Intelligent Agents in Virtual Enterprises

Klaus Fischer, Jörg P. Müller*
German AI Research Centre — DFKI GmbH
Stuhlsatzenhausweg 3, D-66123 Saarbrücken

Ingo Heimig, August-Wilhelm Scheer
Institut für Wirtschaftsinformatik (IWI)
Im Stadtwald, 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 cooperative process of setting up virtual enterprises through the Internet by performing tasks such as presentation, information retrieval and extraction, and the participation in auctions 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 by presenting the main objectives of the new research project AVE (Agents in Virtual Enterprises).

1 Introduction

After visiting a restaurant of an American fast-food chain, Bill Gates decides — inspired by the striking contrast between the colourful, appetising burger adverts and between the taste of the limp, tasteless something he had actually eaten — to use the marketing-power of the gastronomy enterprise for his purposes and determines to initiate a virtual enterprise *(VE) to build the* MacDonalds Computer. *He contacts his secretary, who — supported by her intelligent assistant agent — defines the product requirements by using the VE-management system, and who orders her agent to announce the result in the Internet to interested parties. The next day already — after a short electronic negotiation between potential partners — the enterprises INTEL, IBM, BurgerKing, and Woolworth are chosen as the corporations that fit the goal specification best. The product is quickly renamed* BurgerKing Computer, *the enterprises receive electronic contracts and just one month later BurgerKing starts with the marketing of the product. The computer-supported creation of virtual enterprises in a "plug and play" fashion allows INTEL, IBM, and Microsoft to start the production within one month, whereas employees of BurgerKing design the marketing strategy and Woolworth's sales department prepares for the new product.*

*E-mail: jpm@dfki.uni-sb.de

Certainly, this example is still a vision and some years away from reality; however, it illustrates the immense potential for the computer-supported assembly and operation of virtual enterprises before the background of the rapid development of *Information Highways* and the worldwide cross-linking of information systems.

A notion currently gaining exploding interest is that of *intelligent agents*, i.e., according to [Wooldridge & Jennings 95], 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. Recently, the key role of agent technology in the gigantic information system Internet has been recognised (see [Etzioni & Weld 94], [Maes 94]).

In this paper, we advocate the use of agent technology 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 an electronic virtual enterprise market using the Internet as a vehicle. We describe a new research project AVE (**A**gents in **V**irtual **E**nterprises) the result of which shall be an agent-based methodology and set of tools that enable the human product manager(s) 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. Thus, 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

Currently, the time interval within which a product can be marketed successfully is dramatically decreasing; therefore, enterprises are faced with hard terms of competition. 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. These cooperations enable enterprises to share skills, costs, access to one another's markets and resources and, at the same time, decrease the risk of investments. Presently discussed and designated as the corporate and cooperation model of the future [Byrne et al. 93] is the so-called *virtual enterprise* (VE). (See [Arnold et al. 95] for a distinction between virtual enterprise and other cooperation models.)

2.1 Definition

The term *virtual enterprise* is generally attributed to Mowshowitz [Mowshowitz 86], who established a parallel between it and the term *virtual memory* as it is used in information technology. 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 virtual enterprise as a temporary network of independent companies to share skills, costs and access to each other's market. Throughout this paper, we define a virtual enterprise along the lines of [Arnold et al. 95, p. 10]:

> A virtual enterprise is a cooperation of legally independent enterprises, institutions or individuals which provide a service on the basis of a common un-

derstanding 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]

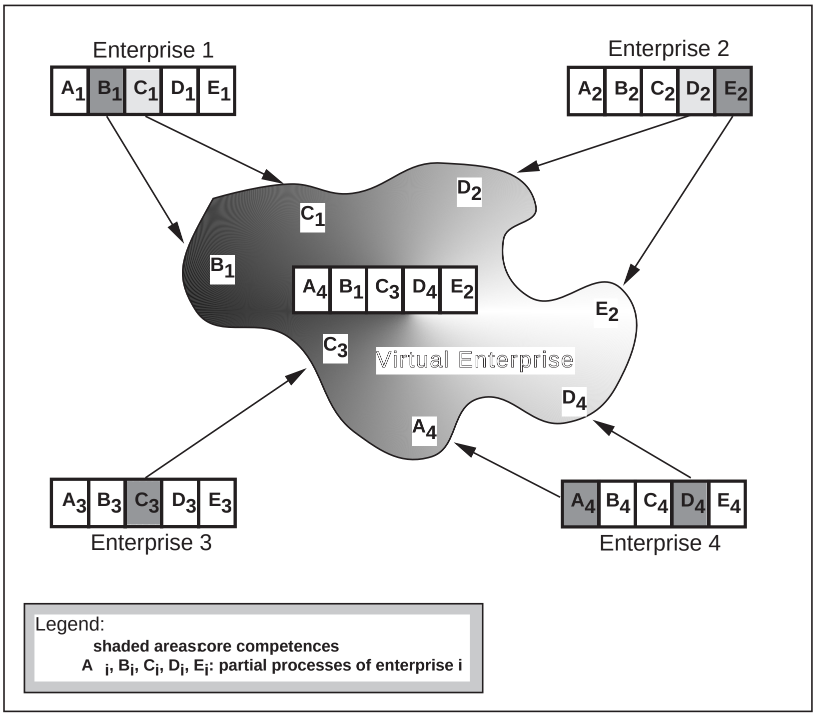

Figure 1: A virtual enterprise

Figure 1 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 1 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.

