Cooperation Models Between Humans and Artificial Self-Organizing Systems: motivations, issues and perspectives

Gabriel Zambrano Rey[1,2,3]
[1]Department of Industrial Engineering
Pontificia Universidad Javeriana
Bogotá, Colombia
gzambrano@javeriana.edu.co

Marco Carvalho
Human-Centered Design Institute
Deparment of Computer Sciences
Florida Institute of Technology
Melbourne, FL 32901, U.S.A.
mcarvalho@fit.edu

Damien Trentesaux[2,3]
[2]Univ. Lille Nord de France
F-59000, Lille
[3]UVHC, TEMPO Lab, PSI-Team
F-59313, Valenciennes, France
damien.trentesaux@univ-valenciennes.fr

***Abstract*— In this paper we introduce and discuss some of the concepts, motivations and requirements for cooperative models between humans and artificial self-organizing systems. After a brief review of alternative cooperation methods, an implicit cooperation-based approach through parameterization and incentives is proposed for human-agent collaboration. Our case study is based on self-organizing, flexible manufacturing control systems.**

Keywords—Flexible Manufacturing Control, Human-Agent Interactions, Artificial Self-Organizing Systems, Human-Automation Teamwork

I. INTRODUCTION

Artificial self-organizing (ASO) systems are often inspired by biological (ants, flocks, swarms and immune systems) or social (markets) systems. The global patterns and behaviors of such systems normally emerge from local interactions between their constituent artificial autonomous entities (AEs). Artificial entities are endowed with specific rules and atomic behaviors that are applied using only local information, or information gathered, directly or indirectly, from local entities.

One of the main reasons for the growing interest in self-organizing systems is their inherent distributed nature, which allows the natural decomposition of complex tasks into smaller entities that work together to achieve a desired goal or find a global solution. For some tasks, this is generally attained when the entire system reaches an equilibrium or steady state [1].

Self-organizing systems are also commonly recognized by their robustness and resilience and, in many cases also by their level of dependability. For the purpose of this work, we define an ASO system as robust if the system is able to continue delivering its service within acceptable levels after encountering unexpected failures, external attacks, or unanticipated operational conditions. If the quality of service is temporarily violated due to the magnitude of the perturbation, the system is not robust, but it is considered resilient, as long as it is able to recover and resume operations within acceptable levels of service. The system is said to be dependable if most of the time it provides an acceptable service.

It is important to state that here the notion of robustness is related to the short-term response of the system while dependability deals with the system behavior in the long term. Figure 1 shows an example taken from a manufacturing system. The system is robust to certain internal perturbations (zone A), resilient but not robust in zone B, and dependable but not robust in zone C, for other perturbations. However, after a critical perturbation illustrated in zone D the system cannot provide any of the previous characteristics.

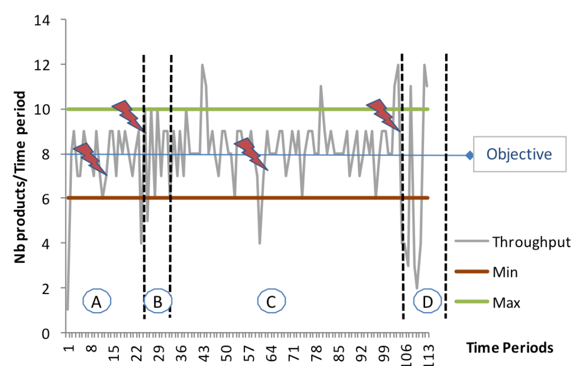

Fig. 1. Level of Service (throughput) of manufacturing system's throughput

Typical ASO system frameworks are based on multi-agent, holonic and bio-inspired paradigms, with applications to complex problems in manufacturing [2][3], supply chain management [4], traffic and transportation [5], security [6], health care [7] and space exploration [8], among others. In most of these applications, and in particular in the context of industrial manufacturing processes, the constituent AEs, e.g. agents or holons, represent physical elements or logical activities designed to act autonomously, with no direct human intervention. Human interventions, actions and controls are normally conceived as external perturbations to the state of the system, which then responds by seeking a new equilibrium. In this context, ASO systems are normally designed to respond to humans in the same way that they respond to changes in the

978-1-4799-0503-4/13/$31.00 ©2013 IEEE

environment, without necessarily considering models that take human-agent cooperation[1], or teamwork into account.

