# The Behavior-Based Firm: An Application of Recent AI Concepts to Company Management

Erich Prem\* The Austrian Research Institute for Artificial Intelligence Schottengasse 3 A-1010 Wien, Austria erich@ai.univie.ac.at

### Abstract

Artificial Intelligence (AI) has a long tradition in developing technological means for the control of complex systems. This paper reviews recent developments in the area of embodied AI and behavior-based robotics and formulates principles as they appear to be applicable in managerial problem domains. We compare these principles to new management concepts such as the horizontal organization and lean production, which exhibit definite similarities to proposals recently made by roboticists. An analysis of these similarities identifies the importance of a tight system-environment coupling. This connection is achieved by a rapid and precise evaluation of external observables from many internal processes. Another important factor is the process orientation of control that marks a clear departure from traditional approaches based on functional decomposition.

## 1 Introduction

The systems under discussion are unthinkable, in the sense that they really are too

in the sense that they really are too complex to fathom.

[Beer 81, p.51]

Artificial Intelligence (AI) has a long tradition in providing a framework for economic and managerial questions. Evidence for this claim is, for instance, provided in the work of Herbert Simon who can be considered one of the most prominent AI researchers while at the same time being nobel laureate due to his contribution to economics. This proximity of AI and economics is surprising only at a first glance. On second thoughts, however, it becomes clear that the two scientific disciplines are in a specific sense complements to each other and of historic kinship as well.

The complementary nature of the economic sciences (including management and organization theory) and the technological science of AI is based on the dual nature of the fundamental questions which the corresponding disciplines are trying to answer. The AI researcher is concerned with the *generation* 

<sup>\*</sup>In: Applied Artificial Intelligence, 11(3), 1997

of intelligently behaving systems. Although there still is a large area of cognitive science research within the AI community-aiming at explanations for intelligent systems—the construction of artefacts is of central importance. Essential constructs of this engineering endeavour are hypothesized elements of rational thought and action that may range from rules in expert systems to neural networks in pattern recognition. The economist, on the other hand, is interested in the explanation of human behavior by means of a very similar hypothesized epistemic construct: rational choice [Simon 65, Elster 89]. In the applied field of the management sciences the constructive aspect is more important for being oriented towards the *generation* of control systems for complex socio-technical systems. It is this applied area where commonalities between both problem domains suggest the existence of mutually promising approaches to solutions.

Historically, the proximity of both disciplines lies in their common ancestor *cybernetics*. While parts of AI (especially its more symbolically oriented branches) have largely abandoned their cybernetic tradition, management science is open to draw parallels between results about control in animals and machines to the task of organizing and controlling enterprises.

The aim of this paper is to discuss the relevance of recent developments in the field of embodied AI to proposals made in the management sciences. The theory of the firm deals with a wide range of problems from financial and production planning to human resource management. The main source of the analogies between AI systems and human organizations is the underlying flow of information between the elements of both systems. According to [Malone 92] the similarity lies in the problem of "how to route information of a system responding to a complex, rapidly changing, and poorly understood situation." Therefore, an application of a conceptual framework taken from AI will mainly contribute to organization and management theory (OMT). A list of common concepts in AI and OMT [Gioia 92, p. 299] contains the following notions for which we can expect mutual support from recent AI and OMT developments:

- information flow
- $\bullet$  control
- process design
- task management

Other major research themes in the (organizational) theory of the firm like resource allocation or decision process management in groups will not benefit from our approach. Novel treatments of learning in the enterprise, however, will be considered in section 4.2.

# 2 Organization and management theories

It was not until 1900 that organizational and management issues turned into a subject for enduring scientific investigation. Research in this field has been directed by rather differing and often incompatible interests. At the outset of organizational theory research was driven by interest in rationalized structural planning. A major issue was to ensure complete power and control of the management over the executive working force. With the advent of higher social standards and an increasing individual liberty other factors such as personal motivation or satisfaction of personal needs became more important. In the increasingly competitive market, research in OMT was driven by the interest in information retrieval and decision processes and how

#### $\mathbf{2}$

they could be supported by different kinds of organizational structures. More recently, considerations about the dynamics of business processes, their rapid adaptation and quality management issues have driven research in this field.

#### 2.1 Traditional approaches

Traditionally, proposals for how to organize the productive and control processes in a firm have often been based on Taylor's ideas of "scientific management" [Taylor 11]. In this physiologically oriented theory, organization and management are regarded as tools for the control and optimization of the individual worker's labour force. Taylor's theory is embedded in the scientific world view at the turn of the century. Physics was the leading discipline at that time and prediction based on calculation was the prevailing technology that also seemed most promising for the control of complex systems. Accordingly, Taylor's recommendations were based on calculatory considerations that should lead to a maximum in efficiency as far as production processes were concerned. Taylor's idea was to regard the manager as a scientist who ought to study the fundamental laws of the enterprise he is in charge of [Freedman 93]. The scientific appraisal for the reductionism of the 19th century prepared the characteristics of industrial management in the early (and sometimes not so early) 20th century: dissection of production processes, enforced standardization and cooperation, total control by the management. As dehumanizing as this approach may be regarded today, it naturally arose from an economic environment of early mechanization and industrialization that was coupled with the availability of mainly unskilled workers.

The fundamentals of this approach to managing enterprises can still be seen in what is nowadays called "functional organization" [McGill & Slocum 96]. Its basis is a tight organization plan which describes the departments for every task in the production process and how resources are provided for the departments. The detailed organizational structure is in need of a considerable amount of coordination. Both the necessary coordinating facilities and the tight organization plan complicate necessary reactions to changes in the enterprise's environment. An additional problem that has been recognized recently is that the separation into company departments prohibits a transfer of (individual or departmental) experiences between divisions, i.e. it impairs company learning.

The natural organizational background structure for the control of the proposed systematic dissection is a strict hierarchy with its classical tree-like fabric. The connections in the hierarchical tree represent a multiplicity of different relations between organizational components such as

- flow of information,
- exertion of control,
- and responsibility.

