionalization - it has no central office, no hierarchy and no vertical integration [Arnold et al. 95a]"*.

The German Research Centre for Artificial Intelligence (DFKI GmbH) investigates AI technologies in the context industrial application scenarios. Application domains investigated in the multi-agent system (MAS) group are oder dispatching in distributed haulage companies, flexible manufacturing, banking and electronic commerce applications (see Figure 1). In this research we observed that in these domains an immense virtualization process is currently taking place: in the highly competitive haulage business shipping companies form temporary inter-regional cooperations, models of *virtual factories* in which software modules control the manufacturing processes in corresponding *physical factories* have been presented [Raulefs 94] and virtual banks are becoming reality [Visser 98].

*E-mail: Klaus.Fischer@dfki.de
†Christian.Russ@dfki.de

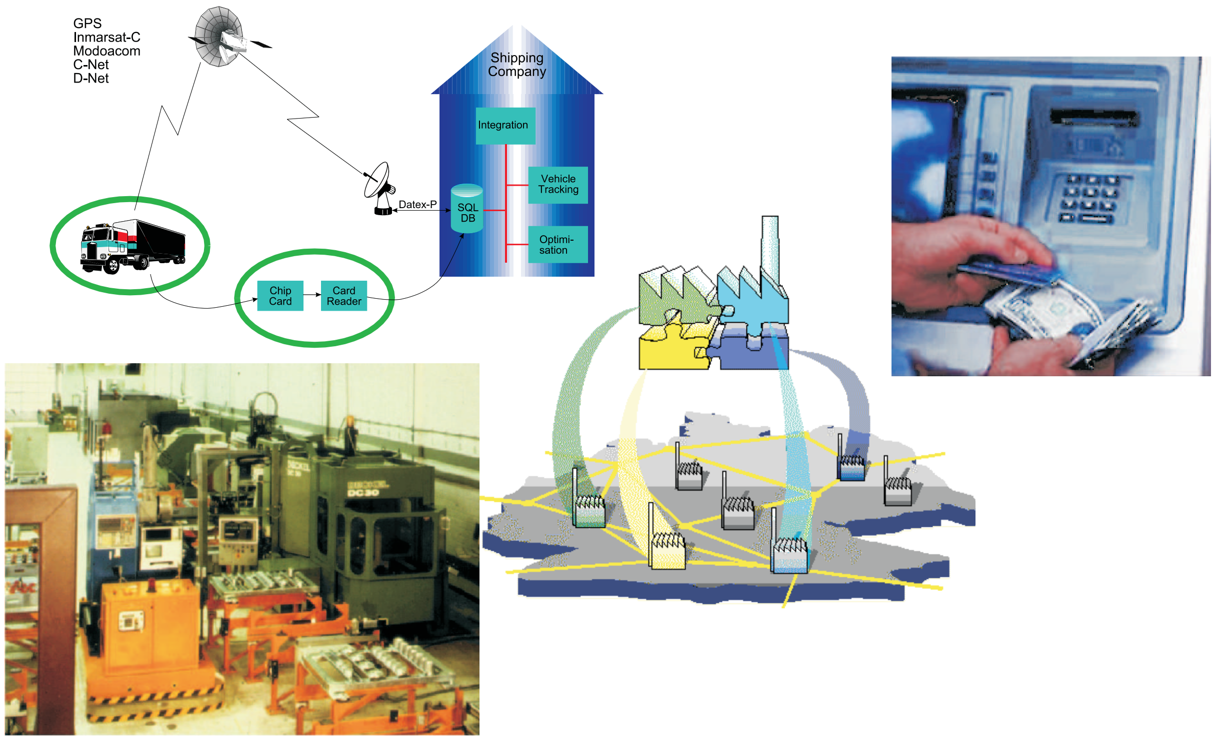

Figure 1: Domains in which virtual structures proliferate.

Our research results show that *agent technology* is the preferable technology for enabling a flexible and dynamic decentral coordination of spatially distributed entities in these application domains. This technology changes the metaphor for human-computer interaction from direct manipulation by the user to an indirect management of background agent processes because intelligent agents can—furnished with knowledge about the goals and preferences of their users and the ability to make decisions—autonomously perform a lot of coordination and everyday tasks on behalf of their users.