In this paper, we propose that human-agent cooperation and teamwork must be an additional and important modeling component within ASO systems. We argue that such considerations will help maintain the human in the loop while still allowing the self-organizing properties of the system to be maintained, improving its overall robustness, dependability and resilience.

One of the main reasons for seeking human-ASO system cooperation is the human's flexibility and ability to contextualize a specific problem or abnormal condition. ASO systems can be designed to be robust, dependable and resilient for a set of perturbations. However, these techniques are highly sensitive and fragile to other unknown and unexpected errors, faults or threats. In these cases, human cooperation can be an important contribution, enhancing the system's ability to recover by influencing AE actions in many ways, such as explicit commands, rule parameter tuning, implicit environmental influence, reinforced learning, etc.

However, human-agent cooperation has its challenges and ASO systems must also be designed to accommodate some of these problems that tend to emerge from the intrinsic different nature of artificial entities and humans. Response times, roles, background knowledge, context understanding, information management, and several other cognitive capabilities are some of the aspects to be considered when modeling an ASO system taking into account human cooperation and integration.

In this paper, we focus on human-ASO system cooperation. In Section II we start by presenting a conceptual human-ASO system cooperation model. Then, in Section III we state some motivations for working on such cooperation and we also present some of the issues that have emerged from human-ASO system cooperation. Section IV reviews some of the literature on the subject and then, in Section V we present a case study as an example of the model introduced in Section II. Some research prospects and concluding remarks can be found in Section VI.

II. Conceptual Human-ASO System Cooperation Model

Generally speaking, ASO systems are composed of an environment (e.g. physical world, networks), artificial entities (AEs) e.g. software agents, holons with information and physical units, self-organizing mechanisms that define the rules that entities use to act on the environment, and artifacts (e.g. information in the form of messages or pheromones, machines, routes, fleet of robots) which are the passive elements supported by the environment and managed by the agents [1]. An AE could also represent a sub-system composed of AEs, with each AE having its own local goals.

In the context of manufacturing control systems, we propose that human-ASO system cooperation can be achieved in different ways in order to offset mutual deficiencies and exploit mutual advantages. Since there are many events that are really hard to anticipate, we need the flexibility of the ASO and the human entities (HEs) to drive the system towards the objectives. In the model presented in Fig. 2, humans in-the-loop are concerned with setting up the system objectives. Based on their analysis of the external conditions and performance indicators, they can cooperate with the system implicitly and/or explicitly. This model applies the open-control framework [9] for human-ASO system cooperation. The difference is that instead of having "implicit control", we use the term "implicit cooperation" in order to make clear that, most of the time, HEs and AEs should pursue mutually acceptable actions. In addition, we would also like to state that AEs autonomy should be encouraged as much as possible in order to guarantee the adaptability, reactivity and resilience of ASO systems.

Explicit control is executed when a specific behavior is required of the system, and therefore AEs need to follow specific directions. This kind of control can be achieved by

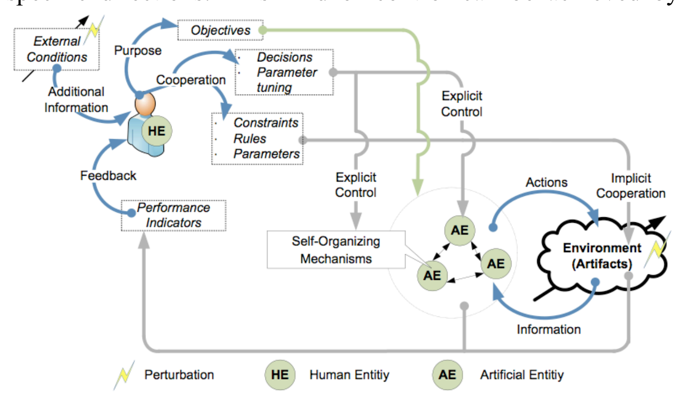

Fig. 2. A Human-ASO System Model

imposing decisions on AEs, imposing behavioral rules/policies and/or fine-tuning the parameters of the self-organizing mechanisms.

