The Knowledge Engineering Review, Vol. 16:4, 375–381.
DOI: 10.1017/S0269888901000200 Printed in the United Kingdom

Holonic and multi-agent systems in industry*

ROBERT W. BRENNAN

Department of Mechanical and Manufacturing Engineering, University of Calgary, Calgary, Alberta, T2N 1N4, Canada; e-mail: brennan@enme.ucalgary.ca

Abstract

The concept of holonic systems has its roots in the desire to understand the structure of natural systems (e.g. living organisms and social organisations) and in particular their ability to behave in a stable yet flexible manner in the face of change. It is not surprising that the lessons learned from these natural systems could help with the design and control of complex man-made systems. However, a key issue is, how can one translate holonic concepts to real industrial environments? For example, one of the key holonic concepts, the holon, can be described as a self-contained autonomous and cooperative entity; when deciding how to implement holons, software agents appear to be the logical choice. In this paper, we summarise the presentations and discussions from a workshop held at the recent *International Conference on Autonomous Agents* that focused on this issue and brought together researchers from both the holonic systems and the multi-agents systems communities.

1 Introduction

Recently there has been growing interest in the holonic approach to the development of complex industrial and business systems. Motivated by the need to enable these man-made systems to adapt to disturbances while maintaining system stability and efficient use of resources, holonic systems were inspired by Arthur Koestler's (1967) early observations of the structure and behaviour of living organisms and social organisations.

Like multi-agent systems (MASs), holonic systems are composed of self-reliant units that are capable of flexible behaviour. More specifically, though, a holon can be thought of as a special type of agent that is characteristically autonomous, cooperative and recursive, that populates a system where there is no high-level distinction between hardware and software.

Although both approaches share many basic concepts, research in each area has been conducted independently for the most part. Holonic systems research has primarily focused on intelligent manufacturing systems and has been organised around the international Holonic Manufacturing Systems (HMS) consortium (HMS, 2001). In contrast, MAS research is much broader in scope, focusing generally on the development of systems in which "data, control, expertise or resources are distributed; agents provide a natural metaphor for delivering system functionality; or a number of legacy systems must be made to interwork" (Wooldridge & Jennings, 1999). As a result, various groups and consortia have emerged to tackle research issues in this area, such as the AgentLink consortium, the International Foundation for MAS, and the Foundation for Intelligent Physical Agents (FIPA).

The *Autonomous Agents 2001* workshop on "Holons: autonomous and cooperative agents for industry" focused on the close relationship between the concepts of multi-agent and holonic systems. For example, the analogous concepts of the HMS "cooperation domain" (i.e. "a logical space in which

* This paper summarises the paper presentations at the *5th International Conference on Autonomous Agents* workshop "Holons: autonomous and cooperative agents for industry", 29 May, 2001.

holons operate, which provides the context within which holons may locate, contact and interact with each other" (HMS, 2000)) and the FIPA "agent domain" (i.e. "a logical grouping of agents/services defined by membership of a directory maintained by the directory facilitator" (FIPA, 1997)) have emerged independently from both the holonic systems and multi-agent systems communities.

The intent of the workshop was to provide an opportunity for researchers to meet and discuss the best holonic approaches as well as identify similarities and differences between holonic and multi-agent systems. In particular, participants reported on the latest advances in holonic systems and discuss the major issues in this area, such as how and why holons differ from agents, and how autonomous, cooperative agents (i.e. holons) can be best applied in industry.

2 Contributions from the speakers

The workshop was well attended by researchers from academic, industrial and government research centres and consisted of a series of paper presentations that are summarised in this section. These papers demonstrated the wide range of research that is currently being conducted on holonic and multi-agent systems and the application of these approaches to industry. In particular, the first presentation, "Holonic enterprise as an information ecosystem" by Mihaela Ulieru and Scott Walker (Ulieru *et al.*, 2001) illustrated this broad scope by discussing research challenges present at various levels of a manufacturing enterprise. They suggested that, to realise the flexibility and responsiveness required of current manufacturing enterprises, one must view the enterprise in terms of three fundamental levels:

- Global inter-enterprise. This is the supply chain level of the enterprise where the concern is on-line order distribution, deployment and tracking.
- Intra-enterprise. This level is concerned with production planning and control – an area that has been the focus of a considerable amount of research in industrial applications of holonic and multi-agent systems (e.g. McFarlane & Bussmann (2000)).
- Machine (physical agent). The concern at this level is with distributed control of real-time systems and the application of holonic and multi-agent system concepts in this domain.