The top levels of the tree are concerned with company goals and strategies alone. The leaves represent plan execution by the worker, for which extremely detailed plans have to be prepared by a "middle management" of mainly (downwards) controlling and (upwards) reporting clerks. OMTs of the late 1930's addressed the growth in administration that was unavoidable as firms increased in size and complexity and workers became increasingly more interested in prestigious rewards (salary and rank). This lead to the development of OMTs with an approach that was oriented towards the individual worker's motivation. As workers became better trained and shared a greater responsibility for their role in the production

#### $\mathbf{3}$

process, the theory concentrated on how to motivate the worker in a social context. The organizational model that promised to allow for high motivational incentives was a participative model in which groups of well-trained professionals cooperated on the solution of problems. In this model, the basic structure of the strict hierarchy is not removed but serves to carry working groups at the nodes of the tree. The critical points in this approach are the group coordinators who play a double role as head of one and regular member of another group.

Around 1950, the scientific theory of management changed significantly due to the advent of a completely new technology: the electronic computer. This development was parallelized by the development of novel mathematical theories and their application to computer science. Soon the new formal theories and the new technology both influenced managerial science. In this stimulating era, Herbert Simon and others studied management problems in the light of decision-making with limited information. The result was a highly formalized and mathematically well-founded theory of rational choice. However, this theory largely neglected the system environment due to its concentration on internal and information oriented aspects of decision-making.

# 2.2 New challenges and new approaches

The nowadays widely acknowledged new challenges for enterprises are to a great extent due to changes in the firm's environmental conditions. In the last years enterprises have been increasingly confronted with changes in consumer behavior, internationalization and deregulation of markets, decreasing governmental influence, etc. Today's customers demand excellent quality, friendly and individual service as well as minimal delivery periods. These customer requirements are paired with completely new market situations. The internationalization and often globalization of markets goes together with the decreasing possibility or willingness of governments to protect companies from severe market fluctuations. The consequence of these developments is a market situation that is much more competitive than it used to be. Consumer's purchasing decisions become more and more flexible, increasingly based on international comparisons and highly unpredictable. This development is parallelized by increasing pressures on the cost-structure of enterprises and the obligation to ensure shorter terms in order, production and delivery as well as minimal stocks. Therefore, today's enterprises are forced to strive for reductions in production cost while at the same time having to improve product quality.

The continuously increasing rate and amount of technological change provides an additional challenge for modern companies (cf. [Beer 81]). Not only are the markets for products changing rapidly, production and communication technology are another force for the management to permanently improve on the organization of the firm. Technological progress touches upon all aspects of the enterprise from production, to distribution, from communication to new materials, and from management to marketing. The new environmental challenges for enterprises can be summarized as consisting in

- increasing complexity,
- high uncertainty,
- rapidly changing markets, i.e. high variability,
- high rates of technological change,
- and shorter terms in order, production, and delivery.
- 4

These characteristics are challenges for modern enterprises because traditional managerial principles, like those of Taylor, have lead to rigid forms of organization which are inappropriate for dealing with rapidly changing markets.

The acknowledgement of the importance of environmental factors in the theory of the firm happened in parallel with the development of cybernetics that culminated around 1960. Unlike a decision theoretic framework that has a great proximity to information theoretic and mathematical considerations, the new science of cybernetics pursues a much more integrative program. The integration consists in the general acknowledgement that the object under study, e.g. the firm, is a complex system in a complex environment. The cybernetic tradition regards it as subject to different specific observers' views and as an object of several distinct sciences. This view of the enterprise is a holistic and system-theoretic one. In a cybernetic approach to the theory of the firm, the enterprise is considered as a socio-technical system that possesses a tight coupling between the system and its environment, adapts to its environment and is an inherently open system.

This very general picture has only recently been refined to new theories that deal with internal and external complexity and rapid environmental change. Terms from systems science and physics such as "chaos, dynamics, attractor, holism," etc. have been taken up by managerial economists and are among the new buzz-words that occur in articles on management issues, be they scientific or directed towards a broad public. Starting point for this development was the scientific discovery that very small perturbations in dynamical systems may lead to completely novel and unpredictable system behavior. Management scientists have taken this result from dynamical systems theory and argued that such effects also happen in markets or during phases of the development of new technology. Due to their inherently dynamic and unpredictable nature these developments cannot be perfectly controlled but must be taken into account. Unfortunately, managerial techniques which have been proposed to deal with these issues have remained rather vague in the past. One frequent recommendation is to *accept* these challenges rather than to counteract on them. In order to gain control of dynamical changes the manager must adopt a *systemic* view of the *whole*.

While these recommendations are certainly well intended, their direct application is far from being straightforward. In fact, usage of terms from dynamical systems theory can mainly have an epistemic value rather than providing technological tools. This is not surprising to the engineer who tries to control such systems. The reason why nearly all technical systems work at rather low interaction dynamics or ensure severe restrictions on turbulances, etc. is based on the fact that tools for controlling such systems are comparatively poorly developed. On the other hand it is true that nature has developed a whole set of complex systems which seem to cope perfectly with dynamic environmental changes and high interaction dynamics. These complex systems are organisms or living creatures. The intuition is, of course, that we can learn from organisms how to survive in complex, changing environments. However, we do this with the intention to generate artificial control structures. In the following section we therefore take a closer look at new technological approaches to the creation and control of creatures that are not living. Still, they exhibit some important characteristics of biological organisms: autonomy, survival, and high interaction dynamics.

#### 2.3 Metaphors and theories

Management theories often suffer from the problem that they make extensive use of metaphors as opposed to exact models. A metaphor is a transformation of one area of discourse to another one with certain general structure-preserving characteristics. New management proposals regard the enterprise as an organism or living creature. Previous (e.g. Taylor's) suggestions centered on the machine metaphor, on functional task separation, planning and control. Today's approaches regard an animal's autonomy, creativity, stability as concepts that describe qualities which enterprises should pursue also.

