The Living Value Chain

Coordinating Business Processes
with Artificial Life Agents[1]

Torsten Eymann, Detlef Schoder, Boris Padovan

Albert-Ludwigs-Universität Freiburg
Institut für Informatik und Gesellschaft, Abteilung Telematik
Friedrichstrasse 50
79098 Freiburg im Breisgau, Germany
{eymann|schoder|padovan}@iig.uni-freiburg.de

Abstract

The development of IT systems for coordinating business processes mostly starts with the definition of coordination mechanisms and organizational structures. This is usually based on assumptions about the cooperational behavior of the participating individuals. This top-down approach fails if the human participants behave in a competitive way, which can severely affect coordination as a whole.

Artificial life concepts maintain global coordination by modeling local interaction rules and building on competition between individuals. This basic idea is exploited here in an economic context to show the possibility of coordinating business processes, like value chains, by autonomous self-interested software agents. We set up an artificial world where agents, representing human participants, cooperate and communicate using economically interpretable protocols. Starting from a given initial state of an economic agent population, the development of this population by applying local agent interaction rules and the observable coordination patterns will be investigated. We are currently in progress of implementing the simulation named *Avalanche*.

1 Human Behavior and the Coordination of Business Processes

Internal and external communication and cooperation processes are supported by more and more efficient IT applications. The main research field addressing business process coordination is CSCW (Computer-Supported Cooperative Work). Applications derived from this research include e.g. Groupware like Lotus Notes, Group Decision Support Systems and Bulletin Boards. The functionality of these systems still grows with every update. In parallel the usage of these applications has substantially changed from yesterday's automation of workflow processes (well-defined and repeating) to today's support of cooperative and coordinative processes, which may occur on a solitary and ad-hoc basis, thus posing a problem when benevolent cooperation between participants can not be presupposed.

[1] in: Nwana, Hyacinth S., Ndumu, Divine T. (Eds.): Proceedings of the 3rd Intl. Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'98), The Practical Application Company: Blackpool, UK, 1998. ISBN 0 95255 5484. p. 111-122

If conflict and competition exist, the introduction of any CSCW application faces a hard task and in most cases leads to non-acceptance. In previous research work [Hummel 1996] we have shown that this is not only caused by technical shortcomings, but also by application misunderstandings. A clear example of such a failure is given by the demise of shared calendars: information-strategic behavior leads to people not putting all of their appointments in the calendar which results in the uselessness of the application as a whole. This observation holds regardless of the formal governance structure of the process, be it completely inside a large company or in a virtual enterprise just created from linking previously separated processes together. Developing agent-based solutions to a coordination problem will share these experiences, when human participants are to be represented by benevolent software agents, thus ruling out a large part of real human behavior.

In order to support these difficult requirements we need to think about IT applications which are fault tolerant against human behavior in that conflict and competition are neither ignored nor artificially suppressed, but where the energy of the participants is converted to a driving force for the success of the application as a whole. In other words, a clever application architecture should use the local optimization effort of every participant to enhance the overall profit.

The development of agent-based IT systems can in principle follow two paradigms: "(cooperative) distributed problem solving" (DPS) or "multi-agent systems" (MAS) [Durfee/Rosenschein 1994]. In short, DPS is a top-down approach while MAS is bottom-up with regard to the strategies the agents are to pursue. DPS begins with the coordination problem to be solved and imposes an interaction (cooperation) protocol and a strategy to use for each agent. In MAS, each agent may be designed independently, using the same protocol but pursuing different strategies. The coordination problem is not explicitly addressed, and the solution is expected to emerge from the local interactions of the agents.

A common CSCW-like top-down DPS approach would start with the coordination mechanism, and then describe the necessary and allowed individual interaction rules and the ways to transmit information from one system unit to the other. This is feasible for an environment where the behavior of the participants (or the agents representing the participants) can be enforced by organizational measures, but not for an open network without binding rules (like the Internet) where self-interested agents and humans prevail. An agent-based solution for coordinating value chains will therefore have to use the MAS approach than DPS.

