

Enhancing Service-oriented Holonic Multi-agent Systems with Self-organization *

José BARBOSA ^{a,c,d}, Paulo LEITÃO ^{a,b}

^a Polytechnic Institute of Bragança, Campus Sta Apolónia, Apartado 1134, 5301-857, Bragança, Portugal

^b LIACC - Artificial Intelligence and Computer Science Laboratory, R. Campo Alegre 102, 4169-007 Porto, Portugal

^c Univ. Lille Nord de France, F-59000 Lille, France

^d UVHC, TEMPO research center, F-59313 Valenciennes, France

Abstract

Multi-agents systems and holonic manufacturing systems are suitable approaches to design a new and alternative class of production control systems, based on the decentralization of control functions over distributed autonomous and cooperative entities. However, in spite of their enormous potential they lack some aspects related to interoperability, migration, optimisation in decentralised structures and truly self-adaptation. This paper discusses the advantages of combining these paradigms with complementary paradigms, such as service-oriented architectures, and enhancing them with biologically inspired algorithms and techniques, such as emergent behaviour and self-organization, to reach a truly robust, agile and adaptive control system. An example of applying a stigmergy-based algorithm to dynamically route pallets in a production system is also provided.

Key words: Multi-agent Systems, Holonic systems, Service-oriented systems, Self-organization

1 Introduction

Nowadays, manufacturing should adapt to strong changing conditions imposed by the markets. The greater variety of products, the possible large fluctuations in the demand, the shorter lifecycle of products expressed by a higher dynamics of new products are some challenges that manufacturing companies have to deal with to remain competitive. The traditional approach (i.e. centralised or hierarchical) to control complex systems (such as manufacturing systems or administrative systems) splits the overall problem into hierarchically dependent functions that operate within decreasing time-ranges, such as planning, scheduling and/or monitoring. This approach has produced interesting results and near optimal solutions, but only when hard assumptions are satisfied, for example, without the occurrence of external perturbations (e.g., urgent orders) or internal perturbations (e.g., machine breakdowns). However, a real (industrial) system is often complex with the

* This paper was not presented at any other revue. Corresponding author J. Barbosa. Tel. +351.273303065. Fax +273.273313051

Email addresses: jbarbosa@ipb.pt (José Barbosa), pleitao@ipb.pt (Paulo Leitão).

occurrence of deviations or perturbations. Consequently, this traditional approach rapidly becomes inefficient when the system must deal with such dynamic stochastic behaviour.

In this context, holonic manufacturing systems and multi-agent systems paradigms have introduced a new alternative to design control systems addressing flexibility and robustness. They are based on the decentralization of control functions over distributed autonomous and cooperative entities, being the overall control system achieved by the cooperation among these entities. The combination of these paradigms with new emergent technologies, such as service-oriented architectures, grid computing, wireless sensor networks and RFIID technology, allows reaching a new dimension regarding the interoperability and IT-vertical integration within enterprises and running in truly distributed and ubiquitous environments.

After combining multi-agent systems with these new emergent complementary paradigms and technologies, methods regarding the truly self-adaptation and self-organization of complex systems are still missing. Having this in mind, this paper explores the use of biologically inspired concepts, and particularly emergent behaviour and self-organization, combined with the service-oriented multi-agent systems to design innovative, flexible, adaptive, responsive and reconfigurable production control systems that better addresses the current industrial requirements.

The rest of the paper is organized as follows: Section 2 overviews the concept of holonic multi-agent systems and the Section 3 introduces the concept of service-oriented holonic multi-agent systems, based on combining the holonic multi-agent systems with the service-oriented principles. Section 4 discusses the introduction of self-organization mechanisms in such systems and exemplifies with the application of a self-organization mechanism for an agent-based control system. Finally, Section 5 rounds up the paper with the conclusions.

2 What Holonic Multi-agent Systems Can Offer

The multi-agent system (MAS) paradigm derives from the distributed artificial intelligence (DAI) domain, being characterised by decentralisation and parallel execution of activities based on autonomous entities, called agents. Despite the some definitions and interpretations for agents, a possible definition is [11]: “*An autonomous component that represents physical or logical objects in the system, capable to act in order to achieve its goals, and being able to interact with other agents, when it doesn't possess knowledge and skills to reach alone its objectives*”. An agent exhibits autonomy and cooperation, and may have reasoning and adaptability capabilities. In the manufacturing domain, an agent can represent physical resources, such as machines, robots, pallets and products, and logical objects, such as schedulers and orders.

