# Center For Integrated Manufacturing Decision Systems

## An Overview

**RamaKrishna Akella, Mark S. Fox & Stephen Smith**

Center for Integrated Manufacturing Decision Systems
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

## 1. Introduction

In order for American industry to compete in the global market, it must continue to enhance both its product quality and organizational efficiency. Enhancing product quality is not limited to the factory floor, but begins with better design decisions, and continues throughout the production cycle. Similarly, increasing organizational efficiency must not only address the problem of how to utilize the greater flexibility provided by production automation, but how to make better, more integrated decisions throughout the enterprise. Manufacturing, computer, and communication technologies continue to evolve at a rapid rate, yet the cornerstone of these technologies continues to lag in its development: software to support and make decisions. If industry is to fully utilize these technologies, there must be a major effort in developing the next generation of decision making software for the planning and control of the enterprise.

The Center for Integrated Manufacturing Decision Systems has been created to perform research in intelligent decision systems for engineering and manufacturing problem solving. The Center's research has two goals: to increase the quality of decisions at each stage of the production manufacturing life cycle—design, planning, production, distribution and field service, and to achieve an order of magnitude increase in decision making quality by focusing on planning and control systems, which integrate decision making throughout the enterprise.

Successful pursuit of the Center's research agenda will facilitate: optimum product/process/facility design, automated machine set-up and control (rapid/accurate), optimum scheduling (dynamic), and fast, flexible response to changing requirements (Parts-on-Demand). In other words, highly competitive production systems.

Historically, the computerization of manufacturing has proceeded in stages. The first generation introduced flexible manufacturing systems; the second generation focused on integration of manufacturing through use of communication networks and data sharing. In both generations, computerization has focused on the *implementation* of manually made decisions, and the less than satisfactory quality of those decisions has resulted in poor resource utilization and higher costs.

Decision quality must be improved, and therein computer integrated manufacturing has reached a turning point. Such an improvement can be made possible by providing better decision support and smarter decision making systems.

Systems that provide decisions or decision support (for planning, scheduling and control) can be considered a third CIM generation development. Looking beyond, the fourth generation will integrate decision making throughout the entire product manufacturing cycle.

The Center's goal is the realization of third and fourth generation computer integrated manufacturing through focused research which combines conventional and knowledge based systems. The integration of these projects will provide the necessary theories and methods to create fourth generation environments.

The center has created seven laboratories, each focusing on a portion of engineering/manufacturing decision making:

- Design: The rapid and integrated design of electro-mechanical parts.
- Production Planning: Integrating process and facility planning.
- Rapid Manufacturing: Automating unit process planning, programming and control.
- Automated Factory Scheduling: Predictive and reactive detailed scheduling.
- Manufacturing Logistics: Parts on demand.
- Intelligent Measurement and Control: Multiple process focus.
- Manufacturing System Architecture: Representation and communication of factory knowledge and control.

The rest of this paper describes issues and goals of for each laboratory.

## 2. Design Laboratory

In the design of mechanical parts and assemblies, much of design remains a manual task where the role of computers is limited to drafting (i.e., elaboration of design details) and validation (e.g., Finite element analysis). For a number of reasons, it is becoming necessary to provide decision support systems for the designer which enables a more efficient exploration of the space of design alternatives.

One reason is that the design of mechanical and electro-mechanical parts and assemblies is becoming more complex due to the complexity of functional specifications, increasing performance requirements, and the variety of materials and technologies with which to implement the design. Consequently, the generation of "good" designs is becoming more difficult because it exceeds the capabilities of any one person.

Secondly, design decisions usually focus on the translation of function into form and tend to ignore the impact of these decisions on activities performed downstream in the product life cycle:

- Planning,
- Fabrication,
- Assembly,
- Test,
- Distribution,
- Field Service,
- Training.

What may appear to be a good design from a functional perspective, may not be manufacturable, nor serviceable.

