ADCHEM 2006

International Symposium on Advanced Control of Chemical Processes
Gramado, Brazil – April 2-5, 2006

ADAPTIVE AGENT-BASED CONTROL OF PRODUCT GRADE TRANSITIONS IN REACTOR NETWORKS

Eric Tatara, Cindy Hood, Fouad Teymour and Ali Cinar *
Michael North **

* *Illinois Institute of Technology, Chicago, IL*
** *Argonne National Lab, Argonne, IL*

Abstract: Large-scale spatially distributed systems provide control challenges because of their nonlinearity, spatial distribution and generally high order. The control structure for these systems tend to be both discrete and distributed. A layered control structure interfaced with complex arrays of sensors and actuators provides a flexible supervision and control system that can deal with local and global challenges. An adaptive agent-based control structure is presented whereby local control objectives may be changed in order to achieve the global control objective. Information is shared through a global knowledge environment that promotes the distribution of ideas through reinforcement. The performance of the agent-based control approach is illustrated in a case study where the interaction front between two competing autocatalytic species is moved from one spatial configuration to another. The multi-agent control system is able to effectively explore the parameter space of the network and intelligently manipulate the network flow rates such that the desired spatial distribution of species is achieved.

Keywords: Agent-based control, distributed systems.

1. INTRODUCTION

Large-scale spatially distributed systems provide a difficult control challenge because of their nonlinearity, spatial distribution and generally high order.

The control structure for these systems tend to be distributed and contain discrete and continuous elements. Hybrid control systems that combine process dynamics and discrete control elements and include multiple models for different operating points are one way to develop control systems for spatially distributed systems (Morari *et al.*, 2003; Christofides and El-Farra, 2005). An alternative approach is based on a hierarchical agent-based system with local and global control structures (Tatara *et al.*, 2005*c*; Tatara *et al.*, 2005*b*) that has been demonstrated on a network of interconnected continuous stirred tank reactors (CSTRs). Reactor networks exhibit highly complex behavior with multiple steady state operating regimes and have a large pool of candidates for manipulated variables (Tatara *et al.*, 2004).

The operation of highly nonlinear systems like autocatalytic replicator networks may benefit from evolutionary self-organizing control because the optimal operating regime and the required control strategies may not be known a priori. Agent-based control systems provide the capability for localized and global control strategies that are both reactive in controlling disturbances and proactive in searching for better operational solutions (Jennings and Bussmann, 2003). An adaptive agent-based control system for a CSTR network is proposed. The performance of the agent-based control approach is illustrated in a case study where the interaction front between competing autocatalytic species is moved from one spatial configuration to another.

2. AGENT-BASED CONTROL FRAMEWORK

Multi-agent control system architectures have several properties that make them particularly attractive for use in supervising large, complex systems (Lesser, 1999). The first, and usually most important in critical systems, is a high level of reliability. Modularity, scalability and adaptability are also attractive features of multi-agent systems. The adaptive and self-regulatory nature of agent systems has only recently been investigated for solving control problems that are normally solved with traditional methods.

2.1 Design Process

The design procedure used is a derivative of recent agent design methodologies based on the concept of the agent-services-acquaintance model (Wooldridge *et al.*, 1999) and the application to manufacturing control (Brueckner *et al.*, 1998). The goal of the design process is to develop an agent-based control system for physically distributed industrial processes. Certain parts of the agent-based control system are generic because they are based on general concepts of industrial control systems and operation of distributed processes.

Comprehensive studies of the physical process domain provide information regarding the expected normal operating conditions of the the processes, types of faults and disturbances that may occur, and control strategies. Additionally, the desired process operation and/or optimal conditions are expected to be known by the designers. Required agent types and roles are identified based on the requirements for controlling the physical system. The details of the hierarchical agent-based architecture (Tatara *et al.*, 2005*b*; Tatara *et al.*, 2005*a*) will not be repeated in detail here. The focus will rather be on the specific agent synthesis and instantiation for the presented examples.

2.2 Agent synthesis

There is nearly a one-to-one mapping of roles to agent types. The number of control agents is variable, depending on the number of reactors (I) in the network, as well as the complexity of the control actions being performed. Generally, a single control agent can be used for controlling each reactor. While multiple control agents can be applied to each reactor, from a software design point of view, it makes more sense to encapsulate the functionality of several control concepts (ie, temperature, level, etc.) into a single software agent, as long as the control algorithms are not extraordinarily complex.

