Integrating Intelligence and Real-Time Control into Manufacturing Systems

(Extended Abstract)

David J. Musliner Edmund H. Durfee Kang G. Shin

Computer Science and Engineering Division
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, Michigan 48109-2122

{djm,durfee,kgshin}@eecs.umich.edu
(313) 936-2495

Among other requirements, intelligent manufacturing systems must provide methods by which the low-level control of ongoing manufacturing processes can be subjected to the higher-level influence of intelligent processing (either human or artificial). That is, the real-time control of individual devices must be modulated both by considerations of dependencies between devices and by broader considerations of longer-term goals and expectations.

Our research has focused on the problems that arise when trying to integrate classical AI planning methods with the rigid performance guarantees required by real-time domains. Manufacturing domains clearly require control systems that provide real-time response guarantees to ensure, for example, that machines on an assembly line will process arriving parts in a timely fashion and avoid dropping, missing, or damaging materials. At the same time, the move towards flexible, intelligent manufacturing has made it imperative that these systems have the ability to deal with dynamic or uncertain environments that may include changing production goals and schedules, changing deadlines, uncertain part positions, etc. Unfortunately, the various AI and scheduling methods that have been developed to deal with these types of difficult problems are not suited to real-time guarantees. These methods generally involve heuristic search in exponential search spaces, so that, in the worst case, their processing time requirements are exponential. As a result, it is not possible to allocate sufficient computational resources to guarantee that these large-scale search problems can be solved within rigid deadlines.

To address this problem of integrating real-time and AI processing, we have investigated the Cooperative Intelligent Real-time Control Architecture (CIRCA) [1, 2], which is designed to support both hard real-time response guarantees and unrestricted AI methods that can guide those real-time responses. Figure 1 illustrates the architecture, in which an AI subsystem (AIS) reasons about high-level problems that require its powerful but uncertain reasoning methods, while a separate real-time subsystem (RTS) uses its predictable performance characteristics to deal with low-level problems that require guaranteed response times.

Prototype implementations of CIRCA have been applied to two domains. In the first version, the system piloted a Hero 2000 mobile robot through the hallways of our building, providing guaranteed collision avoidance reactions and search-based navigation. A newer version of the system controls a simulated

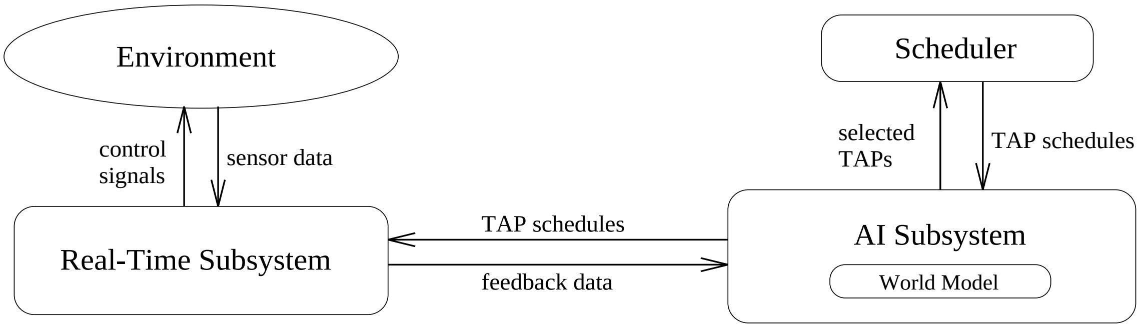

Figure 1: The Cooperative Intelligent Real-Time Control Architecture.

Puma robot arm, which must pack parts arriving on a conveyor belt into a nearby box. The parts can have several shapes (e.g., square, rectangle, triangle), each of which requires a different packing strategy. The control system may not initially know how to pack all of the possible types of parts– it may have to perform some search algorithm to derive an appropriate box-packing strategy. The robot arm is also responsible for reacting to an emergency alert light. If the light goes on, the system must push the button next to the light before a fixed deadline.