Coordination in this context can be described as "the act of managing interdependencies between activities" [Malone/Crowston 1990], given either a resource to be shared or a timing interdependency. The coordination mechanism can either be mediated (centralized, e.g. auctions) or unmediated with the use of subsequent bilateral or multilateral cooperation between all agents [Wellman 1996]. Examples for an unmediated bilateral mechanism in a value chain context are concepts for the coordination of manufacturing processes (e.g. KANBAN), where produced goods are sequentially and independently forwarded between autonomous workgroups as input factors for the next value chain step.

Accordingly, *Cooperation* is viewed as the bilateral dissolution between discrete agents of what we observe as a coordination mechanism and that determines the behavior of those units that are part of an organizational structure. A classical hierarchical organization will have a distribution of power that increases to the top, which may at the extreme turn into a cooperational behavior of command and obedience like in a military context. A cooperational manner of negotiation between two agents would not fit in this context, but works perfectly in a market-like coordination setting. Cooperation in a wider sense not only covers benevolent interaction, but also tricking, cheating and open conflict.

In this chapter we described the requirements that the IT support of business processes has to face and illustrated different approaches (top-down or bottom-up) of designing agent-based IT support for business processes. Next we discuss the use of Artificial Life (ALife) for simulating economic processes. The focus of our current research, the bottom-up development of a multi agent system which decentrally coordinates a multiple step value chain by applying simple local interaction rules for the cooperation and communication of the agents, comes next. The paper concludes with an outlook to apply the concept of our Artificial World as a first step towards modeling real value chains and why Multi-Agent Systems (MAS) especially are a prime technology to apply.

2 Artificial Life Agents and the Simulation of Economic Processes

We regard our project in the research context of agent-based computational economics (ACE), which is roughly characterized as the computational study of economies modeled as evolving decentralized systems of autonomous interacting agents [Tesfatsion 1997]. A central concern of ACE research is then to understand the apparently spontaneous appearance of global regularities in economic processes, such as the unplanned coordination of trade in decentralized market economies that economists associate with Adam Smith's "invisible hand". The challenge is to explain these global regularities from the bottom up, in the sense that the regularities arise from the local interactions of autonomous agents channeled through actual or potential economic institutions rather than through fictitious top-down coordination mechanisms such as a single representative consumer. ACE is thus a specialization to economics of the basic Artificial Life paradigm [Tesfatsion 1997].

The essence of Artificial Life (ALife), as described by Chris Langton, is "the realization of lifelike behavior on the part of man-made systems consisting of populations of semi-autonomous entities whose local interactions with one another are governed by a set of simple rules. Such systems contain no rules for the behavior of the population at the global level, and the often complex, high-level dynamics and structures are emergent properties, which develop over time from out of all of the local interactions among low-level primitives by a process highly reminiscent of embryological development, in which local hierarchies of higher-order structures develop and compete with one another for support among the low-level entities. These emergent structures play a vital role in organizing the behavior of the lowest-level entities by establishing the context within which those entities invoke their local rules and, as a consequence, these structures may evolve in time." [Langton 1989]

Our value chain (see below) can be described as a man-made system with a population of entities (organizational units) with local interaction protocols (simple rules for buying and selling goods). Higher-order structures, which appear to the eye of the beholder as markets and hierarchies, then are considered to evolve. Seizing a suggestion of Jennings *et al.* [Jennings *et al.* 1996], we thus "view the business process as a collection of autonomous, problem-solving agents which interact when they have interdependencies".

Avalanche[2] can also be regarded as an implementation of an Artificial World (AW) [Lane 1993]. AWs are computer-implementable stochastic models which consist of a set of agents that interact with each other and an environment in prescribed ways. They are designed to give insight into processes of emergent organization and coordination: "What makes a set of economic agents organize into an economy?" [Lane 1993].Examples for existing Artificial Worlds are Tierra [Ray 1992] and Sugarscape [Epstein/Axtell 1996].

