Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Intelligent Product = Intelligent Agent + Intelligent Being

Paul Valckenaers, Bart Saint Germain, Paul Verstraete,
Jan Van Belle, Hadeli and Hendrik Van Brussel

Mechanical Engineering Department, K.U.Leuven, Leuven, Belgium (e-mail Paul.Valckenaers@mech.kuleuven.be)

Abstract: The notion of an intelligent product calls for research into *suitability for integration* as it aims for coordination and integration on an unprecedented scale. This is particularly challenging where it affects the core businesses of companies. Moreover, these integration issues go well beyond IT interoperability. This paper presents a novel concept—the intelligent being—as a means to achieve suitability for integration at this scale. The paper argues for a major role of the intelligent being in production and logistic systems analogous to the undeniably prominent role of maps in navigation systems. The concept of an intelligent being contributes to the design and development of complex adaptive systems by enabling integration and coordination at an unprecedented scale: the corresponding reality is the limit.

1. INTRODUCTION

"Intelligent agents and multi-agent systems" constitute an important research domain (AgentLink.org). In this domain, the decision-making capabilities of an agent are central. This is reflected in the linguistic roots of the word agent. Its Latin origin—agere—means to act. In English, an agent represents a person or organisation to act on his/her/its behalf. An agent mediates, negotiates, manages… Research results reflect this and emphasize goal-oriented reasoning, game theory, negotiation protocols, etc.

However, real-world applications need more than decision-making. This has prompted researchers to put forward the environment, addressing the non-agent dimensions in a multi-agent system, as a primary design concern (Weyns *et al.*, 2007). Yet, addressing non-agent dimensions explicitly merely begins answering the key complexity issue: *how to avoid and resolve conflicts in software design?* Importantly, answers must address the integration of software that has been developed without coordination by different parties at different times and locations.

Part of the answer is given by Shakespeare's most famous writing: "To be or not to be…" The Latin word essere, which translates in to be, precedes agere in a fundamental manner. To exist is a necessary precondition for being able to act. Today, much remains to be explored in software development regarding essere, which is the origin of the words essence and essential. This renders the investigation of the intelligent being or I-Being—a concept based on essere—a promising research direction.

Note however that intelligent beings are not some form of artificial existence; they add artificial intelligence to something that already exists. Therefore, the creation of an I-being comprises the identification of the corresponding reality—the being—followed by the development and maintenance of the corresponding intelligence.

This paper starts with a discussion of the historic forerunners for software beings: maps. Subsequently, the discussion reveals the additional functionalities and services made possible by modern ICT relative to the paper and ink of conventional maps. Next, intelligent products are presented as an aggregate of an intelligent being, an intelligent agent and the actual product. An appraisal of related work by others and conclusions wrap up the discussion.

2. MAPS: THE OLDEST I-BEINGS

Historically, the word intelligence denoted the gathering of information—as in Central Intelligence Agency. In this sense, maps add intelligence to some part of the world, thus creating an intelligent being. Maps share a number of key properties with the software I-Beings put forward by this paper:

- Maps mirror to some part of the real world: roads, depths in the water near the coast, the topology of a mountain range, the location of conduits in a sewage system, etc.
- Maps provide useful information (services, functionality) about a corresponding segment of the real world. Importantly, most of the map creation is choice-free where reality dictates what must be on the map. Choices are confined to representation issues (symbol selection), range and scope.
- The corresponding part of the real world is sufficiently stable (i.e. slow-changing).
- A map is robust relative to its correctness, completeness and usage/interpretation. Small errors do not render a map useless. And, if a map fails to provide all relevant information, it suffices to add this. Importantly, such

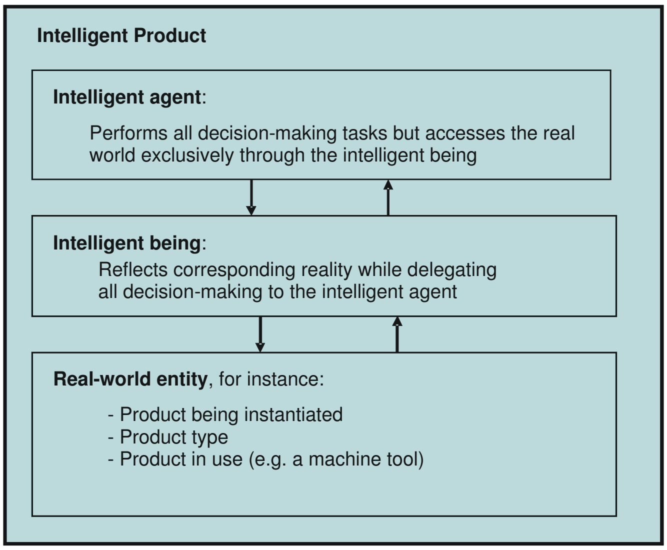

Fig. 1. Structure of an Intelligent Product

extension activity does not invalidate what already exist and the required effort is unaffected by the existing map's size and attributes.