A multi-agent system can be defined as a set of agents that represent the objects of a system, capable of interacting, in order to achieve their individual goals, when they don't have enough knowledge and/or skills to achieve individually their objectives [11] (note that each agent has a partial view of the system and none agent has a complete view of the system). The interaction between agents requires that the agents can understand themselves, using a proper agent communication language, ontologies and interaction protocols. Multi-agent systems suggest the definition of distributed control based on autonomous agents that account for the realization of efficient, flexible, reconfigurable and robust overall plant control, without any need for centralized control. These systems have the capability to respond promptly and correctly to change, and differ from the conventional approaches due to their inherent capabilities to adapt to emergence without external intervention [23].

A similar paradigm is the holonic manufacturing systems (HMS) that suggests the idea that manufacturing systems will continue to need a hierarchical structure besides the increased autonomy assigned to individual entities. HMS translates the concepts developed by A. Koestler for living organisms and social organizations into a set of appropriate concepts for manufacturing domain. In middle of sixties, Koestler introduced the word holon to describe a basic unit of organization in living organisms and social organizations [10], based on the H. Simon theories and on his observations. Simon observed that complex systems are hierarchical systems formed by intermediate stable forms, which do not exist as auto-sufficient and non-interactive elements but, on the contrary, they are simultaneously a part and a whole. Since in the domain of life, parts and wholes do not exist, Koestler proposed the word holon to represent this hybrid nature, being a combination of the Greek word holos, which means whole, and the suffix on, which means particle. Holons are simultaneously self-contained wholes to their subordinated parts and dependent parts when seen from higher levels, known as the Janus effect. This feature allows the structural development of production control systems through the encapsulation of manufacturing functions and components. As example, a holon representing a manufacturing cell is simultaneously the whole, encapsulating holons representing the cell resources, and the part, when considering the shop floor system.

A holon is an autonomous and co-operative entity that can represent a physical or logical activity, such as a robot, a machine a flexible manufacturing system, an operator or an order. The holon has information about itself and the environment, containing an information processing part and a physical processing part when the holon represents a physical device [22]. The HMS is a holarchy that integrates the entire range of manufacturing activities, i.e. a society of holons that can co-operate to achieve a goal, being organized in a hierarchical structure, cooperating to achieve the system goals, by combining their individual skills and knowledge. In HMS, the holons behaviours and activities are determined through the cooperation with other holons, in opposition of being determined by a centralized mechanism.

The multi-agent systems and the holonic manufacturing paradigms were developed under the same fundamental principles of autonomy and cooperation, exploring the distribution and decentralisation of entities and functions [11]. In spite of the similarity of the holon and agent concepts, some distinctions can be pointed out [11], namely in terms of origin (the agents have their roots in the computer science and the holons in the computer integrated manufacturing domain), in terms of concept (the holon is a concept and an agent is both a concept and a technology), in terms of integration of physical resources (agents represent software components and the holon supports the integration of physical resources) and in terms of decomposition (the holon can be composed by several lower-level holons, but an agent cannot).

3 Combining Holonic and MAS with Service-oriented Principles

The interoperability, the knowledge sharing during the interaction processes and the reconfiguration of the control components, by removing, adding or modifying the skills they provide, constitutes, amongst others, barriers to the easy development of such kind of systems (note that for example interoperability is crucial in distributed and heterogeneous environments, as manufacturing systems are).

Service-oriented architecture (SoA) paradigm is a way of building distributed systems [21], originally designed for electronic commerce and business, but progressively adopted in other domains. SoA is based on the concept of providing and requesting services. A service is a software piece that encapsulates the business/control logic or resource functionality of an entity that responds to a specific request. In such systems, the entities that want to offer their functionalities, encapsulate them as services and offer to the other entities by publishing them in a central repository. Using discovery mechanisms, service consumers find the services they need, and interact directly with those services. A main concern in service-based systems is how the services “play” together, emerging the concepts of orchestration and choreography. Orchestration is the practice of sequencing and synchronizing the execution of services, which encapsulate business or manufacturing processes [8]. An orchestration engine implements the logic for workflow-oriented execution and sequencing of atomic services, and provides a high-level interface for the composed process. Service choreography is a complementary concept, which considers the rules that define the messages and interaction sequences that must occur to execute a given process through a particular service interface.