The number of arbitration agents is probably the most flexible variable in the agent model. As with the control agents, the run time environment will set the number of arbitrators required. Larger networks with more control agents will subsequently require more arbitrators to handle the setpoint change requests coming from the controllers. Ultimately, the number of realized arbiters will be determined by the supervisory level agents.

The simplest implementation will have at least one supervisory level agent to coordinate efforts between the control and arbitration agents. While local interactions between agents are intended to serve as the primary driving force in the control system, the supervisor needs to maintain an overview at all times. More complex control schemes can include multiple supervisors for each high-level function such as help in mediating disputes, setting spatial concentration patterns, and supervising process recovery from disturbances that are too complex for the local controllers.

Finally, there will be exactly one instance each of the data collection and data acquisition agents. The data acquisition essentially functions as a bridge between the agent system and the physical domain. The acquisition agent will read values from either a simulator or a hardware data acquisition system and write the numeric values to objects that can be read and manipulated by the control and arbitration agents. The data collection agent encapsulates the roles for both data collection and file I/O since these roles share very similar tasks. Any agent in the system will likely have a small memory space for storing local information relative its specific tasks. The data collection agent will, however, be responsible for cataloging relevant data for the entire network, such as average concentrations, or possible even the concentration histories in each reactor. This data will be written to a file stream in chunks at some variable rate.

2.3 Global Knowledge Environment

Considering the nonlinearity of reactor networks, it is difficult to predict how the behavior of the system changes when the system parameters are manipulated. Consequently, one cannot easily predict how to change operating conditions of the network by manipulating the flow rates nor what

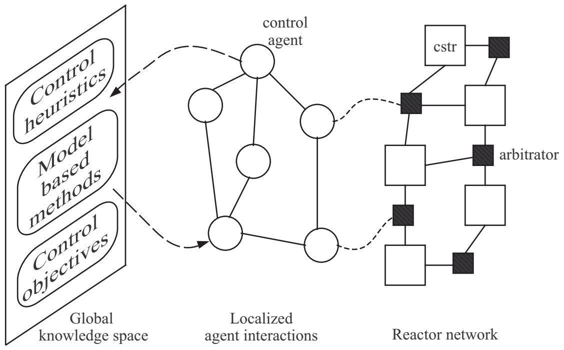

Fig. 1. Coupling of control agents with global knowledge space.

the localized operating conditions should be, in order to satisfy a global objective. Several methods can be used to guide the decision agents in planning their control strategies including dynamic exploration of the parameter space, rule-based heuristic models, or first-principles based models.

Although information is exchanged between agents via arbitrators, these interactions are local and limit the amount and quality that can propagate through the system. The global knowledge representation (Figure 1) serves as an environment for indirect communication between agents. This concept builds upon the hierarchical structuring of the control system by adding a mechanism for communication and reinforcement of ideas. The information in the knowledge space is divided into categories including local control objectives, control heuristics, and data-based models.

Information exchange occurs indirectly between agents because agents asynchronously read/write information from/to the knowledge space. For example, a particular agent may discover a local control strategy that works particularly well in meeting an objective set by a supervisor. This strategy is cataloged in the knowledge space by the originating agent. Other agents may read this strategy from the knowledge space and implement it to satisfy their particular control objective. The value of the strategy is then rated by the agents that adopt this new strategy such that its value relative to others is promoted. Similarly, outdated information in the knowledge space continuously decreases in value and eventually may be deleted from the knowledge space.

Although the stability of the agent dynamics cannot be guaranteed for every scenario, this methodology helps to reduce or prevent the emergence of dysfunctional agent dynamics by reinforcing "good" agent behavior, while punishing undesirable agent behavior. Furthermore, the agent system has been designed with the assumption that the agents' decision delays are small compared to the time scale of the physical process. The importance of this assumption becomes apparent when examining the consequences of its impact on process performance. If the agents' computing time is very long with respect to the process time scale, control of a continuous process becomes difficult due to the reduced data acquisition, control action computation and implementation rates. This assumption generally holds for chemical processes in which operating changes are introduced infrequently and process dynamics are represented with time scales of tens of minutes or even hours. Traditional controllers are normally used in the event of very rapid localized dynamics and, while the agents may modify the setpoints of such controllers, the time-critical (first response) control actions are strictly outside the domain of the higher-level decision making agents.