In their presentation, Ulieru and Walker described a mapping between the concepts of holons and agents that places the mediator agent (Maturana & Norrie, 1996) in a central role, i.e. it acts as an interface between software agents within the holon, and other holons in the manufacturing enterprise. They also suggested that the holonic enterprise concept could be thought of as a merging of the HMS and FIPA models, i.e. HMS traditionally focused on the bottom two levels ("intra-enterprise" and "machine") while FIPA traditionally focused on the top two levels ("inter-enterprise" and "intra-enterprise"). Finally, they noted that the key to integrating these concepts across the entire enterprise is the identification of patterns of holonic collaboration at each of the fundamental levels.

The presentation of this very general model provided a good foundation for the presentations that followed on holonic and multi-agent system design and industrial applications of these systems. The two sub-sections that follow briefly summarise the presentations in each of these areas.

2.1 Holonic and multi-agent system design

2.1.1 Paul Valckenaers

In his presentation "Stigmery in holonic systems", Paul Valckenaers described a unique approach to load forecasting in production systems that is based on indirect coordination techniques (Valckenaers *et al.*, 2001). Rather than designing agents to explicitly interact and negotiate as in the Contract Net, Valckenaers's approach parallels that used by ant societies, i.e. agents leave information at physical locations in the manufacturing system that is used, like pheromone in an ant colony, to influence other agents' behaviour. The benefits that natural systems gain from this indirect coordination approach were noted to cross over very well to the manufacturing domain, i.e.

- global information can be made available locally,
- the environment is part of the solution and

- global/remote information can be made available without exposing individual agents to the overall system.

The presentation led to a discussion on the implementation of this approach in real manufacturing environments in which Valckenaers described a four-step process that can be used to allow the manufacturing control system to make short-term load forecasts available throughout the system:

- "agentify" the environment by means of PROSA (Van Brussel *et al.*, 1998) resource agents,
- design PROSA order agents with the capability to propagate their intentions throughout the factory (i.e. the order agents create mobile agents to accomplish this task),
- provide resource agents with the capability to back-propagate their load forecasts (also by creating mobile agents) and
- allow order agents to use these forecasts to make their routing decisions.

2.1.2 Martyn Fletcher

Martyn Fletcher next discussed general holonic systems modelling concepts and, in particular, the extension of these principles to the realm of real-time control in "Principles underpinning a unified architecture for holonic manufacturing systems" (Fletcher & Deen, 2001). His presentation started with an overview of holon and holonic system architecture and the role of the cooperation domain. In particular, it was noted that an HMS could be modelled as a "holarchy", i.e. a heterarchy of cooperation domains that are used to solve a set of manufacturing tasks.

To implement this approach, either a functional or object-oriented perspective could be taken, depending upon the application domain. When the focus is distributed real-time control, he suggested that the most appropriate engineering model is the emerging IEC 61499 industrial process control model, which is firmly rooted in the functional approach (i.e. the key stereotype in this model is the function block). To illustrate this point and demonstrate the feasibility of holonic systems, an engine assembly test-bed developed during phase I of the HMS project was used as an example.

Using this test-bed to set the stage for how real-world holonic control systems could be developed with IEC 61499, Fletcher next walked the workshop participants through a detailed function-block-based model to handle information and support cooperation in an HMS. This led to a discussion of the open issues that must be addressed in order to construct a complete holonic model:

- What other holons are needed to complete the model?
- How can a cohort holon be dynamically created and how should these dynamic holons be managed?
- How does a common data/knowledge base join a cooperation domain and interact with these holons?