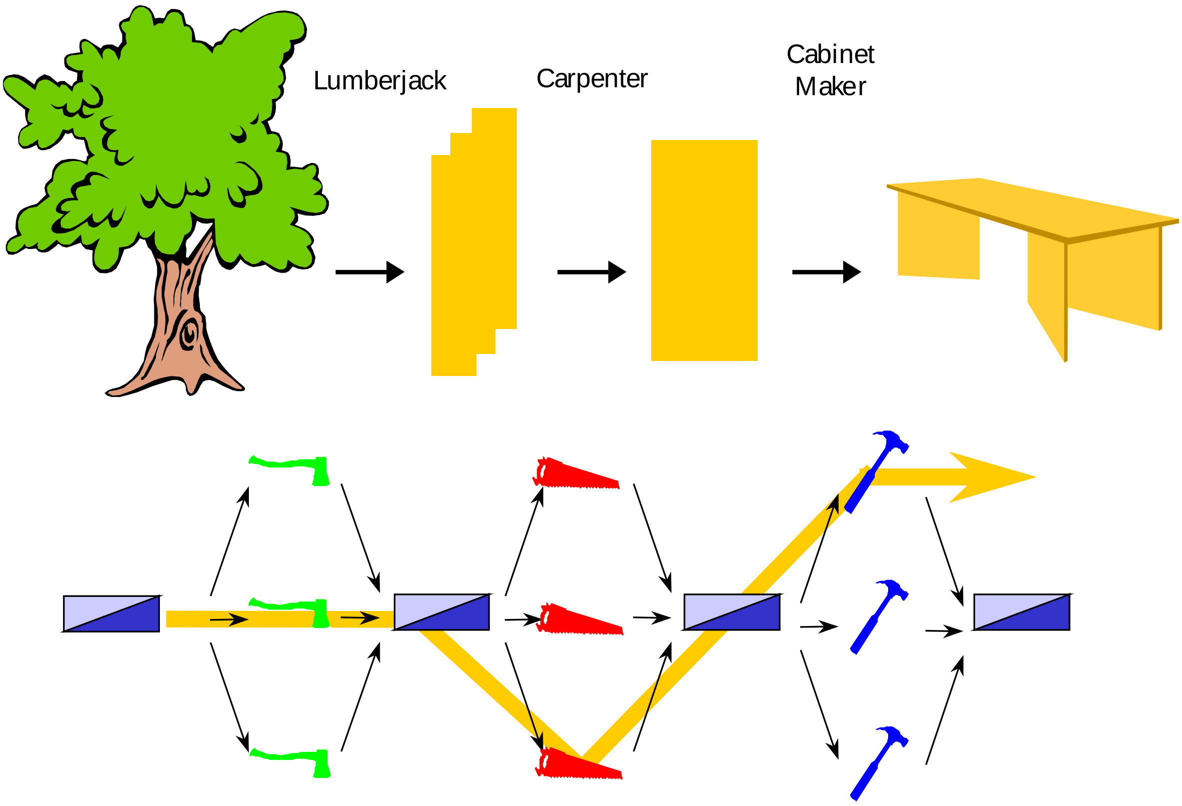

Figure 1: An example of a typical value chain from tree to desk

To reflect a real world setting, our Artificial World consists of several electronic marketplaces, between which mobile agents, representing the participants lumberjacks, carpenters and cabinet makers, roam in search for respective offers or demand. Typically, a value chain consists of several interconnected organizational units [Porter/Millar 1985], which use raw materials as input and by adding labor and knowledge enhance the value of their manufactured product. This output is then taken as input by the next organizational unit in line until the product reaches the buyer. This value chain can be extended from the

[2] ***Avalanche*** is a subsidiary project of the TELOS (Telematics, Coordination Mechanisms and Organisational Structures) project, which is conducted at the Institute for Computer Science and Social Studies in Freiburg. Research topics of TELOS are the interdependencies between the use of Telematics systems in organisations, the development of the organisational structure and the applied coordination mechanisms. The impact of the use of computers, especially networks and CSCW (computer supported cooperative work) systems is important research of the past which we build upon [Schoder 1995, Hummel 1996, Schoder *et al*. 1997].

simplest raw materials (wood, oil, gas) to more complicated products (cars, houses), even if the description resembles more of a network than a linear chain. But for any particular material-product relation it is possible to describe a linear chain, as modeled in the picture for trees and desks.