3. PRODUCT GRADE TRANSITION IN A CHEMICAL REACTOR NETWORK

Product grade transitions may be used to schedule the production of various composite compounds at different points in time. The overall product quality is determined by selecting one or more exit streams from the network and mixing them in the desired proportions. Therefore, if one would like to produce a grade consisting of the majority of one chemical species and only a small amount of another, it makes sense to mirror this grade composition in the network.

The supervisory level agent sets the grade composition by specifying the fraction of each species desired in the network. The composition is then transformed into a set of behaviors via the utility function that determines each control agent's goal. The fraction of desired species determines the agent's willingness to change or remain the same. If the agent's average composition value is set very high, and it surrounded by competitor agents, then its desire to remain unchanged is high. Conversely, if the agent's composition value is set very low and it is surrounded by cooperator agents, its desire to remain unchanged is low. This behavioral programming results in the network self-organizing to meet the global composition goal. A few isolated agents controlling reactors with a trace compound will contribute a lesser amount of their species to the product grade, while many closely clustered groups of the majority will contribute to the primary product in the desired grade. This self-organizing behavior arises from some aspects of the rules governing the local interactions as well as the open loop behavior or the network, since local clusters of one autocatalytic species will be more stable than single isolated reactors.

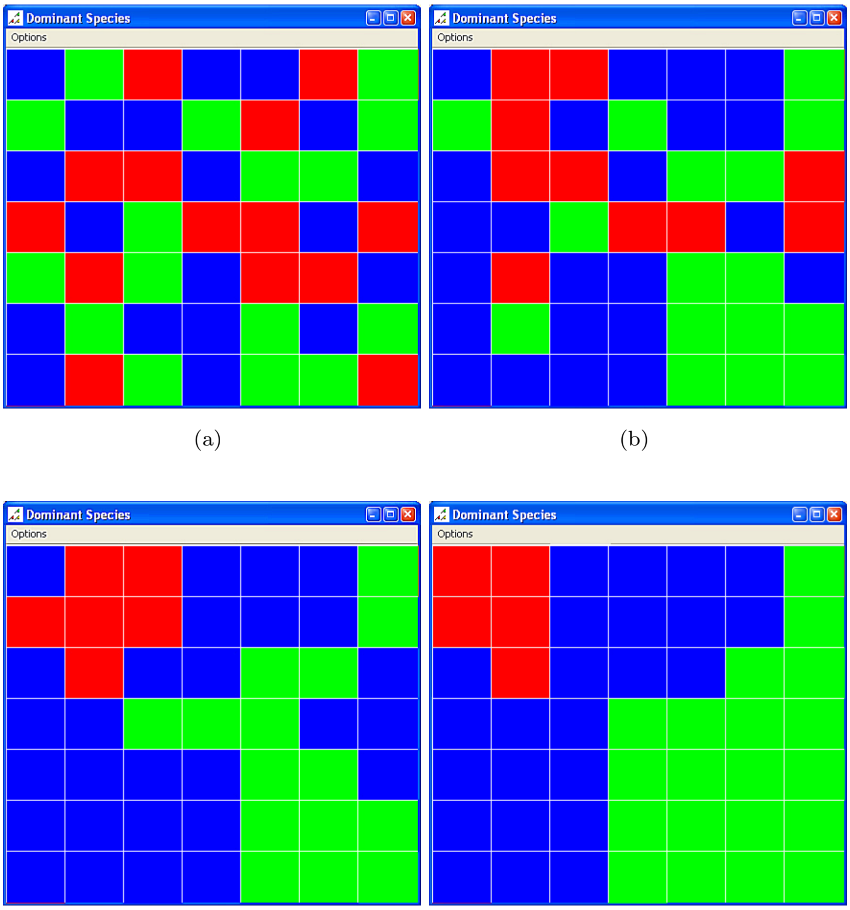

🟦	🟩	🟥	🟦	🟦	🟥	🟩
🟩	🟦	🟦	🟩	🟥	🟦	🟩
🟦	🟥	🟥	🟦	🟩	🟩	🟦
🟥	🟦	🟩	🟥	🟥	🟦	🟥
🟩	🟥	🟩	🟦	🟥	🟥	🟦
🟦	🟩	🟦	🟦	🟩	🟦	🟩
🟦	🟥	🟩	🟦	🟩	🟩	🟥