- Maps are conflict-free. Information can be missing, outdated, wrong, and representations can be differing. But, there will be no conflicting choices embedded in maps relative to the real world. When multiple maps are available, any contradiction can be resolved by observing the corresponding reality. People are able to use multiple maps, retrieving information from the most convenient source.
  In contrast, traffic rules can be conflicting and it is impossible to use them simultaneously or merge them without undoing embedded choices. Importantly, such undoing triggers the need to revalidate the set of traffic rules (i.e. to check whether the new version is consistent and coherent). I-Beings have no such issues.
- Maps offer only interim solutions. Some basic navigation and map-reading skills are required to generate complete answers/solutions.

Because maps reflect reality, they are amongst the most successful, complicated and feature-rich artefacts, created by men, that have no integration problems worth mentioning. This is the prime motivation to develop intelligent beings next to intelligent agents.

Like maps, I-beings will be interim solutions only. They do not replace agents; they provide additional services that, among others, significantly augment the services of internal agent components such as their embedded and private utility functions and world models. Overall, maps constitute a good example of where I-Beings are going. The main difference is the information technology that is available to build useful cyber-brains for suitable stable parts of reality.

3. INTELLIGENT BEINGS

Figure 1 shows the overall structure of an intelligent product. By definition, a conflict with an I-Being, within an intelligent product, is simultaneously a conflict with reality[1]. The designer of an I-Being never introduces restrictions—in the application domain—that are absent in the corresponding reality. This explains the robustness and *suitability for integration* of I-Beings. Reality provides shelter because it is coherent and consistent. The agents take all the decisions.

Nonetheless, conflicts at the level of information representation may persist. For instance, one map may use dashed lines whereas another map employs a colour coding. Airplane pilots and air traffic controllers need to speak the same language. Providing the necessary adapters to overcome such integration problems is rather trivial. These minor conflicts are not considered further in this paper.

The hard constraint for the design and development of an intelligent being is to preserve a suitable shelter provided by a real-world counterpart. In this respect, state-of-the-art ICT offers significantly more opportunities to transform a "being" into an "intelligent being" than the paper, ink and printing press available to Mercator and his colleagues some centuries ago. This section discusses these novel opportunities and their implications.

3.1 Benefits of Modern Information Technology

The superior expressive power of ICT permits reflecting reality without unwanted distortions. The deceiving pictures of world maps, e.g. exaggerating the surface near the artic, can be avoided without the inconvenience of a spherical medium. ICT supports numerical expressions, aggregation, specialization, association, symbolic manipulation, etc. The information representation, its structure and the convenience offered to its users, represent a quantum leap over map technologies.

Intelligent beings offer services beyond the presentation of simple information concerning their corresponding part of reality. For example, intelligent beings may provide virtual sensors to compensate the temporary lack of sensor data or to reduce the cost of sensing. An intelligent vehicle may compensate a momentary absence of the GPS signal by means of its inertial platform and route matching software.

I-Beings reside in computer networks and provide access to associated information. An intelligent being corresponding to a machine tool may provide (controlled) access to the user

[1] A simple tale communicates more effectively than intricate explanations. This story recaps a conversation over the maritime radio waves:
- *Intelligent being*: "Ship ahoy. This is CL273. You are on a collision course with us. Please change your heading."
- *Intelligent agent*: "This is USS129. This is the United States Navy. You change your course."
- *Intelligent being*: "This is Canadian Lighthouse 273…"

manual, the maintenance log, the CAD drawings, etc. Likewise, this I-being provides links to the other I-beings corresponding to its constituents (spindle, turning table, tool magazine), its neighbours (transport system delivering and fetching products on pallets), its environment (section of the factory infrastructure), its visitors (products on pallets on the turning table, repair men) and residents (tools in the magazine, operators). Furthermore, ICT may turn the intelligent being into an information storage location for other software entities (a cyber version of post-it notes) resulting in a stigmergic infrastructure. Valckenaers *et al.* (2007) provides a more exhaustive discussion.