Hence, AE autonomy is compromised and their behavior is highly constrained by the HE controller, who in turn becomes highly responsible for system resilience. This could be the case in which, after a critical perturbation, the system is not capable of guiding itself back to a stable state, and it could create deadlocks or reach situations that are impractical or undesirable and cause severe damage. Implicit cooperation, on the other hand, explores an indirect type of intervention, avoiding master-slave relationships and endeavoring to constrain AE autonomy as little as possible. The purpose is to obtain bounded self-adaptation and self-evolution behavior that is capable of absorbing unexpected events and successfully adjusting the system structure to attain system objectives. Implicit cooperation can then be possible by changing/adding information, constraints or parameters used by artifacts residing in the environment. HEs do not interact directly with AEs; they do so indirectly by setting behavior bounds in the environment.

Consequently, resilience under this kind of cooperation is achieved cooperatively between HEs and AEs. An example of this kind of cooperation is pheromone adjustment in *stigmergic* approaches. Since pheromones guide the solution search process executed by ants, ants can be guided towards a

[1] In this paper, the word cooperation should be understood as any human action, decision, intervention or influence on the system.

particular solution without explicitly telling them what to do or how to do it. Implicit cooperation just changes the ants' perception of the current environmental conditions.

Explicit control can also be a very particular case of implicit cooperation, when the bounds set up by HEs are so tight that they do not allow AEs any autonomy. But the purpose of implicit cooperation is twofold: first to guide the AEs towards a specific behavior, and second to remain robust in the face of possible perturbations, which can be dealt with by AEs if they have enough autonomy to do so.

Three critical issues determine the usage of implicit cooperation or explicit control. First, the current state of the system is a triggering factor, especially states that might lead to a system deadlock or a potential risk. In these cases, explicit intervention might be necessary and will make the system respond more rapidly.

The second issue concerns the required response time. While explicit control most likely generates faster, clear responses from AEs, implicit cooperation requires more time for the system to adapt and evolve towards the desired state. However, explicit control may trigger oscillations and other emergent behaviors that could take the system to instability.

The last issue concerns the appropriate information required by HEs in order to cooperate with the system. In some cases, it could be necessary to count on additional tools, such as simulation-optimization tools, in order to define the course of action and the best kind of cooperation. This would depend on the kind of situation and its costs in terms of performance indicators. In the next section, we present other motivations and issues that encourage further research on this topic, and that could help complement the model presented previously.

III. Motivations and Issues of Human-ASO System Cooperation

As mentioned, human-ASO system cooperation is becoming a key research issue, not only because ASO systems are currently being used to model high-profile applications such as space exploration missions and robot fleet control, with applications in rescue and military missions, but because they can be found in many other everyday applications such as networking and pervasive computing [8]. For the resulting organization, it is imperative to explore the structural and functional features that are required to support such cooperation [10]. In fact, for some authors the first step towards modeling such cooperation is the recognition of HEs and AEs as not being interchangeable but complementary and cooperative entities. In this section, we first explore some motivations to continue working on this topic and then we present some important issues to be considered.

A. Motivations

In the model presented in the previous section, the ASO system is only configured by artificial entities that represent physical objects or information. In order to complete such a model, one important aspect is the integration of human entities as components of the ASO and not as external individuals. Such integration requires the representation of the characteristics, parameters, competencies and even preferences of the HEs in conjunction with those of the AEs (except for preferences which are more feeling-based parameters). Therefore, the system can be set up to look for multiple objectives, maximize HE profits and satisfaction as well as economic profits, security levels, or other performance criteria. For example, this is particularly clear in personnel scheduling where humans are organizational resources, and aside from the organization's economic objectives, human operators also seek their own particular benefits and profits [11].