Considering this, we view the different professions (agents) in the value chain as typical "low-level primitives" in an ALife sense, semi-autonomous agents that are driven by local interaction rules. We can distinguish between agents which chop trees to produce boards (lumberjacks), those that use boards to produce plates by fixing the boards together (carpenters) and finally those that buy plates and build a fine table out of five plates (cabinet makers). The different types of agents are denoted by different icons in this figure. The lumberjack does not know about the desk finally made of the wood he cuts, neither does the cabinet maker know about the particular tree. Nevertheless, all agents produce at best, from a consumer's point of view, a probably high quality desk made of high quality wood - all by local interaction, in sequence of the distinct value chain steps.

The organizational structure and the coordination mechanisms applied by the different professions is not addressed in this picture, as the final product is always the same: a desk, whether all workers along the chain are independent firms (market-like coordination) or part of one large company (hierarchical coordination). On a continuum between market and hierarchy we find a world of discrete coordination mechanisms where agents interact and which may emerge, if attracted, to appear as parts of distinct organizational structures [Kauffman 1995].

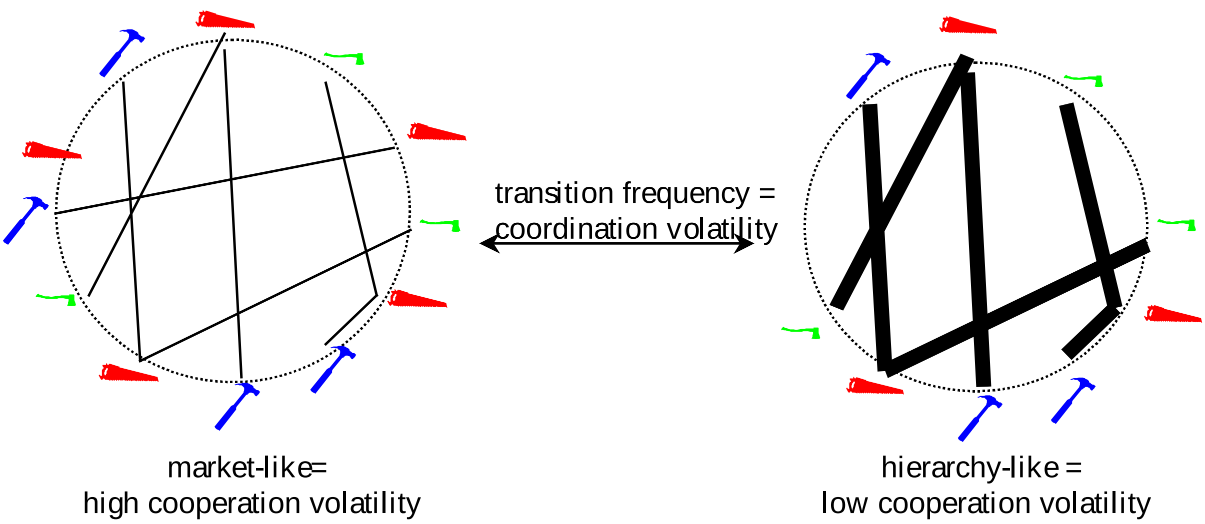

Figure 2: Visualization and Implementation of the Artificial World

Trade relations between the agents are visualized by connecting lines. A thin line between for example a "carpenter agent" and a "cabinet maker agent" means a cooperational relationship where the cabinet maker has as his last action successfully purchased some boards from the carpenter. If his next purchasing action will again involve the same carpenter agent, the line is drawn thicker to show the repetition. This allows us to give a measure for the decisive aspect of change regarding the interaction of the cooperation partners: *cooperation volatility* measures the frequency of repeating an observed cooperational relationship and can be easily measured, aggregated and displayed. The variation speed of the medium cooperation volatility (the *coordination volatility*) can, as transition frequency between two relatively stable states, also easily be measured. Our interpretation matches frequent reconfiguration to a (relatively chaotic) market, while the more stable structure of low cooperation volatility resembles a fixed organization like with dominating hierarchical relations.