In this paper we advocate the use of agent technologies as a metaphor and an engineering paradigm for a variety of issues centred around the design and the operation of virtual enterprises. We present a distributed architecture for the electronic management of virtual enterprise using the Internet as a vehicle. We propose an agent-based methodology that enables a human product manager to generate and formalise goal descriptions for virtual enterprises, to find potential partners in the net, and to agree on a structure of a virtual enterprise through a negotiation process. Our claim is that all of these processes can and should be substantially supported by agent technology. However, this paper does not describe a piece of completed research, but it rather offers a new perspective of how intelligent agents technology can be used in one of tomorrow's key industrial areas.

2 Virtual Enterprises

The term *virtual enterprise (VE)* is generally attributed to Mowshowitz [Mowshowitz 86]. The term obtained its current importance for business economics from Davidow and Malone's landmark book [Davidow & Malone 92]. In [Byrne et al. 93], Byrne defined the

VE as a temporary network of independent companies to share skills, costs, and access to each other's market. Throughout this paper, we define a VE along the lines of [Arnold et al. 95b, p. 10]:

> A *virtual enterprise (VE)* is a cooperation of legally independent enterprises, institutions or individuals which provide a service on the basis of a common understanding of business. The cooperating units mainly contribute their core competences and they act to externals as a single corporation. The corporation refuses an institutionalisation e.g., by central offices; instead, the cooperation is managed by using feasible information and communication technologies.[1]

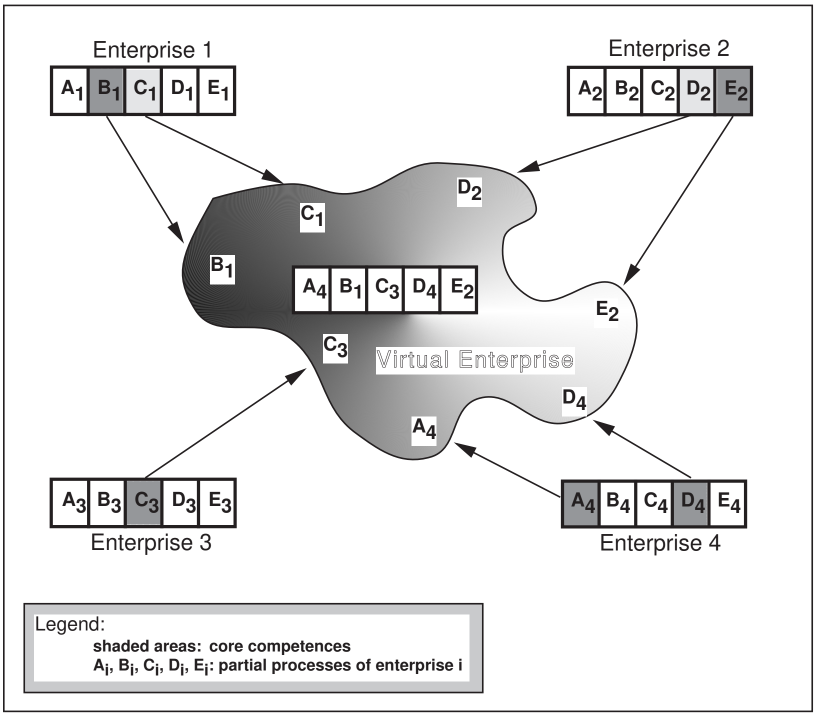

Figure 2: Abstract picture of a virtual enterprise

Figure 2 illustrates this definition. Enterprises 1–4 form the independent units, which work together within a value chain of five steps and provide the partial business processes[2] A_i to E_i (the services mentioned in the definition). Partial processes that include the corporations' core competencies are represented in Figure 2 with a shaded background. A *core competence* (see also [Quinn & Hilmer 95]) is defined as the skills, technologies, and know-how that are crucial for the success of an enterprise and that enable the company to produce and offer a marketable product.

We can identify several phases in the mananagement of a VE. First of all there must be an idea which kind of product or service the VE wants to provide. We call the party which defines this idea the manager of the VE. The second step is to design a business process

[1] Translated from the German original.

[2] These business processes shall be called *partial processes* to distinguish them from the *global process* describing the VE.

description which specifies how the product or service should be produced. Already this second step can be done in a distributed manner where several partners negotiate about the best way to design such a process. In the third step the global business process description is broken down into partial business process descriptions, which when put together again form the global business process. The forth and last step in the design of a VE is to assign the partial businiss process to partners which are then all together able to produce the product or provide the service. For this last step we can distinguish tow extrem instances: (1) the assignment of the partial business to partners is done once and for ever or (2) the assignment is done dynamically for each instance of the production of an individual product or service. These are two extreme positions and of course we can allow mixtures where a group of partners is selected which commit themselves to do parts of the business process. Additionally, certain partial business processes are dynamically assigned to changing partners, depending on which of the providers is the best one (e.g. with respect to costs). However, right now we do not have any clear criteia on which this decision could be based.