CIRCA's goal is to be "intelligent about real-time," as opposed to being "intelligent in real-time." That is, CIRCA's AI processing is not constrained to meet deadlines. Instead, the RTS is responsible for executing reactions that are guaranteed to meet the domain's hard deadlines, while the AIS executes less-predictable search algorithms that address higher-level problems without hard deadlines. For example, in our Puma arm domain, the RTS is programmed (by the AIS) with a set of reactions (cast as test-action pairs, or TAPs) that are known to respond in time to emergency alerts and to the arrival of parts on the conveyor belt. While the RTS is executing those reactions, ensuring that the system avoids failure, the AIS is able to execute high-variance heuristic search methods to find the next appropriate set of reactions. In the example domain, the AIS may derive a new box-packing algorithm that can handle a new type of arriving part. The derivation of this new algorithm does not need to meet a hard deadline, because the reactions concurrently executing on the RTS will continue handling all arriving parts, just stacking unfamiliar ones on a nearby table temporarily. When the new box-packing algorithm has been developed and integrated with additional reactions that prevent failure, the new schedule of reactions can be downloaded to the RTS. Thus CIRCA is able to apply unrestricted AI methods to difficult, high-level problems while also guaranteeing low-level control responses that will meet deadlines.

Our investigations of CIRCA to date have focused on two main features. First, we have developed a scheduling module and a structured interface that allow the unconstrained AI subsystem to asynchronously direct the real-time subsystem without violating any response-time guarantees. The scheduling module is given information about the resources available to the RTS, and attempts to build guaranteed schedules of reactions that are suggested by the AIS. If the RTS resources are not sufficient to guarantee all of the desired reactions, the AIS and Scheduler engage in an iterative process of trading off various performance characteristics in exchange for lowered resource requirements. When a schedule is finally produced by this cooperative reasoning, the AIS can download the new schedule to the RTS.

The AIS/RTS interface relies on incremental, non-blocking communication that is scheduled explicitly within the test-action pairs executed by the RTS. Thus communication in and out of the RTS is predictable, and cannot cause any guaranteed RTS reactions to miss their deadlines; the RTS never waits for communication from the AIS. Switching execution from one schedule of reactions to another is likewise a very rapid and predictable operation which is accounted for in the RTS reaction schedules.

Our second research focus has been the methods used by the AIS to derive the set of reactions that should be executed by the RTS. The AIS reasons about an internal model of the world and the actions that the RTS can take to sense and affect the world. We have developed a formal definition of this model of agent/environment interactions, and the assumptions on which the model is based [2]. Within the context of this model, we have shown how CIRCA's reactive control plans can be proven to simultaneously guarantee the system's safety and achieve its high-level goals. We have implemented an algorithm that simultaneously builds up the set of possible world states and plans reactions as needed to avoid system failures and achieve the system's goals. Together with the scheduling module, this world modeling and planning mechanism essentially automates the generation of a stable real-time control plan tailored to the expected world states.

In summary, our research on CIRCA represents a first step towards an integrated system supporting both the real-time response guarantees required for low-level control and the uncertain, search-based AI algorithms used to address high-level, goal-directed planning and scheduling. Integrating these capabilities will be an essential feature of future intelligent manufacturing systems.

References

[1] D. J. Musliner, E. H. Durfee, and K. G. Shin, "CIRCA: A Cooperative Intelligent Real-Time Control Architecture," to appear in *IEEE Trans. Systems, Man, and Cybernetics*, vol. 23, no. 6, 1993.

[2] D. J. Musliner, E. H. Durfee, and K. G. Shin, "World Modeling for the Dynamic Construction of Real-Time Control Plans," submitted to *Artificial Intelligence*, 1993.

NOTE: These papers are available on-line via anonymous ftp to ftp.eecs.umich.edu; see the file "outgoing/djm/README."