3 *Avalanche* - An Implementation of an Economic Artificial World

Our implementation *Avalanche* of an economic multi-agent system will be realized in Java with the basic components of locations (trading places) and agents (suppliers and buyers). We use standardized concepts of Java, mobile agent class libraries (e.g. Mole [Baumann 1996], Voyager [Objectspace 1997], Odyssey [White 1997]), and add very few methods and variables considered to be necessary in an economic sense.

Locations are Java Virtual Machines (VM), server processes which run statically on a given computer. Starting the model world means the same as initializing the locations. Locations are addressed and distinguished by a unique port number on a TCP/IP-network connected computer (e.g. 127.0.0.1:7777), which makes it possible to run several locations on one computer for modeling purposes. Locations are a sandbox environment for the agents and offer a registration list to see which other agents are on the same location, but never actively influence or coordinate the agents' activities.

Agents are also distinguishable by a unique ID. The agents communicate unmediated and bilateral by calling methods or passing a message. The negotiation protocol used will be largely taken from FIPA's Agent Communication Language [FIPA 1997], implementing contract nets [Smith 1980], and modified with elements suggested by [Eriksson *et al*. 1997] and [Sandholm 1996] which reflect the self-interest of the agents. The agents run concurrently and parallel, an influence on the results through the application of a sequential scheduling algorithm is therefore avoided. Not based on rounds, but synchronized by `heartbeat` ticks, every agent will pro-actively execute actions with varying execution time.

The strategy of the agent for the use of the protocol is stochastically determined by parameters like acquisitiveness (the likelihood to change an offer price as supplier or a demand price as buyer), price-prospect (the likelihood to expect a changing price from the supplier or buyer), satisfaction (the likelihood not to request an additional offer from a supplier) or impatience (the likelihood to change the location if not confident with the market situation). Any action of the agent is then specified by "rolling the dice" against these and only a few other parameters and can not be computed in advance. Other, non-stochastical parameters are for example the agent's ID, the production function and the capital, which is not only a means of computing and storing "money" units when buying or selling, but also an indicator of the relative success of the single agent. An agent which is faster and/or fitter in trading as another will have a relatively higher capital account.

The mobility of the agents allows them to roam the net in search for the most suitable offer or demand situation. Success in this task leads to a reward by allowing to reproduce the agent's "genes" (variables and methods) using genetic algorithms, which dynamically leads in time to an agent population which is best adapted to the current supply and demand situation. Changes in this situation will be accompanied by changes in the agent population without having to exogenously interfere.

An example for the effect of the stochastic operations works as follows: the supplier agent announces the availability of finished goods on its behalf on the registration list of the current location. The current price for his products is revealed only on a buyer's request.

Since it is unlikely that the products can be sold off at once, the offer may last several `heartbeat` ticks. To model economic offer behavior with respect to the self-interest of the agents, we use the heuristic concept of Kasbah [Chavez/Maes 1996] and change it to follow a probabilistic decay function, checking with every heartbeat against the supplier's "acquisitiveness" parameter. If the check fails, the supplier will lower its offer price in the next phase, thus making a successful selling act more and more probable (regardless of the buyer's actions). If the goods can not be sold at all (e.g. if no agent of the next value chain step is alive or in reach), the offered price can even fall down to zero.

Those agents which appear as buyers (because they have neither products nor materials) will first have a look on the registration list of the location they currently reside on. After that, the buyer then communicates subsequently with the identified supplier agents, realizing a local market survey (checking against the satisfaction parameter). The most simple preference function now ranks the suppliers by unit price; this list is then used by the buyer to purchase materials until his demand is matched. Since the agents run concurrently, it is intentionally possible that not all materials are available anymore at the moment when the agent finally decides to buy, in which case the buyer agent would have to switch to an inferior supplier or start the whole procedure again.

The use of Genetic Algorithms (GAs) in this economic simulation is motivated by the need to identify those parameter values which lead to an agent's success or demise [Holland/Miller 1991]. The success of a single agent is relatively easy indicated by a high capital account compared to other agents. The success of the population as a whole can be measured in economic (Pareto efficiency, social welfare) or technical terms (computational or communicational efficiency) [Sandholm 1996]. In our value chain environment, we consider an agent population to be superior to another when the productivity per time unit is higher.