Fig. 1 depicts the structure of this metaphor. Living systems and firms are entities in the natural world. Observations and measurements of both systems lead to formalized descriptions of their corresponding behaviors. The process of metaphorical reasoning consists in taking parallels in the behavior and environmental circumstances of organisms and firms (1). The next step is to metaphorically suggest to draw parallels between the organization of the firm and living systems (2). Finally, methods that have proven to be useful in the control of one system are applied to the other based on the assumed similarity in internal organization. An early example for such a strategy is Beer's description of the firm as a biological control system [Beer 79] (taken up, for example, in [Malik 96]).

A metaphor's value lies in its suggestive power to create solutions which are taken up from the better known system and applied to the system of interest. It remains problematic, however, that this transformation is of course not a clear-cut scientific process rather than an artistic technique. Many of the parallels that have been drawn between enterprises and organisms therefore remain on a superficial level without the ability to directly suggest



Figure 1: A system science view of metaphors (cf. [Rosen 85]). Entities in the natural world are mapped on formal descriptions. The metaphor is a process of backwards reasoning from similarities in the description of the systems to similarities in their organization.

the right actions or structures. Very often such metaphors have mainly been used as epistemic tools that would further the understanding of structures and processes in managerial or economic contexts.

The manager, on the other hand, is forced to act and create structures that will allow to solve problems. This is basically an engineering problem, not a merely epistemic one. If we want to take the metaphor of the enterprise as an organism seriously and try to push it beyond its mere descriptive and explanatory capabilities, we have to consider scientific results in the area of the creation and control of systems with parallels to organisms by technological means. One such area is the field of autonomous robotics which studies the construction of artificial creatures that act in unstructured environments. A brief look at requirements for such creatures already indicates the proximity to the managerial problems which have been described here. According to [Brooks 91b] these requirements for

6

creatures include:

- coping appropriately and in a timely fashion with changes in a dynamic environment
- robustness with respect to the environment
- the ability to maintain multiple goals and to pursue a particular one according to current circumstances
- being active and having a purpose in being

The metaphor, which we are going to study, consists in regarding the enterprise in the changing market economy as an artificial creature in its dynamic environment. The hope is to take advantage from technological tools that have been recently developed to control such systems in AI. In the next section we take a look at old and new control techniques for autonomous robots.

# 3 AI and Robotics

#### 3.1 Traditional approach

"Good old fashioned AI" bases its theories and technological approaches to the problem of generating (and understanding) intelligent behavior on the symbol system hypothesis [Simon 69, Newell 80]. Briefly summarized, the implicit idea is that in human reasoning perception delivers symbolic descriptions of the environment. These descriptions are then subject to intellectual processes in which symbolic manipulations are used to solve intellectual problems, e.g. to find the best action in order to pursue one's goals. The outcome of such symbol manipulation procedures is the generation of actions. Fig. 2 (a) depicts the classical structure of these control architectures. The mapping of perceived qualities to symbolic descriptions is performed so as to construct a symbolic model of the world inside the system which is detached from the real world and used instead of it to hypothesize on the outcome of actions, i.e. to *plan* what to do next. It has been pointed out by [Brooks 90] that "symbol systems in their purest forms assume a knowable objective truth." *Knowable* here means that it is, at least in principle, possible to generate true symbolic descriptions of the environment.

In such architectures the typical control loop can be described as sense - model - plan - act (SMPA). This means that sensing takes place only once once during a complete run through the loop. Additionally, planning is completely detached from the real world, since it only acts on the symbolic world model. Consequences of generated actions can only be taken into account for a probably necessary plan adaptation after they have been sensed and become expressed in the model again. This approach successfully separates the dynamics of the system-environment interaction from the dynamics of the control structure by introducing a symbolic image of the external world. This image, i.e. the symbolic world model, replaces the world and enables the usage of planning algorithms that operate on a purely formal system. This operation can happen at a time-scale that is completely different from the system's interaction with its environment. In fact, the original hope was that the planning algorithms would become so efficient that the interaction dynamics would be no problem any more.

Another distinguishing feature of the SMPA-architecture is the functional decomposition into modules of rules which only interact at very low dynamics. The methodological considerations behind this approach have been outlined by Simon as early as 1969 [Simon 69] and criticized amongst oth-



(b) behavior-based approach

Figure 2: Traditional and behavior-based robot control architecture. (After [Brooks 85].)

ers by [Prem 96b]. One assumption in this methodology is that it is possible to generate an appropriate system-environment interaction from such strongly decoupled modules. Simon argues that the behavior of complex systems can be approximated by modules that are driven at low interaction dynamics. This is true, if the system to be modeled is "nearlydecomposable", i.e. if the systems can be described by nearly-decomposable matrices. In this case, the difference between a system and its model is negligible small for short prediction periods. The problem is, however, that this analysis is wrong for open systems which possess relatively high interaction dynamics. In such a case, it is no longer easily possible to capture the observed phenomena by a fixed set of observables nor to fully account for them by nearly-decomposable systems [Rosen 85, Prem 95]. The construction of the traditional robot control architecture was lead in this wrong direction because the fact that robots are open systems had been largely neglected. The simplicity of the control structure and its decoupled dynamics arise from the simplicity of the interface between

the system (the robot) and its environment. Once one assumes an inherently simple relationship between the environment and an abstract system interface, the control problem turns into a purely algorithmic problem of mapping symbols (sensor values) to other symbols (effector output). In this abstraction, time within the system (algorithmic planning and calculation time) is completely separated from time outside the system (the interaction dynamics).

Well-known practical problems with this architecture are that

- the interaction dynamics between system and environment is low,
- symbolic perception is a very complex task,
- the correspondence of the world model and the world is extremely difficult to maintain,
- novel aspects of the environment are hard to cope with.
- 8

In order to overcome these difficulties the new approach of "behavior-based robotics" was introduced in the late 1980's. Embodied AI is its direct successor and described below.