(a)

🟦	🟥	🟥	🟦	🟦	🟦	🟩
🟩	🟥	🟦	🟩	🟦	🟦	🟩
🟦	🟥	🟥	🟦	🟩	🟩	🟥
🟦	🟦	🟩	🟥	🟥	🟦	🟥
🟦	🟥	🟦	🟦	🟩	🟩	🟦
🟦	🟩	🟦	🟦	🟩	🟩	🟩
🟦	🟦	🟦	🟦	🟩	🟩	🟩

(b)

🟦	🟥	🟥	🟦	🟦	🟦	🟩
🟥	🟥	🟥	🟦	🟦	🟦	🟩
🟦	🟥	🟦	🟦	🟩	🟩	🟦
🟦	🟦	🟩	🟩	🟩	🟦	🟦
🟦	🟦	🟦	🟦	🟩	🟩	🟦
🟦	🟦	🟦	🟦	🟩	🟩	🟩
🟦	🟦	🟦	🟦	🟩	🟩	🟩

(c)

🟥	🟥	🟦	🟦	🟦	🟦	🟩
🟥	🟥	🟦	🟦	🟦	🟦	🟩
🟦	🟥	🟦	🟦	🟦	🟩	🟩
🟦	🟦	🟦	🟩	🟩	🟩	🟩
🟦	🟦	🟦	🟩	🟩	🟩	🟩
🟦	🟦	🟦	🟩	🟩	🟩	🟩
🟦	🟦	🟦	🟩	🟩	🟩	🟩

(d)

Fig. 2. Evolution of 2D spatial concentration profile of the dominant species in each reactor, with red (dark gray) representing species 1, blue (black) species 2, and green (light gray) species 3.

The performance of the agent-based control architecture is demonstrated in a case study to control the distribution of autocatalytic species in a network model (Tatara *et al.*, 2004) of 49 (7x7 grid) reactors hosting three autocatalytic species using the interaction flow rates as the manipulated variables. The species that populate the reactor network are characterized by identical growth and death rates, such that one species does not have an unfair advantage over the others. The reactor network model and agent-based control system is implemented with the open source Java agent modeling and simulation environment RePast (Collier *et al.*, 2003; Tatara *et al.*, 2005*a*).

The map of the dominant species in each reactor, with red (dark gray) representing species 1, blue (black) species 2, and green (light gray) species 3. The map shows the distribution of each species in the network as is updated whenever the state of one of the reactors change, ie the dominant species switches from one species to another.

The spatial distribution of autocatalytic species shown in Figure 2a shows the open loop behavior of the system following the temporal evolution from a randomly generated set of initial conditions. The supervisory agent responsible for the product grade composition of the network sets the grade for species (1,2,3) = (0.1,0.5,0.4); Once the control agents receive the specification for the grade composition (instantaneously, unless a higher priority action exists), they initiate the routine for determining the species setpoint based on their neighborhood and grade information. Figure 2b shows the temporal evolution of the network after 20 time (dimensionless) steps. Species 2 and 3 begin to form clusters at the bottom of the network. The aggregation of species in clusters occurs due to the natural open-loop behavior of the network in addition to the control agents, which reinforce the species robustness in a given reactor when surrounded by like species. After 40 time steps (Figure 2c), the network grade composition begins to approach the setpoint and finally arrives to the desired grade after 100 time steps (Figure 2d). The exact spatial boundaries between the species clusters may fluctuate with time if the network displays oscillatory behavior or is subjected to a disturbance. However, the network average values for the species concentrations will