Sensing, tracking and tracing technologies enable automated updating and storage activities. E.g. an intelligent windmill measures its rotation speed and output continually. It may also receive a suitable weather forecast to estimate the feasibility of near-future power production. In addition, the intelligent mill may log this information for analysis purposes and possibly for the updating of a self-model.

Tracking increases the collection of parts of reality that are capable of providing suitable shelter immensely. If the intelligent being is able to observe and track changes in the corresponding reality, and to adapt itself correspondingly, adequate shelter can be guaranteed in the face of moderate dynamics. As a consequence, intelligent beings can be developed and maintained in much more situations than simple maps. Furthermore, tracking allows the composite I-beings to adapt by replacing subsystems, mirroring the corresponding reality (Simon, 1990). Here, on-line, real-time and networked ICT makes a decisive contribution, enabling substantial scaling.

3.2 Advanced Services from I-beings

Researchers only start to understand the functionalities and services that I-being technology may deliver. An important category of advanced I-Being services provide short-term forecasts. These forecasts account for the interactions amongst the intelligent beings, the intelligent agents and the real-world entities. To achieve this feat, the I-Beings include the mental state of their neighbouring agents in their corresponding reality.

Developers cannot design nor encode assumptions about these agents—more precisely, about the decision-making mechanisms—into an I-Being. This destroys its shelter and effectively renders it hard-to-integrate. Yet, assuming that there will be an agent, or equivalent entity that takes decisions, will not destroy the effective shelter from reality. It suffices not to make assumptions about the agents with the exception of the assumption that they exist.

The shelter of an intelligent being is preserved through a relentless delegation of decisions to the affected agents. I-Beings, mirroring the relevant dimension of the corresponding reality, are able to place these agents in fictitious situations and solicit their decisions. In this way, the I-Beings are able to project what the future will bring, given the installed decision-making mechanisms (whatever these mechanisms are).

Concerning stability and reliability of the forecasts, it is advisable to focus on reflecting the more stable aspects of an agent's mental state, or in other words, the serious commitments of the agent. Nonetheless, for the intelligent being it suffices to reflect how much commitment exist, not to enforce it. Evidently, the agent society may enforce commitments to benefit from sufficiently reliable predictions.

A practical implementation of this conceptual scheme needs to address additional issues (covering the search space of the possible solutions, commitment to declared intentions, etc.). Valckenaers *et al.* (2003) provide an elaborate discussion. Most importantly, even such practical and computationally efficient implementations are able to safeguard to shelter by reality for the embedded I-beings. The use of a delegate MAS is a key element in this respect (Holvoet *et al.*, 2006) (Parunak *et al.* 2007).

4. INTELLIGENT PRODUCTS

An intelligent product cannot be wholly implemented by an intelligent being. Indeed, intelligent products need to make choices that expose them beyond their corresponding reality, which is a physical product instance or an immaterial product type. Therefore, this paper proposes an internal structure for an intelligent product comprising both an I-Being and an intelligent agent (and some real-world entity).

4.1 Internal Structure of Intelligent Products

Figure 1 depicts the internal structure for an intelligent product. It comprises the real-world entity, an intelligent being and an intelligent agent. Note that all three components may be aggregates, composed of more primitive components.

The real-world entity can be:

- The physical product instance in its nominal state.
- The physical product instance in some interim state during production.
- The product instance that no longer exists (recycled). This intelligent product provides an archiving reference.
- Primitive or composite.

Moreover, this shelter-providing reality often is virtual:

- The product instantiation activity (production).
- The product type (a production recipe).
- Compositions of all the above.

The corresponding intelligent beings provide as much service and user functionality as possible without losing the shelter from the corresponding reality and without losing 'locality.' The latter means that an intelligent being receives shelter from a section of reality that is long-lived and, often, replicated in many places. This property provides critical mass for the software development concerning its user community.

The intelligent agents provide decision-making functionality. Simple implementation may substitute agents with rules and procedures triggering the need for human supervision, which conceptually adds up to an intelligent agent (with a human brain). The forecasting services from intelligent beings require a computer-based avatar for such human agents to cope with the high frequency at which decisions in fictitious realities have to be made. In theory the human agent may answer all queries from the I-Beings but this is economically not viable.