Another way in which humans can close the loop efficiently is by enhancing situation awareness through a cooperative analysis of complex situations with AEs. AEs can be very useful in gathering and pre-processing raw information, reducing the mental workload and maintaining HEs' activities oriented mainly towards fault anticipation, detection and recovery. This is particularly true in cases where process conditions are individually within normal ranges, but together can lead to an abnormal situation (or vice-versa). ASO systems have a hard time assessing these situations because of their decentralized nature. For example, in a manufacturing facility, a human operator knows the state of all the machines, and by running certain analyses it is possible to predict certain maintenance activities. If task allocation is achieved by a self-organizing system (e.g. multi-agent systems with ant colony optimization [12]), the production plan is constructed on the basis of pheromone levels. A sudden maintenance activity on a machine that has jobs in its queue (i.e. contracts) will require a new production plan to be calculated because the pheromone trail that passes through the machine is broken. Searching for new optimized pheromone paths requires time and will definitely impact the overall performance. Instead, a smoother adaptation is possible if the operator's knowledge is used to influence path planning to gradually isolate the machine. Product agents continue to make their own decisions to seek their local goals, but the overall performance will not be affected as much as if the maintenance event was the one that triggered the recalculation of the path.

Another motivation for working on human-ASO system cooperation is the fact that AEs can implement learning procedures. AEs not only learn from their own actions but also from HEs' interventions. The positive or negative reinforcement of AEs' decisions favors their learning and adaptation, and might reduce the amount of future explicit interventions and the workload for humans.

B. Issues

Humans "in the loop" are an interesting challenge given the complexity of ASO systems, the different nature of both entities and their cognitive abilities. Aside from the other issues that we have described in the previous sections, one of the most critical issues of having humans in the loop is their capacity to abstract, understand and anticipate ASO system behavior. This is mainly due to the number of possible interactions, which grows exponentially with the number of components. Furthermore, the continuous adaptation and evolution due to unpredictable changes also triggers structural changes that in turn require human adaptation to the new organization [13].

In certain environments, actions occur at a fast pace and decisions must be made within tightly constrained time frames, under uncertain conditions and with partial information. These conditions can create lags between HEs and AEs, which can create challenges for cooperation. In addition to these lags, other reasons that could lead to problems in human-agent cooperation include a poor combination of individual efforts and skills, the breakdown of internal processes, i.e. communication, negotiation, information storage, the improper use of available information (which may be outdated, in the wrong format, in unavailable), and problems with the allocation of authority. These problems create tend to create confusion and can lead to severe consequences, including loss of life, missing critical action time and economic inefficiencies [13]. Lags and failures on human-ASO interactions are also likely to affect the level of confidence and trust that humans can grant ASO systems. It is often due to a reduced level of confidence, that most of the cooperation models tend to place HEs as super system controllers (explicit controllers), limiting the self-adaptation and self-organization nature of ASO systems, and often increasing the mental workload of humans [14].

One well-known problem with these kinds of approaches is the allocation of autonomy. Some recent studies have focused their attention on adjusting autonomy levels. However, further research is needed on the appropriate requirements for adjustable autonomy, and the context-based delegation of authority. This is still an active area of research, with little agreement on the principles, algorithms and mechanisms that can be generally applied for autonomy shifting [15]. In addition, it is also important to determine and regulate the amount and nature of information shared between entities (ASO components or humans), which is not an easy task within highly decentralized systems.

In the next section, we briefly introduce some of the related studies that have addressed similar issues, before introducing our proposal for human-agent teamwork on manufacturing control systems.