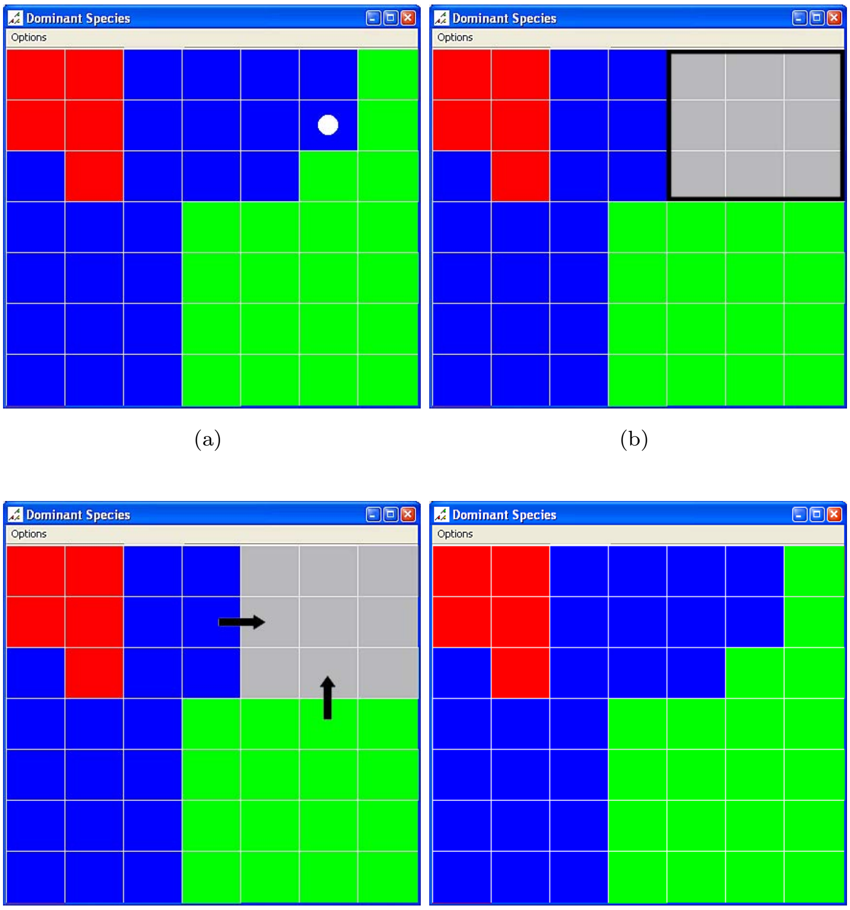

(a) (b)

(c) (d)

Fig. 3. Agent response to species invasion in reactor marked with white circle, with red (dark gray) representing species 1, blue (black) species 2, green (light gray) species 3, and very light gray representing an unpopulated reactor.

remain constant due to local spatial reconfigurations, thereby meeting the grade specification.

4. DISTURBANCE REJECTION

Figure 3 shows a 7x7 grid of reactors hosting three autocatalytic species. The spatial distribution of species is an arbitrary choice, but is treated as the desired operating condition for the network. Therefore, the control agent for each reactor should prevent the invasion of the reactor for which it is responsible. For neighboring competitors, this control function can be achieved by the previously detailed methods of local agent control via the interaction rates (Tatara *et al.*, 2005*b*). However, in the case where a new, more aggressive autocatalytic species is introduced to the network, these methods may fail. A species with a higher ratio of growth to death rates than those of pre-existing species will quickly dominate a contaminated reactor and spread throughout the network, effectively eliminating all other species.

A method for containment of aggressive invading species is modeled after the concept of programmed cell death (PCD) in plants (Liu *et al.*, 2005), in which cells that are infected with viruses die to prevent the infection from spreading to the rest of the plant. The immune system in plants is sensitive to invading species, namely viruses, which causes a chemical reaction in the infected plant cells resulting in their death. The PCD mechanism is encapsulated in the genetics for each cell, yet is only activated when the cells become infected. The mechanism is sufficiently robust to quickly contain infections, yet not so sensitive as to be pathological.