2.2 Setting up virtual enterprises

Because in a virtual enterprise individual enterprises fulfil only parts of the value chain to produce the product, the main problem in assembling a virtual enterprise is the mapping of the partial processes to the individual enterprises. It can be divided into two phases: In the *first phase*, a suitable business process partitioning (see [Hoffmann et al. 95, pg. 11–13]) is worked out. In the *second phase*, the process chain of the virtual enterprises is assembled and instantiated.

[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.

The problem under consideration in the first phase is at what positions of the value chain such a partitioning is sensible. The solution of this problem is facilitated by a uniform description formalism for partial processes. The ARIS methodology developed by Scheer is such a uniform framework for describing and modeling business processes (see [Scheer 92], [Scheer 94], and Section 6). Partial processes are rated based on *classification figures* that quantify the variety of economic and technical realities. One possibility for evaluating the individual enterprises and their partial processes, respectively, is a process-related model for classification figures [Aichele & Kirsch 95]. It is discussed in Section 6.4.

In the second phase, the partners for the virtual enterprise are selected. It can be divided into four subphases:

- Identification of potential partners.
- Generation of alternative mappings from partners to partial processes.
- Evaluation of the strategic interest and risk, respectively.
- Finalisation of partners and mapping to partial processes.

Our focus in this paper is on the second phase and on the possibilities of how agent technology can support this phase. This support is necessary because today's enterprises often do not know when, where, and in which way to contribute to a virtual enterprise. Existing information markets does not offer sufficiently structured and synoptic information. This applies especially to the World Wide Web (WWW). For an efficient usage of the huge amount of information, tools need to be developed to support enterprises in finding markets and partners, and in the decisions as to founding or joining a virtual enterprise.

The information that is necessary for an optimal execution of the second phase is geographically and organisationally distributed. Moreover, some enterprise-specific information (e.g., classification figures and evaluation functions in the third step) are confidential. Thus, setting up a virtual enterprise is a multiagent decision problem, where each enterprise is an autonomous agent.

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.

3.1 Agents and multiagent systems: definition and applications

Agents are autonomous or semi-autonomous hardware or software systems that perform tasks in complex, dynamically changing environments. Autonomy means the ability to make decisions based on an internal representation of the world, without being controlled by a central instance. Agents communicate with their environment and effect changes in their environment by performing actions.

A multiagent system (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. Researchers have made use of this property to build systems for complex applications like airport management [Rao & Georgeff 95], traffic control and transport logistics [Fischer et al. 96], advanced robotics systems [Bonasso et al. 96], and distributed electricity management [Jennings 94].

3.2 Agent Architectures

In order to cope with these difficult tasks, agents need basic capabilities, such as *reactivity*, *deliberation*, *efficiency*, the ability to *interact* with other agents, and *adaptability*. Over the past few years, the development of control architectures that allow to design agents which can meet these requirements and which can deal with the trade-offs involving e.g., the reconciliation of reactivity and goal-directed reasoning for resource-bounded agents has become an important theme of research. Examples are [Firby 89], [Hayes-Roth 95], and [Bonasso et al. 96]. In the following, we shall briefly explain the agent architecture InteRRaP, which was developed at the DFKI to design autonomous interacting agents (see [Müller & Pischel 94], [Müller 96]), and which will form the basis of the design of agents in the AVE project (see Section 6).

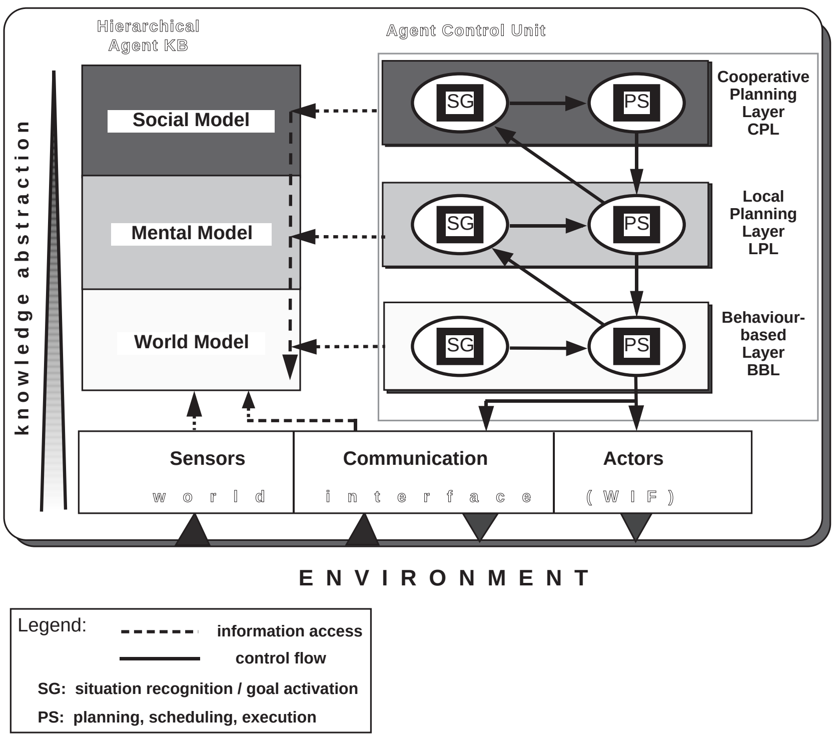

Figure 2: The InteRRaP Agent Architecture

Figure 2 illustrates the structure of the architecture. InteRRaP describes an agent by a world interface, a control unit, and a knowledge base (KB). 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.