A third problem with design is that our ability to predict the performance of a design may be limited. In such a situation, it may be necessary to construct a prototype in order to test its behavior. The slowness with which a prototype may be constructed can dramatically impact cost, quality of result and the time to market.

Lastly, the design of complex systems tends involve more than one designer. It has been our experience that system designs tend to fail at the interface of component designs, and design projects tend to extend beyond their due dates because of the difficulty in coordinating the individual efforts.

Following describes the goals of the Design Laboratory.

***Design Theory.*** In the domain of mechanical and electro-mechanical design, a design theory has yet to evolve. We intend to develop such a theory, focusing on the representation of function and form, and reasoning about how functionality is synthesized out of form in the context of combined spatial, causal, and temporal actions. This theory will smoothly integrate variety of design styles, including:

- Selection,
- Parameterized design,
- Configuration,
- Extrapolation (incremental design),
- Routine design, and
- Discovery.

***Integrated Design.*** Design methodologies, to the extent that they exist, focus only on the transformation of function into form. Evaluation of designs with respect to their impact on downstream activities is usually in the form of a post design evaluation where the knowledge is in the form of rules. The postponement of this evaluation until after the design is completed increases design time while providing only partial guidance to the designer. Our goal is to extend the design theory mentioned above so that the constraining information derived from activities throughout the product life cycle can be used directly in the generation of new designs. The premise of the *Product Design For Integration* project is that by including knowledge of downstream activities into the functional design process, designs will be created which will satisfy functional specifications and optimize the performance of activities across the product life cycle.

***Rapid Design.*** A third goal of the project is to achieve rapid design in two senses. First, the time to generate an acceptable design should be small. This can be achieved by using knowledge of previous designs:

parameterizing the design of a class of products or making incremental changes to an existing design. In addition, the concept of integrated design reduces time by removing alternatives which contradict constraints introduced by downstream activities.

The second sense in which design can be made more rapid, is when the design process is *open*. In open design, there exists insufficient knowledge to predict the performance of a design. Consequently, a prototype must be constructed and tested, with the results being fed back into the design systems so that it can *learn* from its mistakes.

***Interactive Design.*** Our intent is to develop a mode of interaction in which either the designer or the system can control the design process. The design system will be able to operate in a mode of critiquing a design, guiding its development, or automatically generating when enough knowledge is available. It will be able to smoothly move between these modes as desired/required by the user and the state of the design. Secondly, it will be able to keep track of the design process so that the knowledge gained can be used in subsequent design situations.

***Managed Design.*** Experience in large system design, in which the design is spread across many engineers, has shown that often problems arise at the interface where individual designs are integrated rather than with the individual designs themselves. Many times the cause of these problems can be traced to changing definitions of design components and their associated design activities. In turn, these changes are not identified nor communicated in a timely manner to enable redirection of affected portions of the design activities. Simply, the complexity of distributed design exceeds our ability to manage the individual design activities so that the results can be integrated. In this project, we will explore the creation of a distributed design system in which the process and management of design is distributed throughout the system; the dynamics of a project, i.e., changing product specifications and activities, requires closer management of the design process.

## 3. Rapid Manufacturing Laboratory

Automated manufacturing is only made possible after large efforts at machine and process setup are expended by human experts. The human craftsman discovers potential problems by trial and error and then uses these experiences to design the system so that problems are explicitly avoided. Unfortunately, this problem avoidance approach neither accounts for unsuspected situations or new part designs that have different process requirements.

The craftsman's job is to dynamically match the process parameters to the ever changing requirements posed by different part designs. The adjustments must be made in a dynamic environment, under varying conditions: speeds and feeds must be adjusted under normal operating conditions, more drastic adjustments are required in marginal situations caused by worn tools or excessive vibrations, and a major planning effort is required when the machine breaks down due to broken tools or other unexpected events.