2.1.3 Michal Pechoucek

In the final presentation on holonic and multi-agent systems design, Michal Pechoucek described how social knowledge can be maintained in a multi-agent system and proposed a knowledge-based model of agents' mutual awareness (Marik *et al.*, 2001). This presentation generated a considerable amount of discussion on the issue of "local ownership" of global information in a multi-agent system, such as:

- What should be the extent of global knowledge that is owned by individual agents?
- How much should individual agents know about global rules and knowledge ontologies and about peer agents?
- How should locally owned knowledge be structured?

To address these concerns, Pechoucek offered a tri-based acquaintance model that is designed to limit the potentially explosive communication space in multi-agent (and holonic) systems. The viability of this approach was then demonstrated by comparing the required communication traffic for various sizes of agent communities using (i) broadcasting agents that use the classical contract net protocol and (ii) tri-based agents. The substantial savings that result from tri-based reasoning led Pechoucek to suggest that this approach would have considerable potential in application areas such as production planning and resource allocation, supply chain management, planning coalition operations and intrusion detection.

2.2 *Industrial applications of holonic and multi-agent systems*

The final three presentations of the workshop concentrated on the application of holonic and agent concepts to specific industrial problems and, in particular, to the intra-enterprise level (Ulieru *et al.*, 2001) of the manufacturing enterprise.

2.2.1 *Hidehiko Wada*

First, Hidehiko Wada presented a prototype holonic system that uses autonomous and mobile agents to determine resource allocation in a manufacturing test-bed (Wada & Okada, 2001). This research was motivated by the need for greater flexibility in production planning and scheduling systems than what is offered by extant hierarchical systems (i.e. these systems are not responsive to change and are difficult to modify).

To meet this need, Wada proposed a flat architecture manufacturing system that uses three agent types that are each similar to the holons proposed in the PROSA architecture (Van Brussel *et al.*, 1998), i.e. product agents (PROSA product holons), machine agents (PROSA resource holons), and management agents (PROSA staff holons). This agent-based system was then implemented on the Shape Deposition Manufacturing Laboratory at Carnegie Mellon University (SDM, 2001) using mobile agents to represent the product agents. This allowed product agents to transfer across devices (e.g. robots, computer numerically controlled machining centers) in order to accompany the movement of the real products through the manufacturing system.

In his presentation, Wada noted that the main benefit of this approach was the flexibility that resulted from agent autonomy. For example, the system is very robust to changes such as the addition and removal of resources (e.g. as a result of a machine failure) since product agents can adapt quickly to changes in system configuration. Further, by simply instantiating product agents with the new process plans, new products can be easily introduced to the system.

2.2.2 *Weiming Shen*

Next, Weiming Shen proposed a combined genetic algorithm-based search and agent-based negotiation approach for scheduling manufacturing systems (Shen, 2001) which he noted is intended to overcome the difficulties with applying traditional scheduling methods in real-world situations. For example, traditional scheduling methods are typically based on simplified theoretical models and are often implemented in a centralised fashion.

The key feature of Shen's approach is its enhancement of an optimal or near-optimal schedule generation approach (i.e. genetic algorithms) by combining it with agent-based scheduling. As a result, the system's flexibility and responsiveness is increased, allowing both advanced scheduling (i.e. generating a schedule to allocate manufacturing resources among parallel and sequential activities) and dynamic scheduling (i.e. on-line modification of the schedule in response to unanticipated events) to be managed in an integrated manner.

The prototype system that Shen uses to demonstrate this approach is built on his earlier mediator agent-based manufacturing scheduling system (Shen & Norrie, 2001). In this architecture, cooperation is realised through a combination of mediation mechanisms (i.e. by various levels of mediator agents) and the Contract Net protocol.

Discussion centered on the issue of implementation and, in particular, on the performance of the system in a stochastic manufacturing environment. For example, it was noted that the system would be required respond quickly to changes such as machine failures and rush orders that would invalidate the current schedule. Shen pointed out that, although genetic algorithm-based global search does require more time than reactive approaches (e.g. the Contract Net), experiments with a typical manufacturing system showed that schedules could be generated in a reasonable time (in the order of minutes). Further, the combined approach clearly outperformed pure bidding-based approaches in terms of cost. The next step in this research however, requires more experiments to be conducted with this system to evaluate real-world factors that may not have been captured in the simulation, such as communication latency.