IV. Short Literature Review on Cooperation models

There are very few studies dealing with human-ASO system cooperation and those that exist focus mainly on the allocation of authority. Even though rigid hierarchies are not desirable in ASO systems and any type of centralized structure should be avoided, for some researchers the human role must evidently be the highest authority [16]. For instance, in Urlings' [17] teaming principles for human-centered automation (that also apply to human-ASO system design), HEs are responsible for the outcomes in HE-AE teams and must be in command and actively involved in such teams (i.e. adequately informed and capable of controlling the system). HE-AE cooperation is achieved by allowing AEs to monitor the performance of the HEs, and by making each team member (HEs and AEs) inform others of their intentions. If these principles are only translated into explicit control, then the main features of ASO systems are compromised.

For other researchers, rather than positioning humans in a supervisory role, other interaction mechanisms such as mediation or cooperation through mutual understanding are also possible [18]. With more of a mediation role, the collaborative management agent (CMA) in [19] is proposed to support human-agent, human-human, and agent-agent interaction and collaboration. The role of the CMA is to assist other agents and support their negotiation processes, also deciding when human assistance is required. Similarly, in [15] the human-agent roles are reasoned by the system, and task allocation is shaped dynamically on the basis of task criticality, task frequency, participants' cognitive abilities and the resulting cognitive workload. The purpose is to assign tasks to those entities, humans or AEs, that are the best suited and that count on the appropriate resources (i.e. information, artifacts) to accomplish the task. For tasks where these mechanisms are appropriate, the reasoning system determines who has the control and the level of autonomy for that task. Although these approaches try to balance authority, they all explore switching mechanisms that determine when explicit control by HEs is required without introducing other types of cooperation.

With a more cooperative approach, and searching for joint sense-making, in [20] a coactive emergence approach for human-agent teamwork is proposed in which task execution activities are performed by analysts and software agents in tandem. In coactive emergence, humans are analysts that evaluate agents' findings and create policies to direct agent activities. The term coactive is used to emphasize the joint, simultaneous and interdependent nature of the collaboration among analysts and agents. This approach can also be classified as an explicit control approach, but this time it works through self-organizing mechanisms. Instead of telling AEs what to do, they are told how to do it, giving them a certain level of autonomy to adapt and evolve towards a specific objective.

In several approaches, a supervisory-type intervention is intended to be accomplished by AEs with supervisory or coordinative roles. Instead of HEs, there are specialized AEs that count on tools such as iterative simulation, neural networks (NN), data mining and heuristic optimization, and which are in charge of changing the subordinate AEs' policies or developing the policies by fine-tuning their parameters, e.g. weighted combinations of decisional rules [21][22][23]. These could become additional tools that HEs could use in order to be able to cooperate effectively with the system.

As far as the authors' know, no HE implicit cooperation applications have been reported. As in the previous cases, it is the task of supervisor AEs to influence subordinate AEs in the system. Implicit cooperation can be achieved, for instance, in bio-inspired approaches in which information is spread throughout the environment in the form of pheromones in ant colonies or waggle dances in bee societies. Reference [24] reports a survey of bio-inspired applications for complex systems in which implicit cooperation can be implemented.

V. A Manufacturing Control Case Study

A. *System description*

In this paper we consider a potential fields-based ASO system for dynamic task allocation and dynamic pallet routing within a flexible manufacturing system, see Fig. 3 [25]. In this system, machines (artifacts) emit attractive fields for each manufacturing service they provide, and products (AEs) dynamically choose the machines according to their manufacturing sequences.