Starting with an initial population of more or less randomly initiated agents the population develops by ruling out inferior and promoting superior agents by applying evolutionary mechanisms. Literature on genetic algorithms [Goldberg 1989] distinguishes four steps (evaluation, reproduction selection, recombination and mutation), which are implemented in our economic artificial world as follows.

In the *evaluation* phase each agent is judged by a fitness score, which is represented by the capital account of the agent in our simulation. At the beginning of the simulation, each agent is equipped with the same amount of money. All purchase and selling actions have direct consequences on the capital - on one hand purchased goods are paid from the capital, thus diminishing it, on the other hand the income from a sale leads to an increase of capital. In any `heartbeat` the subtraction of money from his capital is inevitable, reflecting costs of storage, CPU time and telecommunication use. If the capital reaches zero, the agent is considered to have died and will be removed from the model world.

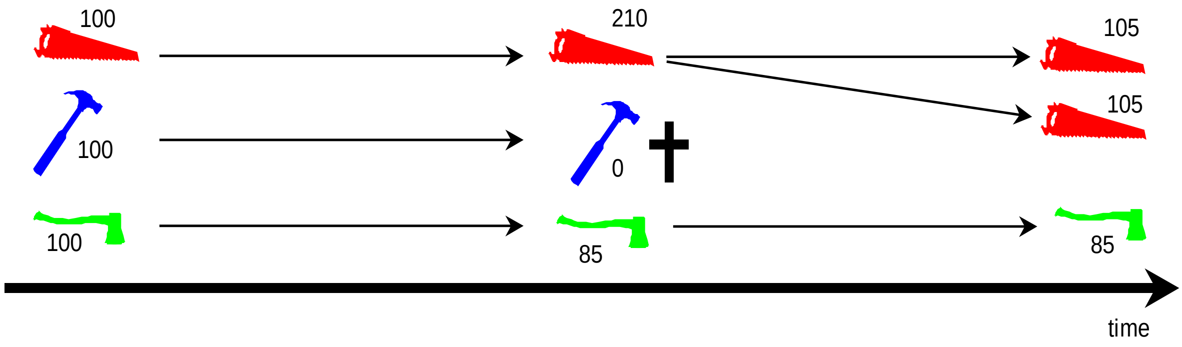

Figure 3: Fitness evaluation of agents using their capital account as score

The r*eproduction selection* phase selects by means of the fitness score those agents for reproduction whose current capital reaches a certain factor (e.g. twice) of his initial capital, reflecting the relative success. The next step of *recombination* is usually done by taking attributes (genes) from paired agents of the same value chain step and having the "children" implement a mixture of the attributes of these parent agents. In *Avalanche*, the attributes are the stochastic parameters referred to above. The fourth and last phase is *mutation*, where the parameter values of the new agents are slightly and randomly modified.

The final two phases (recombination and mutation) will not be implemented in *Avalanche*; to keep things simple, genetic diversity is brought in by the addition of randomly attributed agents anywhere in the model world in equidistant periods. Therefore, a certain amount of agents will always exist in the simulation which are diverse from those currently "employed" and working and which can grasp the evolutionary chance of changing the system in their favor. This also offers the possibility to add new production functions (reflecting technological change) to the population and see how they perform. Instead of pairing, reproduction can then use a simple biological cleavage algorithm - the parent agent may inherit some capital to a perfect clone (by means of attribute values) of itself.

The use of Genetic Algorithms (GA) in economic MAS simulations is controversial [Hillebrand/Stender 1994], since it lacks economic theory. In contrast to living beings, economic institutions and organizational units possess a theoretically unlimited lifetime and do usually not reproduce or generate offspring in a biological sense. In our simulation we use the following economical argumentation for removing from and adding agents to the population: the removal of unsuccessful agents by means of zero capital is interpreted as bankruptcy. The sudden appearance of clones of successful agents can be rooted in microeconomic theory with the argumentation that obvious success of one firm in a market leads to entries of new firms who then adapt to the observed successful behavior [Varian 1993].