The use of the service-oriented paradigm, implemented through web services technologies, enables the adoption of a unifying technology for all levels of the enterprise, from sensors and actuators to enterprise business processes [2]. In consequence, the service-oriented principles can be integrated with MAS to enhance some functionalities and to overcome some of its limitations, namely in terms of interoperability and IT-vertical integration. In spite of being promising paradigms to be used in industrial automation, and based on the same concept of providing a distributed approach to the system, MAS (and also HMS) and SoA present some important differences, namely in terms of autonomy and interoperability (see [19] for a deeply study). These differences highlight the complementary aspects of the two paradigms, suggesting the benefits of combining them [19].

This suggestion is not new since services are already part of the agents’ specification, e.g. in the FIPA (Foundation for Intelligent Physical Agents) specification [5], and agents are also present in standard documents of SoA methodologies [18]. Thereof, the under-considered elements (services in MAS and agents in SoA) are vaguely defined and have a more passive and customized role. Traditionally, the combination of MAS and SoA can be performed in different ways, as illustrated in Figure 1 [15]:

- Using gateways to translate the semantics from the agent world to the service world.
- Encapsulate single agents as services (agent-based services) and thus having a direct access to other services.
- Using the concept of service-oriented agents that not only share services as their major form of communication, but also complement their own goals with external provided services.

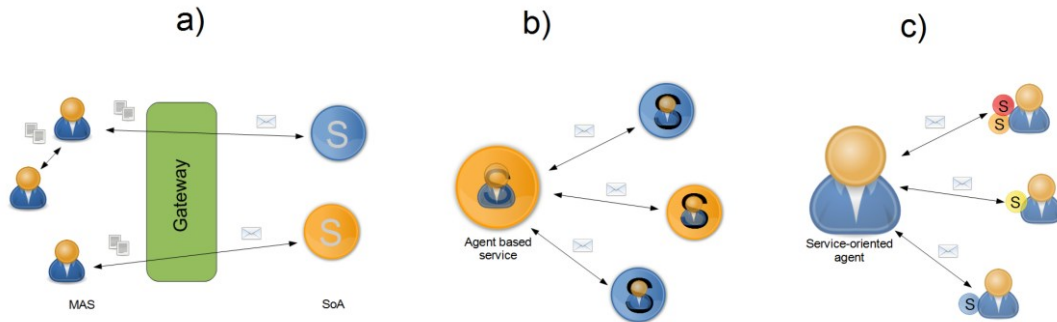


Fig. 1. Three commonly used approaches for integrating SoA and MAS (adapted from [15])

This last option, introduced by [15] and illustrated in Figure 1.c), is characterized by using a set of distributed autonomous and cooperative agents that use the SoA principles, i.e. oriented by the offer and request of services, in order to fulfil industrial and production systems goals. The achieved service-oriented multi-agent systems (SoMAS) approach is different from the traditional MAS mainly because agents are service-oriented, i.e. individual goals of agents may be complemented by services provided by other agents, and the internal functionalities of agents can be offered as services to others agents [15].

An example of using the SoMAS approach is a mediator (agent) of a conveyor that transports pallets, as illustrated in Figure 2. The agent provides a service (called *transfer pallet*) that reflects its own functionality and skills. To achieve its goal, that it is to transfer pallets, the agent needs to request a transfer service from a neighbour device (conveyor or machine) that provides the necessary capability of transferring out the pallet from the conveyor.

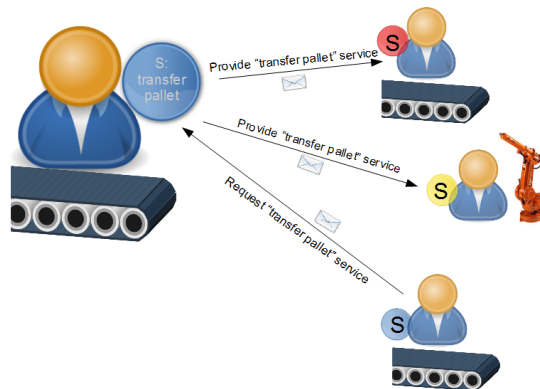


Fig. 2. Example of a service-oriented agent

Besides the combination of service-oriented principles with holonic and multi-agent systems, also interesting is the articulation with others technological domains providing foundations to support distributed and ubiquitous environments, such as grid computing, wireless sensor networks and RFID technology.