#### 3.2 Embodied AI

The credo of embodied AI can be succinctly summarized as

Intelligence is determined by the dynamics of the interaction with the world. [Brooks 91a]

It is based on the idea that in order to build intelligent autonomous systems it is necessary to have it directly and dynamically interact with the world. The departure from traditional approaches is characterized by an increasing importance which the physical structure of the robot body plays. This body is not regarded as a mere box to be moved by the robot's control system. Instead, it is an integrated part of the interaction process between the robot and its environment. To the behavior-based roboticist it is unimportant whether some desired behavior is generated by a computational process in the robot's software or by a physical characteristic of the shape of the robot body. The main effect of this new emphasis on physical aspects is a dramatic departure from the traditional control architecture with the goal to increase the robot's interaction dynamics. Some of these new architectural considerations based on the requirements for creatures are the following [Brooks 91b]:

- In order to cope quickly with changes in the environment it is necessary to sense the environment often by evaluating rather simple predicates as they are needed by the individual subsystems.
- Robustness is achieved by means of multiple parallel activities. There is no central model of the world, individual lay-

ers extract only relevant aspects ("projections of a representation into a simple subspace").

- Each layer of control has its own implicit purpose (or goal), "sometimes goals can be abandoned when circumstances seem unpromising".
- The overall system action is driven by the autonomously acting parts, no central (or distributed) process selects from an explicit representation of goals to decide what to do next.



Figure 3: A simple behavioral module. The applicability predicate p decides, whether the input I is transformed to the output O by means of the transfer function transfer by possibly delaying or suppressing the output in a (cf. [Connell 90, Balkenius 95]).

Fig. 2 compares this approach to the traditional sense – model – act – plan. Instead of a functional composition of the control architecture where only a few parts are in charge of sensing environmental circumstances, the behavior-based approach consists of behavioral decomposition into modules. Each of these "behaviors" is able to sense the world as it is necessary for its task. Higher levels can influence the outcome of lower levels by delaying or overruling their commands to the effectors. Fig. 3 describes how one such simple behavior module works.

A real-world implementation of this architecture looks much more complicated since in-



Figure 4: The structure of a realistic network from real-world subsumption architectures. (Examples can be found in [Angle 89, Angle 91, Connell 88]). Input from sensors feeds into different layers. Different parallel processes (proc.) can suppress output from lower-level processes. Only a few processes directly control effectors.

dividual layers may actually be composed of several simpler sub-systems, an example is depicted in Fig. 4. An important element in these networks is the role that time plays in the system. Outputs of lower-level behaviors can be suppressed, delayed, or changed by higher-level behaviors. In this way, real-world clocks participate in driving the system's dynamics.

Various different robots with behaviorbased control systems that exhibit robust behavior and high interaction dynamics are described in [Brooks 90]. The behavior-based architecture brings with it a departure from the credo of symbolic AI. Symbols and plans are no longer of central importance, at least as long as each layer of this architecture is able to perform its task based solely on the incoming measurements. This is the reason why this approach to building robots has sometimes been called "reactive", which unfortunately is a misleading term. The fact that individual layers of the architecture do not possess representations of any kind is not a central dogma of this approach to embodied AI. It is, however, a central claim that the reaction of a single layer happens fast enough and it is in this specific sense that the term "reactive control" can be used here. In today's behavior-based robot architectures there are hardly any complex forms of symbolic representations. This is because embodied AI is still a very young field that is to a great extent busy with the construction of rather simple animats that perform simple tasks. One notable exception is the work of [Mataric 92], where the robot is equipped with a representational system in order to include map-like spatial navigation. Another embodied AI project that addresses complex cognitive issues and representational systems is the "Cog"-project at MIT's AI laboratory [Brooks & Lynn 94]. In this project the aim is to build an approximation of a humanoid system capable of interaction with its

environment at a human time-scale.

It can hardly be doubted today that the introduction of behavior-based principles into robotic systems has lead to the construction of faster robots which are far more robust than their classical ancestors. The question remains, how far this approach to the generation of intelligent behavior can go. (This is one of the purposes of the "Cog"-project.) One of the central questions here is, whether complex "cognitive" human abilities can be achieved by an approach that starts from such simple actions as "move" and "avoid". It is clear that some functional equivalent to symbolic representations will have to be introduced in one way or the other. It is, however, not clear what these systems will look like and what the special characteristics of this kind of representations will be. Therefore, this question remains to be an empirical one. Many a priori refutations, e.g. [Tsotsos 95], are based on a misunderstanding of the behavior-based approach as strictly "reactive" or solely on the long tradition of symbolic approaches in AI.

The essentials of the behavior-based approach as relevant in our context here and the core of the metaphor can be summarized as follows:

- Behavior-based systems are *active*. They interact autonomously and continuously with their environment based on permanently running parallel control processes.
- Behavior-based embodied AI systems generate *robust behavior* even in dynamic environments. This is because the physics of the system is taken into consideration at design time. Also, the status of the world is repeatedly checked and used without reliance on an abstract model that is temporally and semantically detached from the current status.

• The robots operate with a very high interaction dynamics with their environment.

In addition to these features, the methodology allows for robust prototyping. Even only one or two behavioral layers generate a primitive, but "complete" system in the sense that it interacts robustly with its world. It may not exhibit very clever behaviors and probably not fulfill any clever task. However, a rudimentary animat constructed in this fashion, will allow for higher-level add-ons that can safely build on functioning low levels.

The success of the behavior-based approach to the construction of autonomous robots lies in the following principles.

- The methodology puts a strong emphasis on the dynamics of the systemenvironment interaction of every single module and on the physics of the system.
- There is no separate "input processing" part, every behavior receives inputs directly from the environment.
- The coupling between separate modules leads to a highly interconnected communication and control structure.
- The system acts under lack of information and each behavior is acting independently of others.

# 4 Consequences for the Firm

Taking the approach of traditional robotics as a metaphor for traditional management approaches, the implications of embodied AI for new management theories are straightforward. In order to create something like the "behavior-based enterprise" the following points must be taken into consideration.