4 Conclusion and Outlook

Since the real-time implementation and the spatial distribution of the Artificial World make it impossible for supplier and buyer agents to get a full overview of any current market situation, this research project has many connections to the work of Sandholm concerning automated negotiation among self-interested computationally limited agents [Sandholm 1996]. We also benefit from the market-oriented programming approach of Wellman [Wellman 1996], especially in the design of our goods and mechanism space.

But in contrast to the works of Sandholm, Wellman and Jennings [Jennings 1996], our agents are neither intelligent nor do they act strategically. Their actions are based on stochastic operations on the initial and never changing parameters. The results of the actions do not have any effects on the agents' parameters nor will any change in the behavior occur. Any "strategy" observed is just an ascription to a (within limits) random behavior. Our agents are in effect 0-level agents [Vidal/Durfee 1997], which do not reason directly about the other agents. Compared to the economic environment the agents live in, they are no more intelligent than a single ant compared to the antheap it is a part of.

The use of Genetic Algorithms makes it possible to start the simulation with a non-optimal population of agents, each representing a more or less random distribution of parameter values. Ruling out the unsuccessful agents and promoting the successful, the population gains economic strength over time, and clusters of successful parameter combinations appear which are effective in the dynamic environment the agents live in. Because of the dynamic and non-linear progression of a simulation run agents can also emerge and rule whose attributes are optimal in the environment given even if the initial population did not exploit that evolutionary niche [Kauffman 1995].

At the end of the simulation it is our intention to find the global maximum of the previously unknown fitness landscape, the maybe singular point where the agent population as a whole exhibits maximum productivity, all the while constantly able to adapt fast and flexible to endogenous and exogenous changes of the environment. In that way we fully exploit the advantages of multi-agent systems [Nwana 1996] in the coordination of a business process environment: we provide an emergent solution to a large, inherently distributed problem, in a fashion that is modular, such reducing complexity, that works at high speed due to parallelism, and that attains fault tolerance because of redundancy. Building on current IT standards, we regard the concept illustrated here as not only promising science fiction, but feasible in the near future, and we will show this in our research beginning with our Artificial World simulation *Avalanche*.

5 References

[Axelrod 1985] Axelrod, R.: The Evolution of Cooperation. New York: Basic Books 1985.

[Baumann 1996] Baumann, J: The Mole Mobile Agent Library, White Paper, University of Stuttgart 1996.

[Chavez/Maes 1996] Chavez, A., Maes, P.: Kasbah - An Agent Marketplace for Buying and Selling Goods, Proceedings of the First Conference on Practical Applications of Agent Mechanisms, London 1996.

[Durfee/Rosenschein 1994] Durfee, E.H., Rosenschein, J.S.: Distributed Problem Solving and Multi-Agent Systems: Comparisons and Examples, in: Proceedings of the Thirteenth International Distributed Artificial Intelligence Workshop, p. 94-104, July 1994.

[Englert *et al*. 1996] Englert, J., Eymann, T., Gold, S., Hummel, T., Schoder, D. (1996): Beyond Automation: A Framework for Supporting Cooperation, in: Coelho, J.D. *et al*. (Ed.): Proceedings of ECIS '96, Lissabon, 1996. p.117-126.

[Epstein/Axtell 1996] Epstein, J.M., Axtell, R.: Growing Artificial Societies: Social Science from the Bottom Up, MIT Press/Brookings, Cambridge, MA, 1996.

[Eriksson *et al*. 1997] Eriksson, J., Finne, N., Janson, S. (Ed.): MarketSpace ’96 - An Open Agent-Based Market Infrastructure, UPMAIL Technical Report No. 147, Uppsala, April 7, 1997.

[FIPA 1997] FIPA - Foundation for Intelligent Physical Agents, Agent Communication Technical Committee: Agent Communication Language, FIPA ’97 Draft Specification: Part 2, http://www.fipa.org.

[Goldberg 1989] Goldberg, D.: Genetic Algorithms in Search, Optimization and Machine Learning, Addison Wesley, Reading MA, 1989.