4 Self-organization and Self-Adaptation Mechanisms

The approach that combines MAS and SoA principles allows the achievement of significant advantages to address the current production requirements, but it still lacks the capability to adapt and evolve to unexpected pressures from the environment, product fluctuation and internal/external disturbances. The theories and

mechanisms found in biology, such as emergence and self-organization, may constitute useful insights to design more adaptive and re-configurable systems. The challenge is to go back to the foundations of HMS and MAS theories and combine them with biologically inspired techniques to develop more adaptive and evolvable systems, which can be easily deployed into real environments.

4.1 Emergent biologically inspired theories and fields

Biology and nature provide a plenty of simple mechanisms that solve complex problems, constituting suitable sources of inspiration to support the design of better solutions for self-adaptive and evolvable complex systems. A recent article published in the National Geographic magazine reinforces this idea, stating that “the study of swarm intelligence is providing insights that can help humans to manage complex systems” based on the idea that “a single ant or bee isn't smart but their colonies are” [16].

The concept of swarm intelligence, found in colonies of insects, can be defined as “the emergent collective intelligence of groups of simple and single entities” [3], reflecting the phenomenon of emergency where the global system behaviour emerges from a multiplicity of non-linear interactions among the individual entities. Swarm intelligence offers an alternative way of designing intelligent, complex systems, in which the traditional centralized control is replaced by a distributed functioning where the interactions between individuals lead to the emergence of “intelligent” global behaviour, unknown to them [3]. Examples of swarm intelligence include ant colonies, bird flocking, fish shoaling and bacterial growth [16]. A truly multi-agent system matches these insights, following the swarm intelligence concept. However, in opposition to the biology, where usually these entities exhibit very limited cognitive skills, agents may be enriched with intelligence capabilities, and especially learning mechanisms, to allow the development of more powerful systems.

The theory of complexity is a field of study that tries to characterize complex systems, also providing interesting insights to be considered in the multi-agent systems domain. A complex system is a *system composed of interconnected parts that as a whole exhibit one or more properties (behaviour among the possible properties) not obvious from the properties of the individual parts* [9]. These complex systems have emergent properties that can't be reduced to the behaviour of separated entities (i.e. the behaviour of separate entities does not explain the global behaviour of the system), as stated in [20] “*the complexity of a system increases with the number of distinct components, the number of connections between them, the complexities of the components, and the complexities of the connections*”. This means that the analysis of these systems can't be made by classical methods, e.g. Newtonian mechanics, which are essentially reductionist: divide the global problem by smaller problems, simpler to solve, and to obtain the global one by adding the small solutions. On contrary, in complex systems the global behaviour is greater than the simple sum of all the small behaviours [7]. Examples of complex systems can be found in an organizational society, an ecosystem, the weather or in a manufacturing system.

Self-organization is another promising theory found in biology that can be simply defined as the autonomous adaptation to the dynamic evolution of the environment. Self-organization is not a new concept, being applied in different domains such as economics, computing and robotics. Self-organizing systems do not follow a rigid and estimated organization, evolving, without a central authority in command, through a non-linear and dynamic process with a constant optimization of the individuals' behaviour. Stigmergy is a form of self-organization, involving an indirect coordination between entities, where the trace left in the environment stimulates the execution of a subsequent action, by the same or different entity. As an example, ants exchange information by depositing pheromones on their way back to the nest when they have found food.

A similar concept is self-adaptation, which can be seen as the capability of an entity to change its behaviour depending on the external conditions. Ants, for example, when searching for food, find paths between the nest and food sources that could be optimal or near optimal [6]. In case of an obstacle (e.g. a rock) is laid down on a path, what is the ants behaviour? Naturally, the first's ants try to find newer paths adapting themselves to the new external conditions. Other self-adaptation mechanism is found on the Darwinian evolution theory where the species suffer from mutation, recombination and selective reproduction of the fittest techniques, giving origin to evolutionary algorithms (e.g. genetic algorithms, evolutionary strategies and genetic programming) that are being used, e.g., in evolutionary robotics.