The laboratory will concentrate on developing an overall control architecture for machine tools that will eventually automate all of the craftsman's duties:

- Process planning
- Process programming
- Set up
- Monitoring
- Progress inspection
- Adaptive control
- Refixturing
- Final inspection

The machine controller that will result from our research must: understand its own limitations, understand the consequences of its actions, verify the results of actions, and have a rich model of the manufacturing domain so that machine actions can be intelligently planned. The automated craftsman must indeed have a broad range of skills. Even a partial list of these skills shows the complexities that must be coped with on a daily basis: Different batches of materials have different properties which cause many subtle problems in machining. The skill is to detect differences in materials and to make adjustments in the process before any scrap parts are made. There are many different tools that can be used for the same job, but each choice carries with it a set of advantages and disadvantages (e.g., more or fewer cuts, better or worse surface finishes, longer or shorter tool lives). The skill is to produce an optimal plan with the available resources before the process begins. While some part geometries are very easy to machine and can be done with the highest degree of confidence (e.g., low tolerance drill holes), other part geometries are very difficult to machine and require a great deal of forethought and follow-up inspection (e.g., very thin walls). The skill is to know how carefully to watch the process for given steps as well as which (sensor) sources provide the most revealing information about the process.

To implement the skills of the craftsman, a model with deep knowledge of manufacturing processes will be constructed. This model will provide causal explanations for each operation and sensation. By using the model, it will be possible to automate corrective actions in an imperfect manufacturing world, which have until now resisted automation.

As our research progresses, three successively more complete prototypes will be built that will document how both newly designed parts and replacement parts can be quickly fabricated with a minimum of human intervention. Each prototype will be tested on a range of aerospace parts made from hard-to-machine alloys with an increasing demand for high tolerances. The final prototype will then form the basis of a commercial product that will be developed by selected affiliate sponsor(s).

## 4. Production Planning Laboratory

Attempts are being made to increase production quality and efficiency through better control. While increases can be achieved in this manner, they are limited by the design of production. In order to achieve major improvements it is necessary to plan production with these goals in mind.

Production planning is a complex process consisting of:

- Process planning
- Process programming
- Facility selection and layout
- Quality planning
- Production software programming
- Demand forecasting
- Training

Process planning is concerned with the selection and sequencing of a set of processes which will result in the manufacture of a part. Traditionally, process planning has been performed by experts who have detailed knowledge of the processes required to manufacture a variety of features. Today, expert systems have been created to automate the process planning of printed wire boards and some sheet metal fabrication for aircraft, but the automation of process planning for complex three dimensional parts and assemblies is beyond the state of the art. In particular, the relationship between a 3D feature and the required processing, and the process of assembling parts is still poorly understood. Secondly, the variety of processes available to achieve the same result continues to expand.

In a similar sense, the design of manufacturing facilities is also an expert task in which simulation is used only to verify or refute design hypotheses. Poor facility selection and layout may result in costly capital investment, high work in process and tardiness in order delivery.

Process planning and facility design are tightly coupled. In most situations, the process planner must plan in the context of an existing manufacturing facility with specific processes and machines. Conversely, the selection of machines to implement a process plan require an understanding of the features to be produced and why a particular process was chosen. The layout of machines is also concerned with the process plan, the forecasted demand and other parts and their process plans. Both tasks share the same knowledge and goals.

Another problem facing planners is how to design quality into the production of a product. In the planning of the processes it is necessary to choose processes so that quality is maintained/increased. In addition, it is necessary to design quality monitoring and repair programs which can track quality, identifying the sources of deviations which may go beyond the four walls of the plant, and correct them.

A third growing problem is the cost of creating and maintaining software systems for production monitoring and control. As new products are introduced, new processes planned, and new quality programs created, it is necessary to update production software to reflect these changes. This task is fast becoming one of the most expensive in the computerized factory.