4 An Electronic VE Market

Today's global communication networks enable information markets for a goal-directed exchange of information, services, and products. This development is directed towards trade-specific information boards that are maintained in specific Internet nodes, so-called *information servers*. An information server can solve the general problems of *information retrieval* and *information filtering* (extracting the relevant part of information from huge amounts of data) (i) by specific search engines or software agents [Daigle et al. 95] that actively search through the net, or (ii) by allowing clients that intend to provide information in the Internet to register at the server.

In the case of VE management, such servers can be used by individual enterprises to announce and to obtain relevant information. Important functions of an intelligent electronic VE board are the following:

- Collecting goal specifications and offers (autonomous off-line information retrieval/filtering).
- Answering requests on-line by appropriate offers.
- Automatic information of requesters in case earlier requests to the VE board could be successfully matched with new offers.

Figure 3 illustrates the system architecture of an intelligent electronic VE market. The heart of this system is a trade-specific information server that maintains the VE board, and the services of which can be accessed by different companies. These can obtain information from the VE board by commercial standard software. Additionally, there is a higher layer of software support for these activities: the human decision-maker is supported by *intelligent assistant agents* (IAAs) [Maes 94]. In Section 5, we discuss possible areas of application for IAAs and in particular those based on the metaphor of an electronic market.

Note that the assistant agents in Figure 3 are optional: we require that those human users who do not have such an agent (e.g., enterprise B in Figure 3) can also use the services of the VE market. This places some crucial requirement on the design of the system; e.g., the design of communication and negotiation has to take into account the fact that humans may be involved in the process and have to understand the messages. At least, the information servers need to provide standard forms that users can fill out by using commercial WWW browsers.

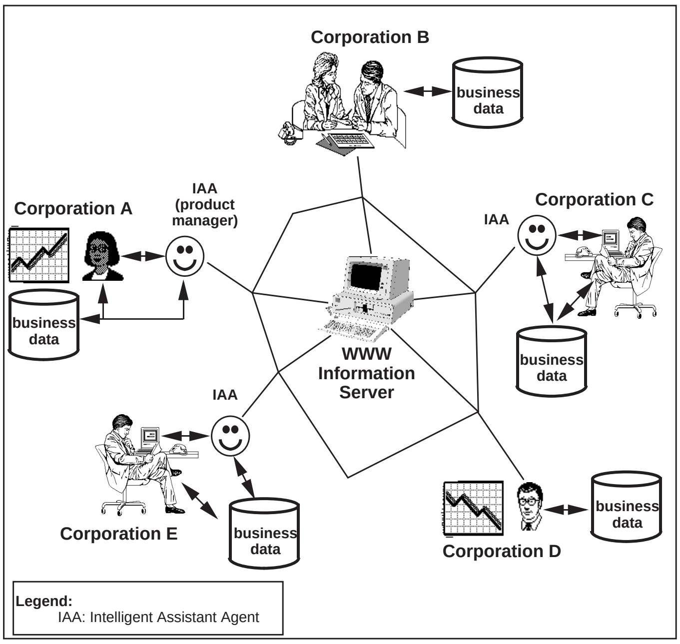

Figure 3: System architecture of an intelligent electronic VE market

An important function in the formation of a virtual enterprise is the *product manager*, who coordinates the selection of partners and the assignment of partners to partial processes. His activities require a high level of coordination and communication with other partners in the virtual enterprise; in particular, it cannot be assumed that all the information required for the decision process is centrally available to the product manager. Thus, his role is that of a coordinator within a decentralised system who is an expert in how to assemble a virtual enterprise rather than that of a central, omniscient instance. In Figure 3, enterprise A takes the role of the product manager. In other examples, this role can be distributed over different partners in the virtual enterprise.

The product manager can be supported by using agent technology. Its main functions can be adopted by IAAs in a (semi-)autonomous fashion: The identification of potential partners and the mapping from partners to partial processes can be implemented by *auction mechanisms* that are initiated by the product manager, who also plays the role of the auctioneer. The participants in the auction recruit themselves from the set of the users who are registered at the server. The auctioneer announces different partial processes of the virtual enterprise to be assembled, and individual enterprises (or: coalitions of enterprises [Shehory & Kraus 95]) can bid for these partial processes; the value of a bid is given by specific classification figures (see Section 6.4).

5 Intelligent Assistant Agents

In [Mullen & Wellman 96], Mullen and Wellman describe three general functions of software agents in virtual market environments: the identification of goods and services (comparable to retrieval, filtering and extraction tasks of assistant agents), the solution of optimisation problems (maximising profits for producers, and maximising utilities for customers), and the management of bidding strategies in negotiation processes. In the VE domain, these functions are extended by additional assistance and presentation tasks. In this section, we focus on the support in setting up virtual enterprises.

5.1 Setting up virtual enterprises

As shown in Section 2.2, the assembly of a virtual enterprise is done in two phases. In the first phase, the product and the partial processes are defined. In the second phases, partners are selected and mapped to partial processes. In the following, the individual problems to be solved in the two phases shall be discussed in more detail.