Intelligent agents exclusively access the real world through the I-Beings. If necessary, some place-holding stubs may be provided in initial implementations. Direct access by the agents to the real world diminishes the services from the I-Beings significantly (Valckenaers *et al*. 2007).

4.2 Intelligent Product Categories

The intelligent products, which render a real-world entity intelligent, belong to two categories: intelligent product instances and intelligent product types.

There will be intelligent products that correspond to product instances and their life cycle—product instantiation, usage, disposal/recycling. These intelligent product instances manage activities. In the PROSA architecture, the order holon implements such intelligent product instances (Valckenaers *et al*., 2005).

The second category comprises intelligent products that correspond to product types. They are recipe experts concerning the instantiation of physical instances. In the PROSA architecture, the product holon implements these intelligent product types. Moreover, each order holon is associated to a product holon and delegates all process-related matters to this product holon.

Furthermore, PROSA also distinguishes a resource holon. This corresponds to a product instance in its usage phase (McFarlane *et al*. 2002). Therefore and conceptually, such intelligent resources are intelligent product instances of a particular kind of product: machines, tools, workers, materials, etc.

What's more, recognizing that intelligent product instances correspond to activities, it becomes attractive to extend the intelligent product concept to intelligent services without a physical product. Intelligent maintenance is an example. This minimizes the differences between normal production and maintenance activities, resulting in leaner designs for plant control systems. Logistic activities, warehousing activities also offer possibilities. Generic solutions become much more doable.

4.3 Motivation for the Subdivision

The above sections propose and present the internal structure of an intelligent product. It also discusses the corresponding reality for the embedded intelligent beings. This section addresses the motivation for the clean separation between intelligent beings and agents.

Section 3 already reveals why intelligent beings are attractive software components concerning integration. Since the neighbouring intelligent agents fail to benefit from the shelter provided by reality, and internal structure has to ensure that software maintenance on the agent subsystem is well-supported. In particular, updating the embedded agent must not affect the intelligent being.

There exists however a motivation to separate agents and beings more radically: they have different habitats. The intelligent being co-resides with the corresponding reality. For instance, the provider of a machine tool is able to deliver the same intelligent being (design) with every machine installation.

In contrast, the intelligent agent's habitat is determined by the embedded decision mechanisms. This does not coincide with the corresponding reality. For example, managing the production activity of a single product type requires different decisions when its context evolves from a supplier-dominated market to a customer-dominated market. The same holds when this production activity migrates from a job shop into a flow shop.

This situation is shared by many man-made systems. Birds have integrated lift and propulsion, produced through biological reproduction. Aircraft have distinct organisations that produce wings and engines. Intelligent products have similar interests. The organizations delivering respectively intelligent beings and intelligent agents have radically different markets. An intersection of these markets, corresponding to an intelligent product, is several orders of magnitude smaller. Thus, a strong separation of these two subsystems within an intelligent product provides critical mass in the respective user communities. Without such a separation, applications will be confined to high-investment, high-return activities (e.g. steel plants).

5. RELATED WORK

This section discusses work by others. Nowadays, intelligent products and product-driven manufacturing are a prominent research topic because of the emergence of enabling technologies, often referred to as RFID. In addition, this section discusses research targeting the laws of the artificial, where these laws share the inevitability of the laws of nature. In addition, agent technologies are covered.

5.1. Intelligent products

Research on intelligent products addresses various concerns and viewpoints. McFarlane addresses manufacturing control, providing a demonstration how Auto-ID systems facilitate the development of a Holonic assembly cell (McFarlane *et al.,* 2003). This basically reveals the potential of the intelligent product concept without providing reusable components in a systematic manner. Wong addresses supply chain issues, and provide a systematic list of services and functionalities that intelligent products may provide (Wong *et al.*, 2002). It does not provide a system architecture in which these can be deployed in a coherent and consistent manner while supporting evolution and adaptation. Both results are setting the first beacons regarding the usage of intelligent products. An interesting point in (McFarlane *et al.,* 2002) is to consider intelligent manufacturing resources (equipment) as intelligent products in the usage phase within their life cycle, which this paper readily adopts.

In contrast, Främling focuses on the data management issues, and combines agent-based concepts with the intelligent product (Främling *et al.*, 2006). The research addresses the naming issues (the URL and ONS for intelligent products), the encoding of information (XML for intelligent products), as well as the communication, security and implementation issues. It addresses the information technology dimension, thus complementing the work by McFarlane's team.