David J. Musliner's Biosketch

David J. Musliner received the B.S.E. degree with high honors in Electrical Engineering and Computer Science from Princeton University, Princeton, New Jersey, in 1988, and is currently a Ph.D. Candidate in the Department of Electrical Engineering and Computer Science at the University of Michigan, Ann Arbor, Michigan. While at Princeton, he received the Charles Ira Young Award for excellence in electrical engineering research. His graduate work has been supported by fellowships from the National Science Foundation and the University of Michigan. His research interests include artificial intelligence, real-time intelligent control, robotic application domains, and computer supported cooperative work. He is a member of IEEE, ACM, AAAI, Sigma Xi, and Tau Beta Pi.

Edmund H. Durfee's Biosketch

Edmund H. Durfee received the A.B. degree in chemistry and physics from Harvard University, Cambridge, Mass., in 1980, the M.S. degree in electrical and computer engineering and the Ph.D. degree in computer and information science from the University of Massachusetts, Amherst, Mass., in 1984 and 1987, respectively. His Ph.D. research developed an approach for planning coordinated actions and interactions in a network of distributed AI problem-solving systems.

He is currently an Assistant Professor in the Dept. of Electrical Engineering and Computer Science at the University of Michigan, where his interests are in distributed artificial intelligence, planning, blackboard systems, and real-time problem solving. He has published extensively in these areas, and is author of the book *Coordination of Distributed Problem Solvers* (Kluwer Academic Press). In his most recent work, he has been designing a framework for coordination based on hierarchical, multi-dimensional behavior specifications, and he has been developing an integrating architecture for combining real-time and intelligent systems. He is a 1991 recipient of a Presidential Young Investigator award from the National Science Foundation.

Prior to joining the faculty at the University of Michigan, Dr. Durfee was a Research Computer Scientist in the Department of Computer and Information Science at the University of Massachusetts. He is an associate editor for *IEEE Transactions on Systems, Man, and Cybernetics*, and has served on a number of conference and workshop program committees, including co-chairing the 1992 Distributed AI Workshop. He is a member of the IEEE Computer Society, the Association for Computing Machinery, AAAI, and AAAS.

Kang G. Shin's Biosketch

Kang G. Shin is Professor and Associate Chair of Electrical Engineering and Computer Science for the Computer Science and Engineering Division, The University of Michigan, Ann Arbor, Michigan.

He has authored/coauthored over 240 technical papers (more than 100 of these in archival journals) and several book chapters in the areas of distributed real-time computing and control, fault-tolerant computing, computer architecture, robotics and automation, and intelligent manufacturing. In 1987, he received the Outstanding IEEE Transactions on Automatic Control Paper Award for a paper on robot trajectory planning. In 1989, he also received the Research Excellence Award from The University of Michigan. In 1985, he founded the Real-Time Computing Laboratory, where he and his colleagues are currently building a 19-node hexagonal mesh multicomputer, called **HARTS**, to validate various architectures and analytic results in the area of distributed real-time computing.

He received the B.S. degree in Electronics Engineering from Seoul National University, Seoul, Korea in 1970, and both the M.S. and Ph.D degrees in Electrical Engineering from Cornell University, Ithaca, New York in 1976 and 1978, respectively. From 1978 to 1982 he was on the faculty of Rensselaer Polytechnic Institute, Troy, New York. He has held visiting positions at the U.S. Airforce Flight Dynamics Laboratory, AT&T Bell Laboratories, Computer Science Division within the Department of Electrical Engineering and Computer Science at UC Berkeley, and International Computer Science Institute, Berkeley, CA.

He is an IEEE fellow, was the Program Chairman of the 1986 IEEE Real-Time Systems Symposium (RTSS), the General Chairman of the 1987 RTSS, the Guest Editor of the 1987 August special issue of *IEEE Transactions on Computers* on Real-Time Systems, a Program Co-Chair for the 1992 *International Conference on Parallel Processing*, and served numerous technical program committees. He also chaired the IEEE Technical Committee on Real-Time Systems during 1991-93, is a Distinguished Visitor of the Computer Society of the IEEE, an Editor of *IEEE Trans. on Parallel and Distributed Computing*, and an Area Editor of *International Journal of Time-Critical Computing Systems*.