The main idea of the agent-oriented approach to the formation of virtual enterprises is to look at this process as a *multiagent decision problem*. The distributed solution of this problem is done in four steps:

1. *Goal specification*: description of the global process, i.e., of the product or service to be delivered by the virtual enterprise.
2. *Decomposition*: mapping of the goal specification into partial processes.
3. *Allocation*: selection of potential partners and mapping of partners to the partial processes identified in the decomposition step.
4. *Synthesis*: definition of the global process by the composition of the partial processes instantiated in the allocation step.

The individual steps are discussed in the following.

5.1.1 Goal specification

As illustrated in the introductory example, virtual enterprises are product-driven. Thus, the description of a virtual enterprise is given by a specification of the attributes of the product or service to be delivered. Relevant attributes can be cost, quality, price, markets, quantity of sales, delivery time, as well as constraints on the participants in the enterprise. The goal specification is an expert task that we expect to be very hard to formalise, and thus to fully automate. However, IAAs can support the human product manager in the design of the goal specification by performing assistant tasks: for example, in most cases, the goal specification is based on a cooperation process itself and involves interaction with other enterprises. This cooperation can be supported by CSCW and Groupware tools.

5.1.2 Decomposition

In this step, starting from the goal specification, the relevant partial processes are identified and formalised using a suitable modelling technique, such as ARIS. Most of this task is to be done by humans being assisted by IAAs; however, there is current research on automating (or at least: on giving better support to) these business process (re-)engineering tasks (see [Scheer 94]). Depending on the complexity of the problem and on the structure of the agents involved, the decomposition step can be carried out either in a centralised or in a decentralised way, the latter involving cooperative engineering activities.

5.1.3 Allocation

After the decomposition step, the partial processes need to be distributed to the agents that could potentially contribute to their implementation, and an allocation of partial processes to appropriate parties must be found. This involves three subsequent activities:

- Selection of potential partners.
- Distribution (announcement) of the partial processes to potential partners.
- Selection of partners to perform the individual processes.

The first activity can be supported by the VE board and by information seeking agents that interact with presentation IAAs of individual enterprises. The latter activities define the allocation problem in a strict sense. A centralised solution to the allocation problems described in economics and multiagent systems literature are auction mechanisms. They shall be discussed in more detail in Section 5.2.

Note that the allocation and decomposition steps cannot be assumed to be independent, in general. On the one hand, the given allocation options (available resources) restrict the reasonable decompositions; on the other hand, decomposition decisions clearly restrict the possible allocations.

In the decomposition and allocation step the human product manager can be supported by a *product manager agent* with functionalities including (i) retrieval of relevant information, e.g., existing VE announcements matching the core competences of the enterprise; (ii) formulation of new VE offers; (iii) support of the bidding and evaluation processes; and (iv) autonomous performance of specific negotiation tasks. In the following, we assume that each virtual enterprise has *one* initiator, who also takes the role of the product manager. The election of the product manager is done a priori by the parties involved[3].

The initiator announces goal specifications on the VE board. Interested parties offer specific partial processes, or deliver bids for partial processes matching parts of a goal specification. Information filtering agents can support the human product manager e.g., by directing their attention to interesting offers.

The allocation step can be done in two ways: In *horizontal allocation*, the individual enterprises complement one another in their service offer, i.e., different enterprises cover different partial processes. In *vertical allocation*, a single partial process is covered by a group of enterprises. Vertical allocation is necessary if the individual parties

- cannot completely cover the functionalities necessary to implement an individual partial process;

[3]This process and the interaction processes that it involves can be run in a computer-supported fashion; this, however, shall not be detailed in this paper.

- cannot provide the required capacity, e.g., the amount specified in the goal description;
- need to support each other in order to provide the required quality.

5.1.4 Synthesis

The last step includes the composition of the partial processes that were mapped to appropriate partners during the allocation process to a process that reflects the total added value of the virtual enterprise. A major practical problem in this context is that the partial processes to be connected may run in different geographical and organisatorial environments at different enterprises within the virtual enterprise. Thus, connecting two partial processes P_i and Q_j performed by enterprise i and j, respectively to a global process $\Pi = f(P_i, Q_j)$ in general cannot be obtained by simple composition of the form $\Pi = Q_j \circ P_i$. Rather, it is a complex business process re-engineering task, as both processes must be modified to fit together smoothly. However, it can be substantially simplified by restricting the form of business processes by a uniform description formalism or by predefined reference models [Scheer 94].

5.2 Auction mechanisms in VE formation

A metaphor of increasing importance is that of a collection of agents in a virtual economic system [Mullen & Wellman 96]. This metaphor was successfully used to a variety of distributed resource allocation problems. Auctions (see also [Murnighan 91, pp. 84ff]) are a useful mechanism of interaction for the agents in such a system.

In the context of the formation of virtual enterprises, auctions can be used to implement the allocation step described in Section 5.1. Here, the product manager has the role of the auctioneer, the subject of the auction are individual partial processes, and the bidders are enterprises that are interested in contributing these partial processes to the virtual enterprise.