- The overall organization of the system must ensure that there is a tight coupling between market demands and company processes.
- Each part of the firm does not need complete information. Individuals or departments only need that information which is necessary to accomplish their tasks. The update of this reduced information, however, must happen as often and quickly as possible.
- Sub-structures should be acting in an autonomous fashion. They *can* be overruled by other structures, but usually are not.
- There is high interaction between the different parts.
- Sub-structures may pursue their own goals. The goals of the whole system are an emergent consequence of the goals of the sub-systems.

Although these principles may seem farfetched, some of them have already been proposed by management theorists. The following section discusses these suggestions that have arisen in real-world management situations. We will see how they relate to our behaviorbased principles of organization.

#### 4.1 Real-world approaches

#### 4.1.1 Lean production

In the mid-80s, under the pressure of extensive automation in Japanese industries, a new approach to industrial production emerged throughout European and US companies. Japanese companies were able to produce consumer goods with an extremely high degree of automation, good quality and at considerably fewer costs compared to their competitors. The new approach was soon to be termed "lean production". The lean company focuses on product-market relationships and is therefore forced to achieve short reaction times to the market. This is achieved by flat hierarchies and production in small organizational units.

This increased efficiency was mainly due to a new organizational structure in the production process. Specialized departments of a producing firm were dissolved and transformed into working groups. The new credo consisted in a focus on small production teams and high standard of technological support to save production costs and improve product quality. Production through these small units happens in teams of highly qualified and motivated workers who are responsible for their product in terms of production time and quality. Division of labour within such teams is kept at a minimum, whereas this division plays a strong role in conventional production processes. Instead of minimizing the time for each single step in the process, lean production seeks to minimize the overall production time per order. This is possible, because production traditionally happens in sections and lots, lean production aims at a continuous process flow. An additional advantage in terms of minimal costs consists in the drastically reduced stock of goods. The lean factory only keeps a minimal stock of products that are necessary for the prodution process. In this way, it forces its suppliers to deliver within very short periods. This does add some fragility to the process. because a single spare part that is out of stock may interrupt the whole production process.

The enabling factors for this new organization is a high degree of automation at the technological side and the availability of well trained professional workers on the other. The problem solving competence of the individual working group member is a distinctive feature of lean production. It allows for rapid problem solving, since the expert is where she is needed. In the traditional assembly-line pro-

duction a mistake by one of the workers or a technical problem can imply an interruption of the whole production process until an expert solves the problem.

Another main advantage of this type of production organization is that the individual members of working groups are involved in a complete production circle and therefore able to assess product quality and learn from failures. Optimization does not happen in terms of improving every single function of the production process or every individual working step. Instead, optimization aims at taking the whole system into account with respect to the product and consumer demands. Additionally, workers are trained on many different abilities and the changing requirements tends to keep motivation high.

Comparing these characteristics of lean production with our previous list, we identify the following realizations of principles of the behavior-based enterprise.

- The tight coupling to market demands is ensured by a flexible production that easily adapts to new products. At the demand level, the working groups are highly involved in consumer's needs. It is not only a marketing or planning department that controls the company goals.
- Information about qualitative requirements or causes of problems in the production process are available to the whole working group. The slogan often mentioned here is that the worker is also the manager of the firm. While this is certainly an exaggeration, it is true that every single worker is supplied with more and more relevant information than in traditional organizations.
- Working groups autonomously pursue their given goal instead of being given a series of detailed working steps.

- High-interaction among the members of a group ensures short communication channels and fast reaction to any changes.
- The departure from an orientation at product lots towards an acknowledgement of the production process ensures that the notion of time plays an integrative role throughout production.

The process-oriented view of lean production has proven so useful that it seems only natural to extend its application to management per se, which is the subject of the next section.

#### 4.1.2 The horizontal organization

It seems straightforward that a flat structure of organization is a useful technique for coping with dynamic environments, not only with respect to production. New techniques for the organization of whole companies were developed with the newly gained insights into the advantages of smashing the management hierarchy [Byrne 93, Stewart 92]. The enabling technology for this development were affordable, integrated information systems capable of accompanying activities throughout all business processes. The *horizontal organization* is characterized by

- a flat, information-based organization,
- "high-involvement" workplaces,
- a management that is based on the definition of business processes
- aligned with customer demands.

The horizontal organization centers around the management process. It concentrates on autonomous decision processes and a reduction of complicated and costly hierarchical levels. Consumer groups define strategic operational areas. The structure is again processoriented as opposed to activity-oriented. The

primary management task is the definition of those processes which optimally satisfy consumer needs. The hierarchy is replaced by working groups who possess a high degree of autonomy and responsibility for their performance. Information transfer happens to those who actually possess the competence for making decisions. This information is not prefiltered by higher levels of management as is typically the case in a hierarchy of managers.

The definition of the business process based on customer needs is the central principle of horizontal management. Fig. 5 depicts a part of the organizational structure of a prototypical horizontal organization. Examples for typical "core" business processes include development of new products, sales, fulfillment, and customer support. Core business processes



Figure 5: The structure of the horizontal organization. Process owners (PO) who are in charge of business processes report directly to the board and the chairman (CM).

may in turn be decomposed into more primitive engagements of enterprise employees. The customer support process typically involves research, advertising, and service teams, thereby leading to a structure which more and more may start to look like Fig. 4. The core business process defines those lower-level processes which are necessary to satisfy the customers. In the same way, resources are allocated to processes as they are needed, not as they happen to be available. Processes are defined to be successfully managed, if they serve to satisfy consumers, not the managers.

Similarly to lean production, horizontal organizations are driven by small, highly autonomous teams, where workers are highly involved in the overall doing of the enterprise. Typically, a worker will be involved in a rather broad spectrum of activities. This is different from the functional organization, where workers would become experts in a single specialized task, e.g. some sub-task of accounting. Just like the worker's skillfulness is an enabling resource for lean production, knowledge and flexibility are essential characteristics for the personnel in the horizontal organization. The high-involvement in the whole firm's doing is ensured by raw, i.e. unprocessed and unfiltered information from the environment in which a business process is embedded. It is worth mentioning the problem that promotion in the horizontal organization is rather different from promotion in hierarchies, due to the lack of ranks. In fact, the elaboration of promotion schemes still is a problem and research issue in this area.