The application to disturbance rejection in the reactor network is straightforward. When a control agent detects the presence of an unknown species above some minimal concentration threshold (for example reactor (2,6) of Figure 3a, it isolates itself from neighboring reactors to prevent the spread of the invading species (Figure 3b). The controller manipulates the residence time of the reactor such that the invading species is washed out below the detection threshold. This procedure may in fact result in the other species being completely washed out as well. However, the containment of the invading species takes precedence to maintaining the original control objective, once the reactor has been compromised. After the invader has been washed out, the reactor can be reconnected to the network and the autocatalytic species can be repopulated using the feed stream or, preferably, resources from adjacent reactors (Figure 3c), such that the original spatial configuration is restored (Figure 3d).

5. CONCLUSIONS

An adaptable, intelligent agent-based control system has been implemented to control the product grade transitions via spatial distribution of autocatalytic species in a reactor network. This methodology has been proposed as a real-time alternative to traditional nonlinear control schemes involving predetermined controller configurations or computationally expensive optimization techniques.

Controlling the overall spatial distribution of autocatalytic species in a network of reactors requires a coordinated effort between local control agents. The practical application of this control scheme is in managing the composition of a product composition. From a production standpoint, the net feed flow rate through the network, that is the total sum of all feed flow and exit flow rates, should be kept constant. Therefore, there is a constraint on the total network inflow and outflow. However, the distribution of interconnection flow rates may be distributed as necessary, as long as the volume constraints are maintained. The distribution of interconnection flows will not be set by a supervisory level agent, since the agent's control algorithm performance would deteriorate quickly as the network is scaled up. Therefore, the set point changes for the interconnection flow rates will be determined by coordination of the local control agents.

In addition to case studies demonstrating the behavior of the agent-based control system under normal operating conditions, methods were designed and implemented to allow the framework to correct disturbances in the form of invading species or reactor malfunction. The disturbance rejection behavior was based on the immune system response of plant species to viral infections. The control agents were successfully able to detect invading autocatalytic species and isolate infected reactors to prevent the invading species from spreading through the network. The infected reactors can be repopulated with the desired species once the invader is washed out of the network.

6. ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. CTS-0325378 of the ITR program.

REFERENCES

Brueckner, S., J. Wyns, P. Peeters and M. Kollingbaum (1998). Designing agents for manufacturing control. In: *Proceedings of the 2nd AI & Manufacturing Research Planning Workshop*. Albuquerque, NM.

Christofides, P. and N. El-Farra (2005). *Control of Nonlinear and Hybrid Process Systems*. Springer-Verlag: Berlin.

Collier, N., T. Howe, and M. North (2003). Onward and upward: The transition to repast 2.0. In: *First Annual North American Association for Computational Social and Organizational Science Conference, Pittsburgh, PA.*

Jennings, N.R. and S Bussmann (2003). Agent-based control systems. *IEEE Control Systems* **23**, 61–73.

Lesser, V.R. (1999). Cooperative multiagent systems: A personal view of the state of the art. *IEEE Transactions on Knowledge and Data Engineering* **11**(1), 133–142.

Liu, Y., M. Schiff, K. Czymmek, Talloczy Z., B. Levine and S.P. Dinesh-Kumar (2005). Autophagy regulates programmed cell death during the plant innate immune response. *Cell* **121**, 567–577.

Morari, M., M. Baotic and F. Borrelli (2003). Hybrid systems modeling and control. *European Journal of Control* **9**(2–3), 177–189.

Tatara, E., F. Teymour and A. Cinar (2005*a*). Control of complex distributed systems with distributed intelligent agents. *submitted to Journal of Process Control.*

Tatara, E., I. Birol, F. Teymour and A. Cinar (2004). Static and dynamic behavior of autocatalytic replicators in reactor networks. *Industrial and Engineering Chemistry Research* **43**, 3972–3993.

Tatara, E., I. Birol, F. Teymour and A. Cinar (2005*b*). Agent-based control of autocatalytic replicators in networks of reactors. *Computers & Chemical Engineering* **29**, 807–815.

Tatara, E., I. Birol, F. Teymour and A. Cinar (2005*c*). Measuring complexity in reactor networks with cubic autocatalytic reactions. *Industrial and Engineering Chemistry Research* **44**, 2781–2791.

Wooldridge, M., N.R. Jennings and D. Kinny (1999). A methodology for agent-oriented analysis and design. In: *Proceedings of the Third International Conference on Autonomous Agents (Agents'99).* Seattle, WA.