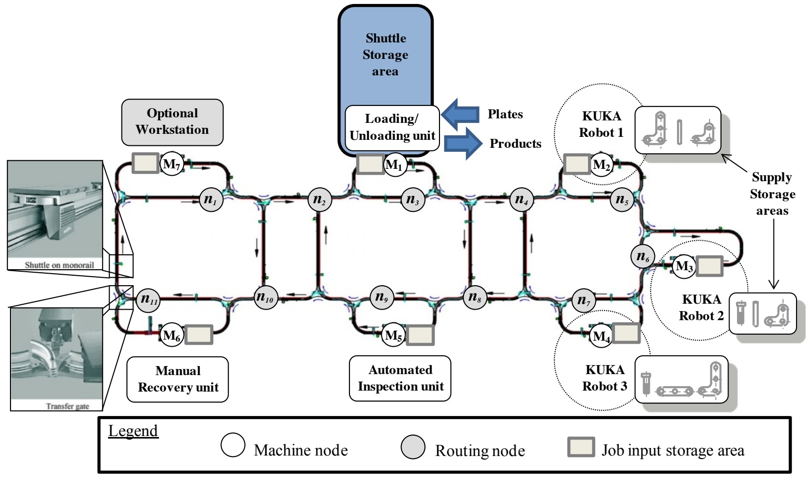

Fig. 3. The Flexible Manufacturing System

Attractive field levels depend on machine availability and are spread throughout the conveyor system (the environment that provides the transportation services between machines and the machine position). These fields are affected by the transfer time and the state of the transfer section (between two nodes ***n*** in Fig. 3). This means that the longer the trajectory, the poorer the quality of its component transfer sections (e.g. traffic jams, breakdowns, deadlocks that increase transfer time), the higher the attenuation of the field throughout that trajectory. The system displays the following self-organizing behaviors:

- HEs can implicitly cooperate with AEs by manipulating the machines' fields. This would progressively decrease the workload of the machine, instead of abruptly when the machine breaks down and products are already allocated to the machine's queue (i.e. job input storage areas in Fig. 3).
- AEs have neither predefined machine allocations nor do they follow specific trajectories. AEs are completely autonomous.
- Manufacturing tasks are allocated stepwise, one at a time and according to the best attractive field sensed by the AE (which also depends on its position in the system).
- Decisions are constantly assessed (at nodes ***n*** in Fig. 3), meaning that AEs can change their choices dynamically according to the state of the machines and trajectories, (e.g. machine breakdowns or maintenance will reduce fields to zero so AEs need to look for another machine). This is an example of self-adaptation.
- AEs do not need to know other products' plans and do not communicate directly with other AEs, but AEs update the state of the transfer sections between nodes on the conveyor according to their transfer experience; this changes the way fields are affected.

B. Human cooperation through explict control and implicit cooperation

Although this system displays interesting features such as adaptability, scalability, high decentralization and reactivity, it can be enhanced by human cooperation in the following ways:

- Explicit control can be used when deadlocks appear. HEs can dictate the machine and route an AE needs to follow in order to get the system out of the deadlock (i.e. in Fig. 3, blockage inside an internal loop such as the one formed by n_4-n_5-n_6-n_7).
- Explicit control can also be used by changing the way AEs assess the machines, for instance by involving other criteria (e.g. repulsive fields of other AEs).
- HEs can implicitly cooperate with AEs by manipulating conveyor section performance; pushing AEs away from a region of the FMS, for example, to balance workloads (e.g. in Fig. 3 the performance of segments starting at n_{10} can isolate M_6 and M_7).
- HEs can implicitly cooperate with AEs by adjusting the parameters that machines use to calculate their attractiveness levels, to integrate other criteria, e.g. energy consumption, QoS, processing speed.

To be able to cooperate with AEs, HEs need to count on updated information and a clear knowledge of the way AEs behave; otherwise this cooperation could make the system nervous and unstable. Further research will study each one of these possible cooperation mechanisms and how they enhance system behavior and affect performance indicators.

VI. Research Prospects and Conclusions

One of the most important conclusions that we can draw from this preliminary study is that the various ways in which humans can cooperate and intervene should be taken into account when designing ASO systems. If they are not taken into consideration in the design process, then humans are either relegated or they end up taking control of the system as super controllers, which actually goes completely against the principle of ASO systems and increases the workload for humans. We presented a model in which two cooperation mechanisms are proposed: explicit control in which the human drives the system by telling artificial entities what to do or how to do it; and implicit cooperation in which humans influence artificial entities indirectly through the environment and/or the artifacts. However, more analysis is required to obtain successful cooperation models, especially regarding information acquisition and analysis in decentralized systems (e.g. using data mining). It is important to note that, in the proposed model, human entities can collaborate with artificial entities at the same level, and an entity in a supervisory role could in fact be modeled as a compound entity reflecting the negotiation between a supervisory artificial entity and a human entity in order to cope with more complex systems.