The degree of efficiency of the self-organization capability is strongly dependent on how the learning mechanisms are implemented. In the design of self-organized systems the key issue is to define powerful intelligence mechanisms, not only including static intelligence mechanisms but also learning capabilities, that enable the system to improve its behaviour in the future as result of its experience. For this purpose, the embodied intelligence concept, associated to the artificial life field, assumes a crucial role. This concept suggests that intelligence requires a body to interact with [17]; in this case, the intelligent behaviour emerges from the interaction of brain, body and environment. This illustrates the existence of new fields of computer science, such as artificial life and evolutionary computing, that try to mimic some biological concepts. Namely, the artificial life [1] is a discipline that studies the natural life in artificial environments, e.g. through simulations using computer models, in order to understand such complex systems. Note that artificial life is not similar or is not included in the artificial intelligence field: the last one is mostly related to the perception, cognition and generation of actions, and the former one focuses on evolution, reproduction, morphogenesis and metabolism processes [4].

Holonic systems embeds already some of this biological inspired concepts such as the ability of evolution and growth to satisfy increasingly complex and changing needs by creating stable “inter mediate” forms, which are self-reliant and more capable than the initial systems.

4.2 Experimental application of a self-organization mechanism

Incorporating self-* mechanisms, such as self-healing, self-adaptation and self-organization, into the agents behaviour can greatly contribute to increase the system performance, flexibility and robustness. Self-healing is another important concept that can be defined as the capacity to diagnose deviations from normal conditions and take proactive actions to normalize them and avoid service disruptions [12]. Self-healing comprises two fundamental steps: self-diagnoses and self-repair. Self-healing and self-repair are standard actions in nature and can be observed in numerous cases; observe for example the capability that the human body skin has to automatically recover from injuries. An example of the self-healing in manufacturing engineering is a grip malfunction. In this situation the agent detects that the grip is not functioning correctly, performs a self-diagnosis and implements self-repair actions, e.g. resetting the configuration parameters of the grip; if the agent isn't able to self-repair it can ask external support.

An experimental case study was used to apply a self-organization mechanism to develop a self-organized agent-based control system. The biological inspired concept of stigmergy was used to improve the routing problem of pallets in a production system, where pallets must dynamically find the best path to reach their goals (e.g. workstations and exit of the system) [14]. The production system corresponds to the FlexLink® Dynamic Assembly System (DAS) 30, illustrated in Figure 3, is composed by two levels: the upper level has 9 conveyors and the lower level has 2 conveyors being these two levels joined by 2 lifters that also serve as the input (i.e. left conveyor) and the exit (i.e. right lifter) of the system. In the upper level some of the conveyors can convey the pallets in more than one direction, being here the part where the routing decision is more critical.



Fig. 3. FlexLink DAS 30 system (located at the Schneider Electric GmbH in Seligenstadt, Germany) [13]

The control system behaviour, using the multi-agent system principles and the self-organization concept, was modelled using the NetLogo modelling and simulating platform. The user interface is illustrated in Figure 4.

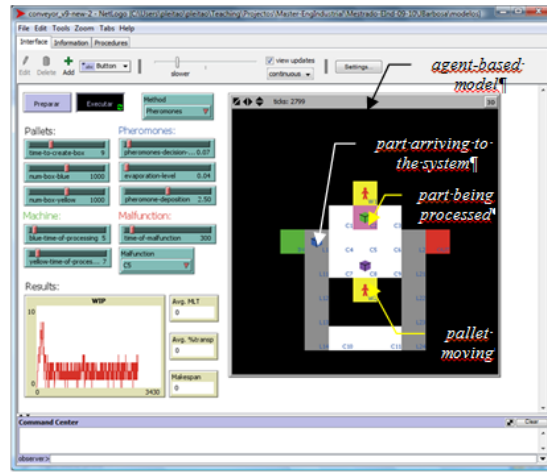


Fig. 4. NetLogo interface for the agent-based production control model

Taking again the previous example of the ants' behaviour, let's recall on the way that they use to communicate between each other, known as stigmergy. Ants lay down a chemical odour, known as pheromone, that can be sense by others ants to know where to go. This way of communication gives the ability to organize them and to react to unexpected external conditions.