3 Agent Technology

The increasing complexity of organisations and computer-controlled technical processes and systems makes it impossible to design them as monolithic entities and to maintain and monitor them by a centralised control. The rapid development of agent-based technologies since the beginning of the 1990s (see [Wooldridge & Jennings 95, Wooldridge et al. 96]) has to be viewed in the light of these requirements.

According to [Wooldridge & Jennings 95] *intelligent agents* are computer systems that autonomously perform tasks in an open environment, that are re- and pro-active, and that are able to interact with other agents. The research in the area of *multiagent systems (MAS)* is concerned with coordinating intelligent behaviour among a collection of autonomous (possibly heterogeneous) intelligent agents [Kraus et al. 95]. A MAS consists of a group of agents that can take specific roles within an organisational structure. The step from isolated single-agent scenarios to open multiagent systems offers the new quality of *emergent behaviour*: the group of agents is more than the sum of the capabilities of its members.

Over the past few years, the development of control architectures that allow to design agents which can flexibly work in a MAS environment, which means that they can behave in a reactive as well as a deliberative way in a dynamically changing environemnt, has become an important theme of research (see [Firby 89, Hayes-Roth 95, Bonasso et al. 96]). In the multi-agent systems group at DFKI in Saarbrücken we developed the agent architecture InteRRaP (see [Müller 96]), which integrates social, deliberative, and reactive abilities in a uniform model. InteRRaP describes an agent by a world interface, a control unit, and a knowledge base. The control unit consists of three layers: the behaviour-based layer (BBL), the local planning layer (LPL), and the cooperative planning layer (CPL). The agent knowledge base is structured correspondingly in a world model, a mental model, and a social model.

The different layers correspond to different functional levels of the agent. The purpose of the BBL is to allow the agent to react to certain critical situations (by so-called *reactor patterns of behaviour* (PoB)), and to deal with routine situations (using *procedure PoB*).

Reactors are triggered by events recognised from the world model that incorporates the agent's object-level knowledge about its environment. The LPL gives the agent the ability of longer-term deliberation. It builds on world model information, but additionally uses the agent's current goals and local intentions maintained in the mental model part of the knowledge base, as well as domain-dependent planning mechanisms available. The CPL finally extends the planning functionality of an agent to *joint plans*, i.e. plans by and/or for multiple agents that allow to resolve conflicts and to cooperate. Apart from world model and mental model knowledge, the CPL uses information about other agents' goals, skills, and commitments stored in the social model of the knowledge base.

From a more intuitive point of view of the context in a VE, agents in a VE, which are designed according to the INTERRAP architecture, use the CPL for negotiation processes with other agents, which possibly belong to other enterprises. They use the LPL to do local problem solving (e.g., finding out if they are able to perform a certain partial business process and, if yes, providing a plan how to actually execute the partial business process) and the BBL to actually execute the actions they committed themselves to in negotiation processes and local problem solving.

4 Conclusions

The increasing importance of VEs and the rapid development of information technology offer a challenging new application area for intelligent agents. In this paper, different possibilities were illustrated how agents can contribute to the management of VEs.

A promising extension of the approach described in this paper is to couple multiagent techniques with tools from Groupware and CSCW, allowing e.g., to support remote human-to-human negotiation as well as electronic meeting scheduling. Such extensions particularly the part of VE management that will be covered by direct interaction among humans and within that autonomous agents merely adopt assistance functions.

We believe that agent technology has much to offer with respect to the operation of a VE. The support of complex operational functions by autonomous agents (e.g., production planning [Fischer 98]) enables direct horizontal coordination of operational functions that are geographically distributed in a VE. It also supports the handling of unexpected events such as machine breakdown, delays, express orders, i.e., the *exception handling* in workflow systems[3]. Thus, local exceptions can often be treated locally without global replanning, and the necessary coordination processes can be implemented by direct interactions among the agents representing the involved operational functions.