Completing the picture are research results that investigate the impact and the novel possibilities of the intelligent product on the planning and decision-making technologies and architectures. Relative to Auto-ID research, these results focus on business management and planning aspects and less on detailed control or information handling. A key issue that emerges is the level of details, the resolution, with which the system looks at the real world (Schuh *et al.*, 2007). Auto-ID enables, compared to yesterday, a high resolution production management that is able to operate in real time. These operating conditions impose new demands on planning systems and call for an innovative distribution of tasks and responsibilities across management and control hierarchies.

Overall, the intelligent product concept bestowed a lot of research activities, which are nonetheless only a beginning. Much remains to be investigated in the domain of product-driven manufacturing and product life-cycle management.

Relative to the available research results, this paper presents a complementary view and contribution. It addresses integrate-ability aspects, which are complementary to the information-handling issues addressed by (Främling *et al*., 2006). And, it is complementary to research focused on decision-making management and control functionality (Schuh *et al*., 2007; McFarlane *et al*., 2003).

5.2. The Science of the Artificial

Herbert Simon pioneered the development of the counterpart of the laws of nature for the artificial (Simon, 1990). Inevitability is the canon in his undertaking. Most manmade artefacts are one possibility out of the vast number of variations that are possible. The laws of the artificial must not suffer such deficit. If a designs violates these laws, under the conditions of their applicability, this design must fail.

The concept of an intelligent being was inspired by Simon's quest. Not nature but artefacts are the subject of study. Holonic systems and the autocatalytic set help to identify which artefacts are sufficiently stable to make the effort of developing an I-Being worthwhile. For such artefacts, the intelligent being incorporates its laws of the artificial. The conditions under which these laws apply are the presence of the corresponding artefact in the real world. This artefact may be a virtual entity but it is manmade and evolving slowly such that the intelligent being may track it.

Overall, this pursuit for inevitability has not (yet) become a mainstream research activity in the intelligent product research area. Major obstacles are that it requires expertise in multiple disciplines and typically solves a problem once-and-forever. The latter implies the absence of an autocatalytic set in the research community itself (i.e. other researchers cannot develop alternative solutions and publish them referring to the original work). The concept of an intelligent being may suffer from this problem at the level of its own concept. However, the range of intelligent beings that can and should be developed is endless and has the potential to gather critical mass in the research community.

5.3 Intelligent Agents

Many researchers have attempted to implement an intelligent product by agent technology alone. In this context, swarms of primitive reactive agents are not (yet) able to answer the design challenge. It is generally accepted that, today, the power of rational agents is required. In this domain, the BDI model in rational software agents is a nearly universally accepted reference for these intelligent agents.

This model incorporates the four key characteristics of intelligent agents: situated, goal directed, reactive and social. In the context of this paper, the second characteristic renders the agent, in this stricter interpretation, less suitable for integration. While attempting to achieve goals, agents make choices that not necessarily enjoy protection from a corresponding reality.

Conversely, attempts to make an agent so universal that it becomes suitable for integration render it computationally intractable. In manufacturing control applications, Bussmann recommends to keep the agent-based system as simple as possible, reflecting that agents provide unprecedented functionality in dynamic environments but do not achieve the suitability for integration that yields long-lived reusable components (Bussmann, 2003).

The intelligent being is the closest related to the Beliefs in this BDI model. In contrast to the agent approach, the informational state of an application based on intelligent beings is distributed and shared across the application's elements. The community of intelligent beings constitutes the informational state (or the valuable part of it) and shares it with all entities of the application (agents, beings, humans, objects). Moreover, this informational state counteracts update, insertion and deletion anomalies analogous to database normalization, thanks to the delegation mechanisms that are rigorously applied to limit exposure and the fact that information duplication is minimized. Any duplicate information has an expiring date (evaporates) and is considered non-committal cached information by default.

5.4 Electronic Institutions, Environments for MAS

Next to agents and I-Beings, there exists a third kind of construct in the agent research domain. Intelligent agents provide decision-making for individual components. But, societies of agents also have their rules, objectives and norms as a group. The instrument that is put forward to handle this concern in multi-agent systems is the electronic institution (Noriega *et al*., 2002, Malucelli *et al*., 2005, Weyns *et al*., 2007). This means that the proposed structure for the intelligent product may become more complex.