2.2.3 Paul Rogers

In his presentation, Paul Rogers described ongoing work on the application of an agent-based dynamic scheduling system for the electronics manufacturing industry (Rogers & Kan, 2001). In particular, he noted that although there has been a considerable amount of work on the development of holonic and agent-based production planning and control systems (e.g. McFarlane & Bussmann (2000)), few of these systems have been proven in real industrial environments, and little guidance has been provided on how this next step can be made.

As a result, this presentation focused on addressing this issue through the development of an agent-based dynamic scheduling system for a specific industrial environment, i.e. a surface mount technology (SMT) production facility. Rogers noted that the decision to choose an SMT facility for their work was based on two main advantages that this domain provides:

- it is inherently volatile, allowing the potential benefits of a distributed agent-based approach to be shown; and
- the system complexity is manageable enough to allow a prototype system to be developed and tested.

To illustrate these points, Rogers described the details of the problem domain and, in particular, the performance goals and planning and control decisions that are specific to SMT facilities. Similar to the systems investigated by others in this workshop, it was clear that in addition to meeting overall performance goals such as on-time deliveries, flexibility and responsiveness are the key desired system characteristics. To achieve these goals, Rogers noted that decision-making should be grouped into three distinct levels for this particular application: area, line and machine level. Further, the agent-based system must also be capable of dealing with common SMT facility disruptions such as product mix variability and component supply problems.

Rogers next described a proposed agent-based architecture that is based on the three decision levels noted above (i.e. the architecture consists of area, line and machine agents) and the FIPA open source (FIPA-OS) platform (FIPA, 1997). He noted, however, that the intention is not to define generic distributed scheduling architectures as others have done (e.g. Fletcher & Deen (2001); Van Brussel *et al.* (1998)); instead, it is to develop a platform to test alternative distributed scheduling system designs in a common test domain.

The workshop discussion focused on the issue of agent autonomy in distributed scheduling systems, i.e. what is an appropriate level of autonomy for agents such as the line agent? Rogers noted that, although line agents possess significant autonomy, this autonomy is not unconstrained (e.g. area agents may force line agents to behave in a particular way to achieve the desired system performance). He added, however, that this may lead some to say that line agent behaviour is "non-holonic" since the system no longer exhibits cooperative autonomy.

3. Summary

The workshop brought together a range of issues on the application of multi-agent and holonic systems to industry and, in particular, focused on the work that is being conducted on the design of these systems (e.g. coordination techniques, system architecture and the management of social knowledge) and the application of prototype systems to specific industrial problems. From the discussions that followed it became clear that a number of areas must be addressed in future research if the ideal of a truly holonic system is to be realised.

For example, although there was a considerable amount of discussion on the concepts of holons and agents, very little work has been done on any formal comparison of these concepts. As noted above, Walker described a conceptual mapping of these concepts that places the mediator agent in a central role but did not provide any theoretical background behind this mapping. This concept, however, appears to be a good starting point for this type of analysis.

A second important area that was raised by Rogers is the need for relevant industrial test cases to prove the concept of holonic and multi-agent systems in the real world. This area of research is not only important for the demonstration of the potential of these approaches, but is also necessary if researchers are to compare the performance of proposed systems in common environments. At first glance, real physical systems offer the best opportunity to test these concepts; unfortunately, they also present a number of practical difficulties to "open" experimentation. For example, experimentation with real industrial systems may be hindered by the proprietary nature of the system design as well as the potential operational risk associated with experimentation (e.g. the expense associated with frequent down times, and danger to people and equipment). As a result, it appears that the best alternative is to define realistic benchmark systems that are based on real industrial applications, but are implemented using a simulation tool (e.g. discrete-event simulation). Although not reported in this workshop, there has been some interesting work in the area of common simulation benchmarks by Cavalieri *et al.* (1999) that is already gaining considerable support from the holonic and multi-agent systems research community.

Finally, it became clear that the majority of the work on the application of holonic and multi-agent system concepts has focused on Ulieru's intra-enterprise level yet not on the inter-enterprise level and the physical machine level. For example, at the physical machine level the basic requirements of application management (i.e. application configuration and reconfiguration) and safety management (i.e. fault detection and recovery) must be met in a distributed, hard/soft real-time environment. Although the general design concepts that were discussed at this workshop apply at this level, these concepts will have to be revisited to ensure that time-critical performance is achieved. As noted above, Fletcher addressed this point in his presentation, though he also made it clear that this work is still very much in its early stages.