In the literature, different auction forms were proposed (see e.g., [Murnighan 91]). A distinction is made between open and sealed-bid auctions. In open auction, bids are made public and the bidders know each other's bids. In sealed-bid auctions, the submissions are known only to the auctioneer. The most important open auctions are the English and the Dutch auction. In the case of sealed bid auction, we distinguish between first-price and second-price auctions.

English auction The highest bidder wins and pays the bid it announced. The auction is finished if no new bids are submitted. In some variations, the auctioneer may specify a reservation price, i.e., a lower bound for acceptable bids.

Dutch auction In this auction form, the auctioneer posts the price to pay. In general, the Dutch auction starts with a high price; it is gradually lowered. The first bidder who accepts the price wins.

First-price auction In this sealed-bid auction type, the auctioneer awards the bidder which gave the highest bid; the bidder has to pay the price it bidded. This auction type is an analog of the Dutch auction, since each bidder has to specify the price it is willing to pay, and it actually pays that price in case it wins.

Second-price auction (also called *Vickrey's mechanism* [Vickrey 61]): As in the first-price auction, the highest bid wins; however, the winning bidder has to pay the second-highest bid price. This is the sealed-bid analog to the English auction: a bidder must only beat the next highest bidder to win; therefore, it may be willing to pay more than its last bid.

Theoretical results state that all of these auction forms are efficient and equivalent in a sense that the auctioneer cannot gain any advantage by selecting a specific form. However, experimental analyses showed that, if bidders are risk-averse, the Dutch or first-price auction often generate the highest selling prices. On the other hand, (human) bidders in a second-price auction generally bid higher than in a first-price auction.

The concrete choice of an appropriate auction mechanism depends on the respective market situation. Thus, in a market with a large number of bidder, the auctioneer will often yield the best result by using an open auction form. Moreover, the result of an auction depends on factors such as whether the bidders are risk-neutral, how the value of the object is evaluated, and whether the auctioneer publicises information about the value of the object (which, in general, stimulates the bidding of agents with a low evaluation of the object). Furthermore, data privacy aspects need to be taken into consideration; they can enforce a sealed-bid auction type in case e.g., individual enterprises are not supposed to know each other's bids (given by classification figures, see Section 6.4).

An interesting phenomenon in auctions are *cartels* or bidding rings aimed at avoiding harmful competition. For example, in an English auction, only one member of the cartel takes part in the auction as a bidder. If this bidder wins, the object is re-auctioned in the cartel. The difference between the internal auction price and the winning price in the main auction is distributed among the cartel members (*in-group profit*). The auctioneer can decrease the incentive for the formation of cartels by specifying a sufficiently high reservation price to minimise the potential for in-group profits.

6 Agents in Virtual Enterprises: Project Description

In this section, we describe a new project investigating the potential of applying intelligent agents technology to the problem of the formation of virtual enterprises. The project AVE (Agents in Virtual Enterprises) is a cooperation between DFKI and IWI. In the first project phase, a prototype is developed which allows users located at different sites to carry out a computer-supported auction process, given a goal specification and an initial decomposition, leading to an admissible specification of the global process (the virtual enterprise). Both the auction process itself and the preceding information phase are supported by intelligent assistant agents implemented in JAVA. In this section, we shall give more detailed descriptions of the individual research problems to be tackled at this first project stage.

6.1 From products to processes

The first interesting task is the mapping of a (rather informal) specification of the product or service to be delivered by the virtual enterprise into a (formal) specification of the business processes that describe the virtual enterprise. Like developing a formal program specification from an informal task description, the definition of a process starting from a product description is a complex procedure. For the first project stage, we assume that there is a given description of the global process. Over the past few years, numerous methodologies and formalism have been developed that allow to formally describe the business processes running in an enterprise. We model business processes using the ARIS (Architecture for Integrated Information Systems) methodology developed by Scheer [Scheer 92], [Scheer 94]. ARIS is a framework for integrated business process management and business process re-engineering. ARIS combines methods for modelling, re-engineering and implementing with their corresponding information systems to form an integrated concept. ARIS is supported by a computer-based tool, called the ARIS Toolset[4].

The Architecture of Information Systems (ARIS) makes it possible to analyse business problems of an enterprise by describing its business processes. The first step involves developing a model for business processes. This model is based on two general principles, namely *views* and *descriptive levels*. In order to reduce the complexity of such a model, it is divided and described in individual *views* that represent discrete design aspects and can be handled (to a large extent) independently by using specialised methods. Data and their relations form the *data view*, the functions to be performed and their relationships form the *functional view*, and the structure and relationships between staff members and organisational units constitute the *organisational view*. The *control view* expresses relationships between the different components (views).

Descriptive levels describe the proximity of information systems to information technology. There are three descriptive levels in ARIS, i.e., the requirements definition level, the design specification level, and the implementation description level. The idea is to develop a life cycle concept in order to ensure a consistent description from the business problem all the way to implementation. We model the examples in this paper by using the *control view* and the highest descriptive level, the so-called *requirements definition level*, where business processes are represented by *event-driven process chains* (EPCs). Appendix A illustrates the global business process of the virtual enterprise of the introductory example.

6.2 Decomposition

A further interesting area of research is how a global process chain can be decomposed into suitable parts that can be allocated to individual problem solvers. We start from the observation that the functionalities given by classical enterprise models (e.g., purchase, production, marketing, sales and distribution) are a good way of how to achieve a vertical decomposition of tasks. More complex decomposition methods are subject to future research.