In this example, a bio-inspired algorithm based on the stigmergy concept was implemented to route the pallets in the conveyor system. Pallet agents while moving through the conveyors store their movements. When a pallet reaches the desired goal it first removes the duplicated routes and then updates the correspondent pheromone level in each travelled conveyor (note that each conveyor (i) has a list of pheromone levels to each possible goal (g)). The removal of the routes loops ensures that loops are not valorised more than once. The update of the pheromone levels takes into consideration the amount of time taken by each pallet to reach their goal (t_i) valorising in this way the most appropriate route (i.e. shortest one). A mathematical representation of the above considerations can be found in equation 1, where the r parameter is the reinforcement pre-adjusted value.

$$\begin{cases} P_{i,g}^k = P_{i,g}^{k-1} + \frac{r}{t_i} & , \text{updated when the pallet arrives to the goal (reinforcement phase)} \\ P_{i,g}^k = P_{i,g}^{k-1} - e & , \text{updated in each simulation step (evaporation phase)} \end{cases} \quad (1)$$

The natural phenomenon of evaporation (e), which occurs in nature, was also implemented by decreasing all levels of pheromone by a pre-set value (see equation 1). This guarantees that during the evolution of the system, the longest paths tend to disappear (i.e. due to the lack of pheromone re-enforcement) and the shortest paths arise, which indicates, for each moment, the best route to a given objective. The system self-adapts in a way that new paths emerges when a disturbance occurs, e.g. a conveyor malfunction, due to the evaporation mechanism and to the lack of pheromone deposition in the direction of the malfunctioned conveyor.

In the implemented approach the correct tuning of the pheromone parameters plays a crucial role, i.e. the importance of the initial conditions of the system, recalling the concept of butterfly effect or Lorenz attractor found in chaos theory. The butterfly effect states that there is a possibility that the wings flap by a butterfly in China can produce a large storm in New England; in a similar way, minimal changes in initial parameters on the Lorenz attractor conducts to very different patterns. In fact, like in the butterfly effect or the Lorenz attractor, if wrong parameters are defined, the stigmergy-based algorithm can behave in a wrong manner, with the system doesn't exhibiting self-organizing or self-adaptation characteristics. In this approach the pheromone parameters

should be dynamically adjusted to the rate of the time arrival of pallets, i.e., the values should be different when the rate is slow or faster.

This stigmergy-based approach gives the system another kind of flexibility and adaptability that is related to the fact of system reconfiguration. Suppose that a conveyor or machine breakdown, or instead of only one system, another system is coupled to the exit of this system. What will be the expected behaviour of this change? The agents, like in the first case, will dynamically adapt to the changes (e.g. broken conveyors, traffic jam or new system) and they will find the best paths to their objectives by re-routing dynamically their paths.

5 Conclusions

This paper discusses the advantages of combining multi-agent systems and holonic systems paradigms with complementary paradigms, such as service-oriented architectures, to design more powerful production control systems, based on the decentralization of control functions over distributed autonomous and cooperative entities. Additionally, aiming to reach a truly self-adaptive and evolvable system, the paper discusses how to enrich the system designed using the previous paradigms with biological inspired techniques, such as emergent behaviour and self-organization, taking into consideration that in nature very complex and adaptive systems are implemented by using very simple behaviours and mechanisms. In this topic, a special attention should be taken in the study of the biological inspiration to solve a determined problem. It may happen that not all of the biological inspired behaviours could be important to solve a given problem and the research must also take into consideration what is really needed, discarding what is not important. This is a crucial issue since the imitation of unnecessary behaviours complicates the implementation and could arise undesired performances of the solution.

Some considerations should also be made regarding the differences between the nature and engineering worlds. In nature, there is time and space to recover after failures, which means, that if something isn't made right at first time there is always more opportunities to get it right. On other hand, in the engineering world, things must be done right at the first time and failures must be avoided at all costs. Another important difference is that the main goal in nature is to guaranty the species survival while engineering has a multiplicity of very specific goals (e.g. costs reduction, ensure quality and customization).

The application of a stigmergy-based algorithm to dynamically route pallets in a production system was described to illustrate how to apply these bio-inspired techniques and the improvements we can get. From the experience of applying these concepts, the expected advantages of such self-organised architectures are related to agility: in the short term, these architectures are reactive and, in the long term, they are able to adapt to their environment. Of course, optimality and context adaptability is not guaranteed, which implies that in future research it is necessary to study how to achieve optimization in distributed and decentralized solutions.