References

[Arnold et al. 95a] O. **Arnold**, W. **Faisst**, M. **Härtling**, and P. **Sieber**. *Virtuelle Unternehmen als Unternehmenstyp der Zukunft?* In: Handbuch der modernen Datenverarbeitung - Theorie und Praxis der Wirtschaftsinformatik, Volume 185, Heidelberg: Hüthig-Verlag, 1995.

[Arnold et al. 95b] O. **Arnold**, W. **Faisst**, M. **Härtling**, and P. **Sieber**. *Virtuelle Unternehmen als Unternehmenstyp der Zukunft?* In: Handbuch der modernen Datenverarbeitung—

[3]Experiences in similar application areas, i.e., transportation scheduling [Fischer et al. 96] and robotics [Müller 96] showed that the ability to deal with dynamics is a substantial property of agent-based systems.

Theorie und Praxis der Wirtschaftsinformatik, volume 185. Heidelberg: Hüthig-Verlag, 1995.

[Bonasso et al. 96] R. P. **Bonasso**, D. **Kortenkamp**, D. P. **Miller**, and M. **Slack**. *Experiences with an Architecture for Intelligent, Reactive Agents.* In: M. Wooldridge, J. P. Müller, and M. Tambe (eds.), Intelligent Agents — Proceedings of the 1995 Workshop on Agent Theories, Architectures, and Languages (ATAL-95), volume 1037: Lecture Notes in Artificial Intelligence, pp. 187–202. Springer-Verlag, 1996.

[Byrne et al. 93] J. A. **Byrne**, R. **Brandt**, and O. **Bort**. *The Virtual Corporation.* Business Week, February 8:36–40, 1993.

[Davidow & Malone 92] W. H. **Davidow** and M. S. **Malone**. *The Virtual Corporation: Structuring and revitalizing the corporation for the 21st century.* New York: Harper Business, 1992.

[Firby 89] R. James **Firby**. *Adaptive Execution in Dynamic Domains.* PhD thesis, Yale University, Computer Science Department, 1989. Also published as Technical Report YALEU/CSD/RR#672.

[Fischer et al. 96] K. **Fischer**, J. P. **Müller**, and M. **Pischel**. *Cooperative Transportation Scheduling: an Application Domain for DAI.* Journal of Applied Artificial Intelligence. Special issue on Intelligent Agents, 10(1), 1996.

[Fischer 98] K. **Fischer**. *An Agent-Based Approach to Holonic Manufacturing Systems.* In: L. M. Camarinha-Matos, H. Afsarmanesh, and V. Marik (eds.), Intelligent Systems for Manufacturing. Kluwer Academic Publishers, 1998.

[Hayes-Roth 95] B. **Hayes-Roth**. *An Architecture for Adaptive Intelligent Systems.* Artificial Intelligence, 72:329–365, 1995.

[Kraus et al. 95] Sarit **Kraus**, Jonathan **Wilkenfeld**, and Gilad **Zlotkin**. *Multiagent Negotiation Under Time Constraints.* In: Artificial Intelligence journal 75(2):297-345, 1995.

[Mowshowitz 86] A. **Mowshowitz**. *Social Dimensions of Office Automation.* Advances in Computers, 25:335–404, 1986.

[Müller 96] J. P. **Müller**. *An Architecture for Dynamically Interacting Agents.* PhD thesis, Universität des Saarlandes, Saarbrücken, 1996.

[Quinn & Hilmer 95] J. B. **Quinn** and F. G. **Hilmer**. *Strategic outsourcing.* The McKinsey Quarterly, 1:48–70, 1995.

[Raulefs 94] P. **Raulefs**. *The Virtual Factory.* In: K. Brannstein and E. Raubold (eds.), 13th World Computer Congress, volume 2, pp. 18–38. North Holland, 1994.

[Visser 98] Volker **Visser**. *Eine Bank auf dem Weg zur Virtualität.* In: Theorie und Praxis der Wirtschaftsinformatik, Volume 200, ISSN 0939-260, 1998.

[Wooldridge & Jennings 95] M. J. **Wooldridge** and N. R. **Jennings** (eds.). *Intelligent Agents – Theories, Architectures, and Languages*, volume 890: Lecture Notes in Artificial Intelligence. Springer-Verlag, 1995.

[Wooldridge et al. 96] M. J. **Wooldridge**, J. P. **Müller**, and M. **Tambe** (eds.). *Intelligent Agents II*, volume 1037: Lecture Notes in Artificial Intelligence. Springer-Verlag, 1996.