Another dimension of human-ASO cooperation that should lead to future studies is the capacity of ASO systems to absorb and handle human errors. Human errors could be necessary for the ASO learning process (adaptability), but they should be contained to avoid serious consequences. This is particularly important when full authority (meaning explicit control only) is allocated to HEs only, because the ASO system does not have enough autonomy to absorb the consequences of bad decisions and recover itself. That is why autonomy should be balanced for the system to be resilient to human errors and implicit cooperation should be further encouraged. As a matter of fact, this is a critical issue that compromises the system's resilience.

Finally, another prospect concerns the learning processes in order to avoid constantly falling into explicit interventions. The challenge continues to be with the interpretative mechanisms because the goals and human expectations must be translated and understood by AEs, and AEs must be capable of assessing their performance to determine if the resulting executions meet the goals.

Acknowledgments

Gabriel ZAMBRANO REY is supported financially by the scholarship program "Francisco José de Caldas – Generación del Bicentenario" of the Administrative Department of Science, Technology and Innovation – COLCIENCIAS and Javeriana University (Bogotá, Colombia). In addition, this paper is part of the "Joint Research Lab of Risk Management in Life Critical Systems" project funded by the Partner University Fund of the U.S. French Embassy.

References

[1] G. D. M. Serugendo, "Robustness and Dependability of Self-Organizing Systems - A Safety Engineering Perspective," in *Stabilization, Safety, and Security of Distributed Systems*, R. Guerraoui and F. Petit, Eds. Springer Berlin Heidelberg, pp. 254–268, 2009.

[2] A. Tharumarajah, A. J. Wells, and L. Nemes, "Comparison of emerging manufacturing concepts," in *IEEE Int. Conf. on Systems, Man, and Cybernetics*, pp. 325 – 331, 1998.

[3] S. Bussmann and K. Schild, "Self-organizing manufacturing control: An industrial application of agent technology," in *MultiAgent Systems. Proc. Fourth Int. Conf. on*, pp. 87–94, 2000.

[4] J. M. Swaminathan, S. F. Smith, and N. M. Sadeh, "Modeling Supply Chain Dynamics: A Multiagent Approach*," *Decis. Sci.*, vol. 29, no. 3, pp. 607–632, 1998.

[5] Bo Chen and H. H. Cheng, "A Review of the Applications of Agent Technology in Traffic and Transportation Systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 11, no. 2, pp. 485–497, 2010.

[6] M. Carvalho and C. Perez, "An evolutionary multi-agent approach to anomaly detection and cyber defense," in *Proc. of the 7th Annu. Workshop on Cyber Security and Information Intelligence Research*, p. 28, 2011.

[7] D. Isern, D. Sánchez, and A. Moreno, "Agents applied in health care: A review," *Int. J. Med. Inf.*, vol. 79, no. 3, pp. 145–166, 2010.

[8] R. H. Bordini, M. Fisher, and M. Sierhuis, "Analysing human-agent teamwork," in *10th ESA Workshop on Advanced Space Technologies for Robotics and Automation, Noordwijk, The Netherlands*, 2008.

[9] T. Berger, Y. Sallez, S. Raileanu, C. Tahon, D. Trentesaux, and T. Borangiu, "Personal Rapid Transit in an open-control framework," *Comput. Ind. Eng.*, vol. 61, no. 2, pp. 300–312, 2011.

[10] F. Heylighen and C. Gershenson, "The meaning of self-organization in computing," *IEEE Intell. Syst. Sect. Trends Controv.-Self-Organ. Inf. Syst.*, 2003.