Researchers have proposed to make the environment of a multi-agent system a first-class subsystem in MAS software designs (Weyns *et al*., 2007). Intelligent Beings are crucial components in such an environment. A more general-purpose IT infrastructure also belongs in this environment (e.g. yellow pages services). And, electronic institutions have their contribution. Overall, these remain open issues concerning the overall design of multi-agent systems comprising intelligent products.

6. CONCLUSIONS

This paper puts forward an intelligent product composed of an intelligent agent, responsible for decision-making, and an intelligent being, responsible for reflecting a corresponding reality into cyber space. This intelligent being is in many ways the descendant of navigation maps, sharing and preserving a most interesting property with its ancient predecessor: if an integration conflict involves an intelligent being, the conflict has the same problem with reality. As with maps, the intelligent being is unable to deliver end-user functionality but delivers a valuable contribution nonetheless. Intelligent beings deliver an infrastructure, of a size and a complexity that is out-of-reach of agent-only designs, on which intelligent agents deliver end-user solutions.

ACKNOWLEDGEMENT

This paper presents work funded by the Research Fund of the K.U.Leuven – Concerted Research Action on Autonomic Computing for Distributed Production Systems.

REFERENCES

Bussmann, S (2003). An Agent-Oriented Design Methodology for Production Control. In: PhD Thesis, University of Southampton – UK.

Främling, K., T. Ala-risku, M. Karkkainen (2006). Agent-Based Model for Managing Composite Product Information. In: *Computers in Industry*, 57(1), pp. 72-81, Elsevier.

Holvoet, T., P. Valckenaers (2006). Beliefs, Desires and Intentions Through the Environment. In: *Proc. 5th Intern. Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pp. 1052-1054, Hakodate.

Malucelli, A., H. L. Cardoso and E. C. Oliveira (2005). Enriching a MAS Environment with Institutional Services. In: *Proc. 2nd Intern. Workshop on Environment for Multi Agent Systems at AAMAS* 2005 (Eds. D. Weyns, H. V. D. Parunak, F. Michel), Utrecht.

McFarlane, D., S. Sarma, J. L. Chirn, C. Y. Wong, K. Ashton (2003). Auto ID Systems and Intelligent Manufacturing Control. In: *Engineering Applications of Artificial Intelligence*, 16(4), pp. 365-376, Elsevier.

McFarlane, D (2002). Auto ID Based Control – An Overview. In: *Auto ID Center White Paper.*

Noriega, P., C. Sierra (2002). Electronic Institutions: Future Trends and Challenges. In: *Lecture Notes in Computer Science*, 2446, London – UK, Springer-Verlag.

Parunak H. V. D., S. Brueckner, D. Weyns, T. Holvoet, P. Verstraete, P. Valckenaers (2007). E Pluribus Unum: Polyagent and Delegate MAS Architectures. In: *Proceedings of the 8th International Workshop on Multiagent Based Simulation at AAMAS 2007*, Hawaii.

Schuh, G., S. Gottschalk, T. Höhne (2007). High Resolution Production Management. In: *CIRP Annals – Manufacturing Technology*, 56(1), pp. 439-442.

Simon, H.A. (1990). *The Sciences of the Artificial.* MIT Press, Cambridge Mass.

Valckenaers, P., H. Van Brussel, Hadeli, O. Bochmann, B. Saint Germain, C. Zamfirescu (2003). On the Design of Emergent Systems: An Investigation of Integration and Interoperability Issues. In: *Engineering Applications of Artificial Intelligence*, 16, pp. 377-393, Elsevier.

Valckenaers, P., H. Van Brussel (2005). Holonic Manufacturing Execution Systems. In: *CIRP Annals – Manufacturing Technology*, 54(1), pp. 427-432.

Valckenaers, P., J. Dauter, C. Sierra, J-A Rodriguez-Aquilar (2007). Applications and Environments for Multi-Agent Systems. In: *Journal on Autonomous Agents and Multiagent System*, 14, pp. 61-85, Springer-Verlag.

Weyns, D., A. Omicini, J. Odell (2007). Environment as a First-Class Abstraction in Multiagent Systems. In: *Journal on Autonomous Agents and Multiagent Systems – Special Issue on Environments for Multiagent Systems* (Eds. H. V. D. Parunak and D. Weyns), 14, Springer.

Wong, C. Y., D. McFarlane, A. Zaharudin, V. Agarawal (2002). The Intelligent Product Driven Supply Chain. In: *Proceedings of IEEE Systems, Man and Cybernetics*, Hammamet – Tunisia.