6 References

- [1] C. Adami. Introduction to Artificial Life. Springer-Verlag, 1998.
- [2] A. Bepperling, J. M. Mendes, A. W. Colombo, R. Schoop and A. Aspragathos. A Framework for Development and Implementation of Web service-Based Intelligent Autonomous Mechatronics Components. Proceedings of the IEEE International Conference on Industrial Informatics, Singapore, pp. 341-347, 2006.
- [3] E. Bonabeau, M. Dorigo and G. Theraulaz. Swarm Intelligence: from Natural to Artificial Systems. Oxford University Press, 1999.
- [4] R. Brooks. The Relationship between Matter and Life. Nature, 409: 409-411, 2001.
- [5] FIPA Abstract Architecture Specification. Standard of the Foundation for Intelligent Physical Agents, 2002 Available at <http://www.fipa.org/specs/fipa00001>
- [6] S. Goss, S. Aron, J.-L. Deneubourg and J. M. Pasteels. Self-organized shortcuts in the Argentine ant. Naturwissenschaften, 76:579-581, 1989
- [7] J. Holland. "Emergence: from Chaos to Order". Oxford University Press, 1998.
- [8] F. Jammes, H. Smit, J. L. Martinez Lastra and I. Delamer. Orchestration of Service-Oriented Manufacturing Processes. Proceedings of the 10th IEEE International Conference ETFA, vol. 1, pp. 617-624, 2005.
- [9] C. Joslyn and L. Rocha. Towards semiotic agent-based models of socio-technical organizations, Proceedings of the AI, Simulation and Planning in High Autonomy Systems (AIS 2000) Conference, pp. 70-79, 2000.

- [10] A. Koestler. *The Ghost in the Machine*. Arkana Books, London, 1969.
- [11] P. Leitão. Agent-based Distributed Manufacturing Control: A State-of-the-art Survey. *International Journal of Engineering Applications of Artificial Intelligence*, 22(7): 979-991, 2009.
- [12] P. Leitão. Holonic Rationale and Self-Organization on Design of Complex Evolvable Systems. *Proceedings of the 4th International Conference on Industrial Applications of Holonic and Multi-Agent Systems (HoloMAS'09)*. V. Marik, T. Strasser and A. Zoitl (eds), LNAI 5696, Springer, pp. 13-24, 2009.
- [13] P. Leitão, J. Alves, J. M. Mendes, A. W. Colombo (2010), "Energy Aware Knowledge Extraction from Petri Nets Supporting Decision-making in Service-oriented Automation", *Proceedings of the IEEE International Symposium on Industrial Electronics (ISIE'10)*, Bari, Italy, 4-6 July, pp. 3521-3526.
- [14] P. Leitão and J. Barbosa. Modelling and Simulating Self-Organizing Agent-based Manufacturing Systems. To appear in the *Proceedings of the 36th Annual Conference of the IEEE Industrial Electronics Society (IECON'10)*, 2010.
- [15] J.M. Mendes, P. Leitão, F. Restivo and A.W. Colombo. Service-oriented Agents for Collaborative Industrial Automation and Production Systems". *Proceedings of the 4th International Conference on Industrial Applications of Holonic and Multi-Agent Systems (HoloMAS'09)*. V. Marik, T. Strasser and A. Zoitl (eds), LNAI 5696, Springer, pp. 1-12, 2009.
- [16] P. Miller. *The Genius of Swarms*. National Geographic, 2007.
- [17] R. Pfeifer and C. Scheier. *Understanding Intelligence*. Cambridge MA, MIT, 2001.
- [18] Reference Model for Service Oriented Architecture 1.0. OASIS Standard. October 12, 2006. Available at <http://docs.oasis-open.org/soa-rm/v1.0>
- [19] L. Ribeiro, J. Barata and P. Mendes. MAS and SOA: Complementary Automation Paradigms. *IFIP International Federation for Information Processing*, vol. 266, Springer Boston, pp. 259-268, 2008.
- [20] K. Richardson. *Managing Organizational Complexity: Philosophy, Theory And Application*, Information Age Publishing, 2005.
- [21] S. Ross-Talbot. *Orchestration and Choreography: Standards, Tools and Technologies for Distributed Workflows*, Pi4 Technology, London, UK and W3C, Geneva, Switzerland, 2005.
- [22] M. Winkler and M. Mey. *Holonic Manufacturing Systems*. European Production Engineering, 1994.
- [23] M. Wooldridge. *An Introduction to Multi-Agent Systems*. John Wiley & Sons, 2002.