[11] M. Sabar, B. Montreuil, and J.-M. Frayret, "A multi-agent-based approach for personnel scheduling in assembly centers," *Eng. Appl. Artif. Intell.*, vol. 22, no. 7, pp. 1080–1088, 2009.

[12] W. Xiang and H. P. Lee, "Ant colony intelligence in multi-agent dynamic manufacturing scheduling," *Eng. Appl. Artif. Intell.*, vol. 21, no. 1, pp. 73–85, 2008.

[13] F. Gao, "Modeling teamwork of multi-human multi-agent teams," in *Procs. of the 2013 conference on Computer supported cooperative work companion*, New York, USA, pp. 47–50, 2013.

[14] M. Ball and V. Callaghan, "Explorations of Autonomy: An Investigation of Adjustable Autonomy in Intelligent Environments", *8th Int. Conf. on Intelligent Environments*, vol., no., pp.114,121, 26-29, 2012.

[15] E. DeKoven and A. K. Murphy, "A Framework for Supporting Teamwork between Humans and Autonomous Systems," DTIC Document, 2006.

[16] J. Hatvany, "Intelligence and cooperation in heterarchic manufacturing systems," *Robot. Comput.-Integr. Manuf.*, vol. 2, no. 2, pp. 101–104, 1985.

[17] J. Tweedale, N. Ichalkaranje, C. Sioutis, B. Jarvis, A. Consoli, and G. Phillips-Wren, "Innovations in multi-agent systems," *J. Netw. Comput. Appl.*, vol. 30, no. 3, pp. 1089–1115, 2007.

[18] G. A. Boy, "Orchestrating Situation Awareness and Authority in Complex Socio-technical Systems," in *Complex Systems Design & Management*, M. Aiguier, Y. Caseau, D. Krob, and A. Rauzy, Eds. Springer Berlin Heidelberg, pp. 285–296, 2013.

[19] J. M. Bradshaw, A. Acquisti, J. Allen, M. R. Breedy, L. Bunch, N. Chambers, P. Feltovich, L. Galescu, M. A. Goodrich, and R. Jeffers, "Teamwork-centered autonomy for extended human-agent interaction in space applications," in *AAAI 2004 Spring Symposium*, pp. 22–24, 2004.

[20] M. M. Carvalho, T. C. Eskridge, L. Bunch, J. M. Bradshaw, A. Dalton, P. Feltovich, J. Lott, and D. Kidwell, "A Human-Agent Teamwork Command and Control Framework for Moving Target Defense (MTC2)," in *Procs. of the 8th Cyber Security and Information Intelligence Research Workshop*, Oak Ridge, USA, 2012.

[21] G. Maione and D. Naso, "Adaptation of multi-agent manufacturing control by means of genetic algorithms and discrete event simulation," in *IEEE Intl. Conf. on Systems, Man and Cybernetics*, vol. 4, p. 6, 2002.

[22] G. Maione and D. Naso, "A genetic approach for adaptive multiagent control in heterarchical manufacturing systems," *IEEE Trans. Syst. Man Cybern. Part Syst. Humans*, vol. 33, no. 5, pp. 573 – 588, 2003.

[23] G. Metan, I. Sabuncuoglu, and H. Pierreval, "Real time selection of scheduling rules and knowledge extraction via dynamically controlled data mining," *Int. J. Prod. Res.*, vol. 48, no. 23, pp. 6909–6938, 2010.

[24] P. Leitão, J. Barbosa, and D. Trentesaux, "Bio-inspired multi-agent systems for reconfigurable manufacturing systems," *Eng. Appl. Artif. Intell.*, vol. 25, no. 5, pp. 934–944, 2012.

[25] N. Zbib, C. Pach, Y. Sallez, and D. Trentesaux, "Heterarchical production control in manufacturing systems using the potential fields concept," *J. Intell. Manuf.*, vol. 23, no. 5, pp. 1649–1670, 2012.