The autonomy of teams at the business process level means a liberty to decide on means for goal achievement, it is not a freedom to select the goals of the enterprise. The new task of the remaining management board is indeed the formulation of clear goals to be achieved by the set of business processes. Only based on these individual process goals is it possible for the working teams to continually evaluate their behavior. Goals at the process level are always in close proximity to consumer needs such that an overly specialization in task achievement is avoided.

There are a few subtle points about these working teams:

- It is necessary to have these teams as opposed to single individuals, because only teams have enough potential to improve on the business process management.
- Individuals within teams are not specialists, but rather generalists with special skills.
- Learning is ensured at the team and individual level (see below).

Lean production and horizontal management have proven to be successful techniques in reengineering firms. The next section serves to provide a deeper and more general understanding of the usefulness of a process-oriented view.

#### 4.2 Learning

OMTs as well well as research in embodied AI have only recently focussed on learning as a means for optimizing system performance.

Learning in autonomous robots to date is rather simple and happens mostly at the level of a single behavior, i.e. the quality of a single behavior is optimized during learning. Examples for learning tasks in this sense are map building [Nehmzow & Smithers 91, Mataric 92] or obstacle avoidance and path finding [Verschure & Pfeifer 93, Millan 92]. The general approach in these architectures is depicted in Fig. 6. In order for any action optimization procedure to function properly, it is necessary that the outcome of actions is quickly and precisely available. In Fig. 6 the connection to the environment is ensured by the learning signal r that should depend on the outcome of actions (cf. [Prem 96a]). This information must be directed to those entities which drive the optimization and should rapidly improve the quality of actions.

Novel OMTs, e.g. [McGill & Slocum 96], put a strong emphasis on individual and or-



Figure 6: The success of a behavior controls the construction of the transfer function. This can be achieved by a training signal (in r) that depends on the outcome of actions and the success or failure of the system's behavior.

ganizational learning in the firm. These theories argue that the learning ("intelligent") enterprise is substantially different from earlier forms. The oldest approach, which can be identified with the classical functional hierarchy, is the "knowing" organization. The competence which individuals acquire in this organization is, of course, a highly specialized one, limited to the functional department in which the individuals make their careers. Replacing a member in this hierarchy is a difficult task that can generate tremendous costs and complications. Again, the specialization is a consequence of the low inter-component linkages in the hierarchy.

Process- and team-oriented approaches, on the other hand, support company-wide learning. The requirement of fast and reliable feedback is met by the tight environmental coupling. The concentration on working teams avoids overly specialization.

- Learning is ensured at the team level because of the existence of clear goals for the whole team.
- Learning at the individual level remains, of course, important. Typically, individuals rotate through the factory's work areas with the goal of getting to know as
- 15

much of the entire enterprise as possible. Since team members usually stem from different fields, intra-team learning is also facilitated.

There is, of course, still a need for an overall system management. Working groups are capable of optimizing the processes and detect customer dissatisfaction. The top-level management, however, will still have to strive for new customers and markets, review new technologies and set strategic goals. This can also be compared to the behavior-based approach, where the top two levels are "reason about behavior of objects" and "plan changes". Many have argued that this is the point where more traditional AI techniques will have to be used (e.g. [Lammens et al. 93]). Whereas it is obvious that some sort of "planning" and consideration of "goals" will play a major role at this level, it is far from clear what these processes will have to look like. Some argue for traditional AI techniques, others believe that abstract reasoning is also based on bodily interaction with the world and therefore centers on processes that are rather different from conventional approaches [Johnson 87].

#### 4.3 Approaching the behaviorbased firm

After this view at new management proposals and robot technologies the question arises as to why these new approaches seem to be more suitable for the design of complex dynamic systems. Table 7 gives an overview of the actions taken in the new methodologies and the effects they had on the corresponding systems.

The first observation when comparing both methods is that supervisory structures are eliminated. The organization hierarchy in the enterprise parallels the classical robot control scheme in the sense that there are strong centralized structures "planning", respectively

"chairman", that are in charge of summing up all information and delivering clear plans for subsequent layers of control. Information flow in the traditional versions of firm or robot control organization is clearly channeled. Fig. 8 depicts this aspect of the metaphor, where the departmental structure corresponds to the encapsulated subroutines of the SMPA-architecture. The top of the organizational tree is reflected in the metaphor as the central processor that dictates the sequence of operations and subroutine calls. The functional decompositions of both approaches tend to maximize intra-component linkages while at the same time reducing inter-component connection. The functional encapsulation seems to be useful from the viewpoint of specialization. The functional entities are designed to be specialists. The price to pay is, however, a complicated and time-consuming communication structure. In both functional systems (robot and firm), the immediate contact to environmental conditions is lost for many instances within the structure.

This problem is a direct consequence of functional decomposition. Firms and robots are designed to function, to attain goals. Problem decomposition in the traditional approach means the generation of subgoals. In both approaches, subgoals are identified with states or objects, *not* with processes. This inherently static view carries the core responsibility for the elimination of dynamics. Processes are only considered as internal state transition operators that transform one state or object into another. In this architecture, a change in the set of goals implies a radical change in the subgoals for the functional experts.

The new process-oriented views operate in a completely different fashion. Overall system goals are described in terms of the processes that lead to their fulfillment, *not* in terms of states or objects. The important fact about these processes, and a necessary re-

action	effect in embodied AI	effect in firms
elimination of	reduced module complexity,	reduction of costs,
functional compartments	faster communication	faster communication
emphasis of	working systems in	adaptive market and
environmental relation	natural environments	consumer orientation
elimination of	frequent evaluation of	unfiltered and frequently
separated "input processing"	environment observables	updated information
independent	robust systems,	highly motivated teams
autonomous behaviors	fast calculation	support learning
tight coupling of	integration of	intrinsic importance
behavioral modules	external time	of time

Figure 7: Effects of the new methodologies on behavior-based robotics and firms.