Thus, the global process chain is functionally decomposed into a sequence of subchains. In the computer example, one of these subchains is the planning of sales. Figure 4 illustrates the EPC representation of the partial process *sale planning*. The process starts with the event *contract signed*. The first activity is the function *create sales forecast*, which has as input data provided from BurgerKing (marketing data), Woolworth (sales know-how), and additional information providers (e.g., statistical market data). The output of the function is a sales forecast. Upon creation of the sales forecast, the function

[4]© by IDS Prof. Scheer GmbH

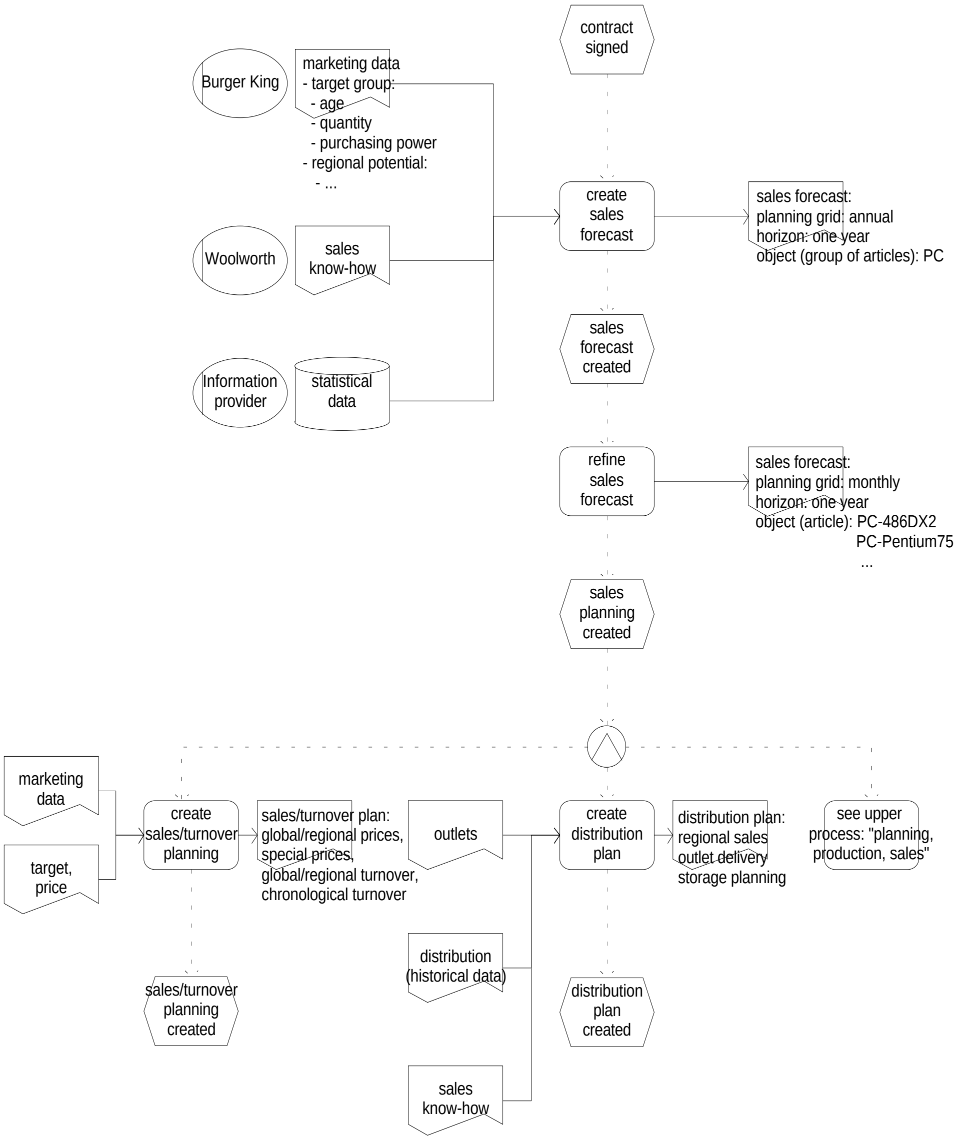

Figure 4: The partial process *sales planning*

refine sales forecast is performed, before the process chain branches into a conjunction. From the example, we see the basic modelling entities: events (diamonds), functions (soft rectangles), organisational units (circles), and data (sheets and database symbols).

Note that the partial processes generated in the decomposition phase are generic reference models that are not yet fully instantiated. The instantiation of processes, i.e., the mapping of functions to organisational units, is computed in the allocation phase described below.