The success of this event is a reflection of the growing interest in the application of holonic and multi-agent systems to industry. This has been demonstrated in a number of recent sessions and workshops such as the highly successful International Workshop on Intelligent Manufacturing Systems organised by Van Brussel and Valckenaers (1999) and the First International Workshop on Industrial Applications of Holonic and Multi-Agent Systems (HoloMAS) organised by Marik and Pechoucek (2000). Further, upcoming workshops in this area are a second HoloMAS in September 2001 (see `http://cyber.felk.cvut.cz/HoloMAS/2001/`) and the Mini Symposium on Intelligent, Distributed Manufacturing Management and Control in October 2001 organised by Duncan McFarlane of Cambridge University and William Gruver of Simon Fraser University (see `http://crater.sie.arizona.edu/idmmc.html`).

References

Cavalieri, S, Bongaerts, L, Macchi, M, Taisch, M, Wyns, J, 1999, "A benchmark framework for manufacturing control" *Proceedings of the 2nd International Workshop on Intelligent Manufacturing Systems*, 225–236.

FIPA,1997, *Part 1 Version 1.0: Agent Management Specification*, Foundation for Intelligent Physical Agents.

Fletcher, M and Deen, SM, 2001, "Principles underpinning a unified architecture for holonic manufacturing systems" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 23–30.

HMS, 2001, "Holonic manufacturing systems overview" `http://hms.ifw.uni-hannover.de/public/overview.html`.

Koestler, A, 1967, *The Ghost in the Machine* Arkana.

McFarlane, D and Bussmann S, 2000, "Developments in holonic production planning and control" *Production Planning and Control*, **11**(6) 522–536.

Marik, V, Pechoucek, M and Stepanokova, O, 2001, "Organization of social knowledge in multi-agent systems" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 39–45.

Marik, V and Pechoucek, M (eds), 2000, *Proceedings of the 11thInternational Workshop on Database and Expert Systems Applications: First International Workshop on Applications of Holonic and Multi-Agent Systems* (HoloMAS) IEEE Computer Society.

Maturana, F and Norrie, DH, 1996, "Multi-agent mediator architecture for distributed manufacturing" *Journal of Intelligent Manufacturing* **7** 257–270.

Rogers, P and Kan, S, 2001, "Agent-based dynamic scheduling of surface mount technology production facilities" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 61–67.

SDM, 2001, Shape Deposition Manufacturing Laboratory Home Page, `http://www.cs.cmu.edu/ sdm`.

Shen, W, 2001, "Integration of agent-based approaches with genetic algorithms for manufacturing scheduling" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 55–60.

Shen, W and Norrie, DH, 2001, "Dynamic manufacturing scheduling using both functional and resource related agents" *Integrated Computer-Aided Engineering* **8**(1) 17–30.

Ulieru, M, Walker, SS and Brennan, RW, 2001, "The holonic enterprise as a collaborative information ecosystem" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 1–14.

Valckenaers, P, Van Brussel, H, Kollingbaum, M and Bochmann, O, 2001, "Stigmergy in holonic systems" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 15–21.

Van Brussel, H and Valckenaers, P (eds), 1999, *Proceedings of the 2nd International Workshop on Intelligent Manufacturing Systems* KU Leuven.

Van Brussel, H, Wyns, J, Valckenaers, P, Bongaerts, L, Peeters, P, 1998, "Reference architecture for holonic manufacturing systems: PROSA" *Computers in Industry* **37** 255–274.

Wada, H and Okada, S, 2001, "Towards holonic manufacturing sytems: an autonomous agent approach" *5th International Conference on Autonomous Agents Workshop on Holons: Autonomous and Cooperative Agents for Industry* 47–53.

Wooldridge, M and Jennings, N, 1999, "Software engineering with agents: pitfalls and pratfalls" *IEEE Internet Computing*, **3**(3) 20–27.