Figure 8: The similarities between the classical approach to robot control and the vertical organization.

quirement for this approach to work properly, is that they come with *conditions* that allow for checking whether a process operates successfully. This verification does not happen by means of a complicated information and control structure. It happens as immediately as possible by operating on the *environmental* observables, *not* on internal transformations and interpretations of these observables. Another main difference is that the direct evaluation of external observables is directly related to the overall system goal, not to the fulfillment of a subgoal that has been generated in an analytic way.

The specific way in which overall system function is decomposed in the traditional approach is based on the physicalist and reductionist approach as it is expressed in the machine metaphor. This metaphor comes with two important principles. First, the task of interacting with the environment is considered to be essentially computation. Secondly, functions are to be separated in modules. The first principle implies that in the machine-oriented approach, the overall system control task is viewed as the problem to generate a mathematical algorithm that transforms symbolic inputs (numbers delivered by sensors, measurements of observables) into symbolic outputs (numbers sent to actuators, commands for actions taken by the firm). Such a computational approach means the generation of a state-transition sequence, where the states contain all the necessary information for the

algorithm to function properly. The very notion of a computational algorithm does not contain the concept of time. This explains the separation of internal system (computational) time from environmental time. The fact that the states are considered to bear all the information necessary for determining subsequent transitions is the ultimate reason for the idea of the generation of world models in robotics (and, similarly, in management approaches that center on detailed planning). The second principle that comes with the machinemetaphor implies the general tacit assumption that function is always localized in structure and that, accordingly, sub-functions must be localized in sub-structures. This stems from the idea that technological products tend to be constructed from modules, each of which has a clearly specified function. A component of such modules then simply has all the functions that all the individual components have together. Theoretical biologists have opposed this position by arguing that in living systems the relation between function and structure can be more complicated [Rosen 91, Rosen 93]. As a simple example a bird's wing may be considered, which compared to an airplane is airfoil and engine at the same time. Such a direct coupling of different functions in one structure seems to be a characteristic of biological systems. This is not true in the traditional design of artefacts, where components carry sub-functions. (However, clever design and new technology tend to blur this difference. For example, an aeroplane's wing also serves as the gas container.) To say the least, it is true that structures in behavior-based robots and horizontal organizations tend to carry more functions for the overall doing of the system. The gripper of the mobile robot "Toto" serves as gripper, sensor, and memory at the same time [Connell 90]. And the elimination of hierarchical structures in the horizontal organization implies that the functions of higher organizational levels must be executed by the new business processes.

The process-oriented methods blur the algorithmic view of the traditional approach. These schemes come with an inherent parallelism in which internal system time can be given a straightforward interpretation since every single process interacts with the world. A reasonable notion of time thus consists in the delay that is needed for every individual module's interaction with the environment. The clear sequence of states is further blurred by an organization that does not show simple one-to-one relationships between organizational institutions. Comparing Fig. 4 and 8, the graph of the horizontal organization contains more edges per vertex than the graph of a corresponding vertical scheme.

The above mentioned parallelism is responsible for another advantage of the behaviorbased approach. The idea to generate complete environment-process-environment loops ensures a certain minimality in design. It leads to robust low-level behaviors. The effect in robots is that higher levels can be designed with a reliance on low-level features that are not changed by adding on top of the hierarchy. In firms, the major advantage consists in customer satisfaction with a minimum of resources.

#### 4.4 Conclusion

In this paper we have investigated a new scientific metaphor. It consists in taking complex artificial systems as a picture of firms. This approach suggests to take approaches to the construction of robot control systems and compare them to recent theories developed for the management of firms. From these observations it follows that management principles for firms must be based on

- a tight environmental embedding of business processes
- that are driven by costumer needs and
- fulfilled by interacting, cooperative working groups
- that are well-informed by rapidly evaluated external observables.

These suggestions are not based on just another computer architecture that happened to overcome disadvantages of previous solutions. Our analysis has shown that the new solutions are based on a rather drastic departure from old system-theoretic considerations. The complexity of living systems is reconsidered as a source for the management of complex artificial systems. The new approaches, however, are much closer to the architecture of these natural systems than previous suggestions. However, our comparatively poor understanding of control strategies in organisms ensures that the study of these similarities will still have a long way to go from here.

#### Acknowledgments: The

Austrian Research Institute for Artificial Intelligence is sponsored by the Austrian Federal Ministry for Science, Transport, and the Arts. The author wishes to thank Rodney A. Brooks for his support and Karl Vieider for generous help with the literature. I am also grateful for comments on an earlier version of this paper by Johannes Fürnkranz, Paolo Petta, and Robert Trappl.

# References

[Angle 89] Angle C.M.: Genghis - a Six Legged Walking Robot, Massachusetts Institute of Technology, Cambridge, MA, Bachelor's thesis, 1989.

- [Angle 91] Angle C.M.: Design of an Artificial Creature, Massachusetts Institute of Technology, Cambridge, MA, Master's thesis, 1991.
- [Balkenius 95] Balkenius C.: Natural Intelligence in Artificial Creatures, Lund University Cognitive Studies, 37, 1995.
- [Beer 79] Beer S.: The Heart of Enterprise, Wiley, Chichester, UK, 1979.
- [Beer 81] Beer S.: Brain of the Firm, Wiley, Chichester, UK, 1981.
- [Brandmeyer & Otte 92] Brandmeyer K., Otte T.: Marken I: Lebende Systeme mit Dynamik, Harvard Manager, 3, pp. 27-32, 1992.
- [Brooks 85] Brooks R.A.: A Robust Layered Control System for a Mobile Robot, AI-Laboratory, Massachusetts Institute of Technology, Cambridge, MA, AI-Memo 864, 1985.
- [Brooks 90] Brooks R.A.: Elephants Don't Play Chess, in Maes P.(ed.), Designing Autonomous Agents, MIT Press, Cambridge, MA, Bradford Books, pp. 3-16, 1990.
- [Brooks 91a] Brooks R.A.: Intelligence without Representation, in Special Volume: Foundations of Artificial Intelligence, Artificial Intelligence, 47 (1-3), 1991.
- [Brooks 91b] Brooks R.A.: Intelligence without Reason, in Proceedings of the 12 International Conference on Artificial Intelligence, Morgan Kaufmann, San Mateo, CA, pp. 569–595, 1991.
- [Brooks & Lynn 94] Brooks R.A., Lynn A.: Building Brains for Bodies, Autonomous Robots, Vol. 1, 7-25, 1994.
- 19