6.3 Modelling agents and agent interaction

The main focus of the AVE project shall be on the functionalities of IAAs in the processes centred around the allocation step described in Section 5.1. The IAAs that assist the human product manager and the enterprises interested in setting up a virtual enterprise act in a dynamic and complex virtual environment. The agents in AVE are designed according to the InteRRaP agent model described in Section 3. The *behaviour-based layer* of an agent contains its reactive skills, e.g., reactor patterns of behaviour that are used to monitor the net for new relevant information, or to notice the attempt of another agent to initiate a (bidirectional) communication process. The agent's *local planning layer* contains descriptions of the local EPCs that represent abstract descriptions of partial processes that the agent's enterprise is able to carry out. This layer has access to reference models and to relevant classification figures. It is able to evaluate instantiated EPCs with respect to the reference models. These local evaluations form the basis for the decisions at the *cooperative planning layer*, where the knowledge is located that enables the agent to participate in auctions. Based on the general negotiation model proposed by Rosenschein and Zlotkin in [Rosenschein & Zlotkin 94], this knowledge includes:

- The generation of the negotiation set, i.e., the set of admissible results.
- The negotiation protocol that characterises external aspects of the auction, i.e., the admissible (possibly role-specific) interactions of the individual participants).
- The negotiation strategy that determines how the agents make their decisions in the auction process

In AVE, different protocols corresponding to the different auction forms described in Section 5.2 as well as different strategies (e.g., as regards the agents' risk-taking behaviour) shall be implemented and experimentally evaluated. For a more detailed description of the individual layers of the InteRRaP model and of the cooperation and negotiation model underlying the architecture, we refer to [Müller 96].

6.4 Modelling and evaluation of processes

Each partial process within the global process chain is instantiated by an auction process in the allocation phase. The instantiation is to a high degree based on the evaluation of alternative modelling alternatives. The development of suitable utility functions that enable individual agents to evaluate the quality of a candidate partial process is a challenging research task, as it includes the development of classification figures that allow reliable quantifications of the quality of processes based on a formal process description. As a result of the first project phase, we expect a restricted model that, for a specific domain, allows to make decisions based on a limited set of classification figures [Aichele & Kirsch 95]. They connect the economic and the technical view of a process. In particular, financial classification figures have to be included like sales profitability or indebtedness as well as non-financial classification figures like market share, grade of service, throughput time per order and set-up time per operation.

Table 1 shows an exemplary set of classification figures for the partial process *sales planning*[5]. The evaluation of the process depends on quantitative economic factors like

[5] This example is due to Christian Aichele, IDS Prof. Scheer GmbH, personal communication, 1996.

the duration of the partial process and its share in the total product costs, but also on qualitative factors such as the quality of customer service. The relevant classification figures are determined by the product manager. Thus, a bid for the partial process consists of a seven-tuple with a sub-bid for each figure. The product manager then computes the best bid as a weighted combination of the sub-bids. To focus the bidding process, the product manager may specify a reservation price specifying upper or lower bounds for acceptable bids, respectively. The rightmost column of Table 1 shows examples for reservation prices.

	Classification figures	**Reservation price**
I	duration of process	$\leq$ 2 weeks
II	share of total product costs	$<$ 1%
III	number of outlets (region specific)	EU: $>$ 100 Switzerland: $>$ 20 Austria: $>$ 10 USA: $>$ 500 ...
IV	deviation of planned sales from actual sales	$<$ 5%
V	dispatching quota (per outlet) = $\frac{\#\ product\ available}{total\#\ customerorders(outlet)}$	$>$ 99%
VI	service quota (per outlet) = $\frac{\#\ answered\ customer\ requests}{total\#\ customer\ requests(outlet)}$	$>$ 99%
VII	# customer requests per outlet employee	$<$ 10 for numerator $>$ 100

Table 1: Classification figures for the partial process *sales planning*

Thus, in contrast to most negotiation protocols considered so far in multiagent systems research, the selection of partial processes in the VE domain is a multi-attribute decision problem (see [Keeney & Raiffa 76]). Therefore, the task of the auctioneer is much more complex than e.g., in the contract net protocol [Davis & Smith 83], where the selection of the best bid is a simple arithmetic process.

6.5 Competitive agents in a virtual market

An additional problem in virtual markets is that the agents that participate in these markets are self-interested local utility maximisers by nature. As there is no central designer of such a system, but rather many local designers of individual IAAs, it is obvious that an enterprise may have an incentive to program their agents such that they will achieve maximal profit for the enterprise. This, however, implies that, whenever this is beneficial, agents may lie (e.g., in auctions). Recently, there have been approaches to apply game-theoretic results to rational agents. The application of these results in the virtual enterprise domain are a further interesting research task in the AVE product. As an example, the second-price sealed-bid auction ensures that all agents bid honestly (see [Vickrey 61] [Ephrati & Rosenschein 95]).

A second problem that has to be tackled is that of *commitment*. E.g., an IAA that is involved in an auction to set up a virtual enterprise may bid for a specific partial process, e.g., for producing 20.000 units of a product a year. Clearly, there is a commitment behind such a decision, and winning the bid implies committing oneself to the content of the bid.

Competences of and commitments among autonomous software agents is a challenging research area that offers a variety of theoretical, practical, and—last but not least—legal issues (see e.g., [Krogh 96]) to be solved within the next years.

7 Conclusions

The increasing importance of virtual enterprises 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 formation of virtual enterprises. A new project was described which aims to investigate the potential of agent technology for this application range.

A promising extension of the approach described in this paper is to couple multi-agent 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.

The focus of this paper was on the use of agents in setting up virtual enterprises. A further important area of usage for this technology is the *operation* of virtual enterprises, i.e., the active support of workflow management. In [Raulefs 94], an approach based on the model of a *Virtual Factory* is presented, and software modules for the control of a corresponding *Physical Factory* are described.