[Holland/Miller 1991] Holland, J.H., Miller, J.: Artificial Adaptive Agents in Economic Theory, in: American Economic Review, Papers and Proceedings, Vol. 81, No. 2, 1991, p. 365-370.

[Hummel 1996] Hummel, T.: Chancen und Grenzen der Computerunterstützung kooperativen Arbeitens. Wiesbaden: DUV, 1996.

[Jennings *et al*. 1996] Jennings, N.R., Faratin P., Johnson, M.J., O’Brien, P., Wiegand, M.E.: Using Intelligent Agents to Manage Business Processes, 1996.

[Kauffman 1995] Kauffman, S.: At Home in the Universe, New York: Oxford University Press, 1995.

[Lane 1993] Lane, David A.: Artificial worlds and economics (Part II), in: Journal of Evolutionary Economics, p.177-197. Springer, 1993

[Langton 1989] Langton, C. (Ed.): Artificial Life, Addison-Wesley, Redwood City, CA, 1989

[Malone/Crowston 1990] Malone, T.W., Crowston, K.: Toward an Interdisciplinary Theory of Coordination, Technical Report 120, Center for Coordination Science, MIT, Cambridge MA, 1990.

[Nwana 1996] Nwana, H.S.: Software Agents: An Overview, in: Knowledge Engineering Review, 11(3), p. 24-29, 1996.

[Objectspace 1997] ObjectSpace (Ed.): ObjectSpace Voyager Technical Overview, White Paper, http:// www. objectspace. com/ voyager/ VoyagerTechOview.pdf, 1997.

[Porter/Millar 1985] Porter, Michael E., Millar, Victor E. (1985): How information gives you competitive advantage, in: Harvard Business Review, July-August , S.149-160. 1985

[Ray 1992] Ray, Thomas S. (1992): An Approach to the Synthesis of Life, in: Langton, C.G.; Taylor, C.; Farmer, J.D.; Rasmussen, S. (Ed.): Artificial Life II, S.371-408. Redwood City: Addison Wesley, 1992.

[Sandholm 1996] Sandholm, T.W.: Negotiation among self-interested computationally limited agents, Dissertation, Amherst 1996.

[Schoder 1995] Schoder, D.: Erfolg und Mißerfolg telematischer Innovationen, Wiesbaden: Dt. Univ.-Verlag, 1995.

[Schoder *et al*. 1997] Schoder, D., Hummel, T., Müller, G.: Interdisziplinäre Modelle für Entwurf und Einsatz telematischer Systeme, in: Jarke, M. (Ed.): Informatik'97. Heidelberg: Springer, 1997.

[Smith 1980] Smith, R.G.: The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver, in: IEEE Transactions on Computers, Vol. 29, 1980, p. 1104-1113.

[Stender 1994] Stender, J.: Many-Agent Simulation and Genetic Algorithms – Adaptation Mechanisms in Economic Models, in: Hillebrand, E., Stender J. (Ed.): Many-Agent Simulation and Artificial Life, Amsterdam: IOS-Press, 1994.

[Tesfatsion 1997] Tesfatsion, L.: How Economists can get alife, in: Arthur, W.B., Durlauf, S., Lane, D. (Eds.): The Economy as a Complex Evolving System, II, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Vol. XXVII, Addison-Wesley, 1997.

[Varian 1993] Varian, H.R.: Intermediate Microeconomics – A Modern Approach, 3rd Edition, New York: W.W. Norton, 1993.

[Vidal/Durfee 1997] Vidal, J.M., Durfee, E.H.: Agents learning about agents – a framework and analysis, in: AAAI-97 Workshop on Learning in Multi-Agent Systems, July 1997.

[Wellman 1996] Wellman, Michael P.: Market-Oriented Programming: Some Early Lessons, in: Clearwater (Ed.): Market-Based Control: A Paradigm for Distributed Resource Allocation, World Scientific. 1996.

[White 1997] White, J.: Mobile Agents White Paper, http:// www. genmagic.com/ agents/ Whitepaper/ whitepaper. html, 1997.