- [Byrne 93] Byrne J.A.: The Horizontal Corporation, pp. 44-49, Business Week, December 20, 1993.
- [Cyert 92] Cyert R.M.: A Behavioral Theory of the Firm, Blackwell Publishers, Cambridge, MA, first published by Prentice-Hall, NJ, 1963, 1992.
- [Connell 90] Connell J.H.: Minimalist Mobile Robotics, Academic Press, San Diego, 1990.
- [Connell 88] Connell J.H.: A behavior-based arm controller, MIT Artificial Intelligence Laboratory, A.I. Memo 1025, 1988.
- [Elster 89] Elster J.: Nuts and Bolts for the Social Sciences, Cambridge University Press, New York, 1989.
- [Freedman 93] Freedman D.H.: Was kommt nach dem Taylorismus?, (Is Management Still a Science ?, Harvard Business Review, 6, 1992), Harvard Manager, 2, pp. 24-32, 1993.
- [Gioia 92] Gioia D.A.: Common Ground? The intersection of artificial intelligence and organization and management theory, in Masuch M. & Warglien M.(eds.), Artificial Intelligence in Organization and Management Theory, Elsevier, Amsterdam, pp. 295-309, 1992.
- [Johnson 87] Johnson M.: The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason, University of Chicago Press, Chicago, 1987.
- [Lammens et al. 93] Lammens J.M., Hexmoor H.H., Shapiro S.C.: Of Elephants and Men, in Steels L.(ed.), The Biology and Technology of Intelligent Autonomous Agents, NATO ASI Series "Computer and Systems Sciences",

Vol.144, Springer, Berlin, pp. 312-344, 1993.

- [Malik 96] Malik F.: Strategie des Managements komplexer Systeme, Haupt, Bern, 1996.
- [Malone 92] Malone T.W.: Analogies between human organizations and artificial intelligence systems: two examples and some reflections, in Masuch M. & Warglien M.(eds.), Artificial Intelligence in Organization and Management Theory, Elsevier, Amsterdam, pp. 21-40, 1992.

[Mataric 92]

Mataric M.J.: Integration of Representation Into Goal-Driven Behavior-Based Robots, IEEE Transactions on Robotics and Automation, 8(3), June 1992, 1992.

[McGill & Slocum 96] McGill M.E., Slocum J.: Das intelligente Unternehmen, Schaeffer-

Poeschel, Stuttgart, (The smarter organization, J.Wiley), 1996.

- [Millan 92] Millan J.del R.: Building Reactive Path-Finders through Reinforcement Connectionist Learning: Three Issues and an Architecture, in Neumann B.(ed.), Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI92), Wiley, Chichester, UK, pp. 661-665, 1992.
- [Nehmzow & Smithers 91] Nehmzow U., Smithers T.: Map-building using Self-Organizing Networks in "Really Useful Robots", in Meyer J.-A. & Wilson S.W.(eds.), From Animals to Animats, A Bradford Book, MIT Press, Cambridge, MA, 1991.
- [Newell 80] Newell A.: Physical Symbol Systems, Cognitive Science, 4(2), 135-184, 1980.
- 20

- [Prem 95] Prem E.: Grounding and the Entailment Structure in Robots and Artificial Life, in Moran F., et al.(eds.), Advances in Artificial Life, Springer, Berlin, LNAI 929, pp. 39-51, 1995.
- [Prem 96a] Prem

E.: Elements of an epistemology of Embodied Artificial Intelligence, in Mataric M.(ed.), Embodied Cognition and Action, AAAI Press, Menlo Park, Technical Report FS-96-02, pp. 97-101, 1996. (ftp://ftp.ai.univie.ac.at/papers/oefai-tr-96-04.ps.Z)

- [Prem 96b] Prem E.: New AI: Naturalness Revealed in the Study of Artificial Intelligence, in Dorffner G.(ed.), Neural Networks and a New AI, International Thomson Computer Press, London, pp. 13-28, 1996.
- [Rosen 85] Rosen R.: Anticipatory Systems, Pergamon, Oxford, UK, 1985.
- [Rosen 91] Rosen R.: Life Itself, Columbia University Press, New York, Complexity in Ecological Systems Series, 1991.
- [Rosen 93] Rosen R.: Bionics Revisited, in Haken H., et al.(eds.), The Machine as Metaphor and Tool, Springer, Berlin, pp. 87-100, 1993.
- [Simon 65] Simon H.A.: The Logic of Rational Decision, The British Journal for the Philosophy of Science, 63 (16), 169-186, 1965.
- [Simon 69] Simon H.A.: The Sciences of the Artificial, MIT Press, Cambridge, MA, 1969 (1981).
- [Stewart 92] Stewart T.A.: The Search for the Organization of Tomorrow, pp. 66– 73, Fortune, May 18, 1992.

- [Taylor 11] Taylor F.W.: The Principles of Scientific Management, Harper & Row, New York, 1911.
- [Tsotsos 95] Tsotsos J.K.: Behaviorist Intelligence and the Scaling Problem, Artificial Intelligence, 75(2), 1995.
- [Verschure & Pfeifer 93] Verschure P., Pfeifer R.: Categorization Representation and the Dynamics of System-Environment Interaction, in Meyer J.-A., et al., From Animals to Animats 2, A Bradford Book, MIT Press, Cambridge, MA, pp.210-217, 1993.