We believe that agent technology has much to offer with respect to the operation of a virtual enterprise. The modelling of complex operational functions as autonomous agents (e.g., production planning [Fischer 93]) enables direct horizontal coordination of operational functions that are geographically distributed in a virtual enterprise. It also supports the handling of unexpected events such as machine breakdown, delays, express orders, i.e., the *exception handling* in workflow systems[6]. 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. The investigation of intelligent agents in the operation of virtual enterprises shall be the second phase of the AVE project.

In summary, the vision of tomorrow's VE management system not only includes the support of setting up a virtual enterprise, but also its active operational control.

References

[Aichele & Kirsch 95] C. **Aichele** and J. **Kirsch**. *Geschäftsprozeßanalyse auf Basis von Kennzahlensystemen.* Management & Computer, 3, 1995.

[Arnold et al. 95] 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.

[6]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.

[Bonasso et al. 96] R. P. **Bonasso**, D. **Kortenkamp**, D. P. **Miller**, and M. **Slack**. *Experiences with an Architecture for Intelligent, Reactive Agents.* In: [Wooldridge et al. 96], 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.

[Daigle et al. 95] L. **Daigle**, P. **Deutsch**, B. **Heelan**, C. **Alpaugh**, and M. **MacLachlan**. *Uniform Resource Agents.* In: Proceedings of the Intelligent Information Agents Workshop at 4th International Conference on Information and Knowledge Management (CIKM'95), Baltimore, Maryland, USA, 1995.

[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.

[Davis & Smith 83] R. **Davis** and R. G. **Smith**. *Negotiation as a metaphor for distributed problem solving.* Artificial Intelligence, 20:63 – 109, 1983.

[Ephrati & Rosenschein 95] E. **Ephrati** and J. **Rosenschein**. *A Framework for the Interleaving of Execution and Planning for Dynamic Tasks by Multiple Agents.* In: C. Castelfranchi and J.-P. Müller (eds.), From Reaction to Cognition — 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'93), volume 957: Lecture Notes in Artificial Intelligence, pp. 139–153. Springer-Verlag, 1995.

[Etzioni & Weld 94] O. **Etzioni** and D. **Weld**. *A Softbot-Based Interface to the Internet.* Communications of the ACM, 37(7):72–76, 1994.

[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 93] K. **Fischer**. *Verteiltes und kooperatives Planen in einer flexiblen Fertigungsumgebung (Distributed and Cooperative Planning in the Environment of a Flexible Manufacturing System).* DISKI, Dissertationen zur Künstlichen Intelligenz. infix, 1993.

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

[Hoffmann et al. 95] W. **Hoffmann**, A.-W. **Scheer**, and C. **Hanebeck**. *Geschäftsprozeßmanagement in Virtuellen Unternehmen.* Veröffentlichungen, Institut für Wirtschaftsinformatik, 1995.

[Jennings 94] N. R. **Jennings**. *Cooperation in Industrial Multi-agent Systems*, volume 43: World Scientific Series in Computer Science. World Scientific Publishing Co. Inc., 1994. ISBN: 981-02-1652-1.

[Keeney & Raiffa 76] R. L. **Keeney** and H. **Raiffa**. *Decisions With Multiple Objectives.* John Wiley and Sons, New York, 1976.

[Krogh 96] C. **Krogh**. *The Rights of Agents.* In: [Wooldridge et al. 96], pp. 1–16. Springer-Verlag, 1996.

[Maes 94] P. **Maes**. *Agents that reduce work and information overload.* Communications of the ACM, 37(3):31–40, 1994.

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

[Mullen & Wellman 96] T. **Mullen** and M. P. **Wellman**. *Some Issues in the Design of Market-Oriented Agents.* In: [Wooldridge et al. 96], pp. 283–298. Springer-Verlag, 1996.

[Müller & Pischel 94] J. P. **Müller** and M. **Pischel**. *An Architecture for Dynamically Interacting Agents.* International Journal of Intelligent and Cooperative Information Systems (IJICIS), 3(1):25–45, 1994.

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

[Murnighan 91] J. K. **Murnighan**. *The Dynamics of Bargaining Games.* Englewood Cliffs, NJ: Prentice Hall, 1991.

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

[Rao & Georgeff 95] A. S. **Rao** and M. P. **Georgeff**. *BDI-Agents: from theory to practice.* In: Proceedings of the First Intl. Conference on Multiagent Systems, San Francisco, 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.

[Rosenschein & Zlotkin 94] J. S. **Rosenschein** and G. **Zlotkin**. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* MIT Press, 1994.

[Scheer 92] A.-W. **Scheer**. *Architecture of Integrated Information Systems - Foundations of Enterprise Modelling.* Springer-Verlag, Berlin u.a., 1992.

[Scheer 94] A.-W. **Scheer**. *Business Process Engineering - Reference Models for Industrial Enterprises.* Springer-Verlag, Berlin u.a., 2nd edition, 1994.

[Shehory & Kraus 95] O. **Shehory** and S. **Kraus**. *Coalition formation among autonomous agents: strategies and complexity.* In: C. Castelfranchi and J.-P. Müller (eds.), From Reaction to Cognition — 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'93), Lecture Notes in Artificial Intelligence, pp. 56–72. Springer-Verlag, 1995.

[Vickrey 61] W. **Vickrey**. *Counterspeculation, auctions and competitive sealed tenders.* Journal of Finance, 16:8–37, 1961.

[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.

Figure 5: Global process chain in ARIS notation (Appendix A)