The first major focus of the *Production Planner* project will be on developing the theories and techniques for process planning and facility design so that they utilize and share the same knowledge. The generation of process plans will optimize the use of the facility design and conversely, the facility design and layout will allow for optimization of process plan generation. Concepts to be explored include:

- Planning techniques for process selection, sequencing and merging, and incorporation of knowledge of plans for similar parts in generating new plans.
- Viewing facility design as a design task where previous designs can be used to guide the generation of new designs.
- Capturing of expertise to provide quantitative and qualitative analysis and knowledge based simulation of facility designs.
- Adaptive repair of process plans and facility designs based upon their performance.

***Planning for Quality.*** In order to achieve greater levels of product quality, it is necessary to design it into the product and the manufacturing processes from the beginning, and to design methods to monitor and control it. As part of the process and facility design task we will explore how quality can be designed into the process.

***Planning for Control.*** Our goal is to explore what it means to build production control software which can be (semi-)automatically altered to reflect a changing environment. This will be done in conjunction with the factory scheduling project.

## 5. Automated Scheduling Laboratory

One of the major deterrents to productivity in US industry today is the inability to effectively coordinate production activities. Factory operations are routinely characterized by high work-in-process (WIP) inventories, tardy orders, poor resource utilization, and other shop floor inefficiencies. Perhaps the single most significant obstacle to improved factory performance is the complexity associated with constructing and maintaining good production schedules. To be useful in practice, a production schedule must reflect the influence of a large and conflicting set of requirements, objectives, and preferences:

- organizational goals - meeting due dates, minimizing WIP time, maximizing resource utilization, maintaining continuity in factory operations
- causal restrictions and functional limitations - operation precedence and separation constraints, resource requirements, resource setup procedures, resource capabilities and capacity restrictions
- resource availability - work shifts, resource maintenance schedules, dynamic restrictions on scheduling alternatives (e.g. machine breakdowns)
- operational preferences - setup minimization, machine reliability preferences, preferred production processes.

Existing computer-based techniques for production scheduling are capable of incorporating only a small portion of this knowledge, and consequently produce schedules that typically bear little relationship to the actual state of the factory. At best, these techniques provide human schedulers with high level decision-making guidelines. These guidelines, however, do little to reduce the complexity of the actual coordination problem.

The factory scheduling problem is further complicated by the unpredictability of factory operations. The factory floor is, in fact, a very dynamic environment. Machines break down, raw materials fail to arrive on time, partially manufactured parts fail to meet quality control standards and require rework, operators call in sick, etc. Thus, even an ability to produce "optimal" solutions (i.e. advance schedules that effectively balance conflicting objectives and preferences in accordance with the restrictions imposed by the production environment) is of limited utility without a companion ability to reactively manage these schedules in response to unanticipated circumstances. Existing computer-based scheduling techniques provide virtually no support for reactive decision-making; reactive scheduling decisions are invariably made myopically, with little or no understanding of the more global implications of these decisions. As a consequence, production schedules released to the factory floor quickly become obsolete, and little global coherence in factory floor decision-making is ever maintained.

The insufficiency of adopting a static view of the scheduling problem raises a larger issue relative to coordination of production activities in a particular manufacturing environment: What level of predictive guidance is most effective in producing a factory *behavior* that optimizes the satisfaction of all relevant constraints? KANBAN-style coordination, which relies only on guidance relating to what to "pull" out of the system next, produces effective results in well-structured environments (i.e. exhibiting predictable manufacturing processes and commonality of processes across products) with fairly steady demand patterns. In less structured environments, greater reliance must be placed on advance planning/scheduling to optimize factory performance. However, the level of precision of this guidance (i.e. the degree of decision-making flexibility) must be calibrated against the uncertainties in the production environment and the costs of maintaining this guidance over time. An effective balance between optimization and responsiveness to changing circumstances must be achieved. Current understanding of this tradeoff is minimal. Existing coordination practices based on the use of production schedules (e.g. MRP-based approaches) unduly sacrifice with respect to optimization, operating with suboptimal expectations of performance (e.g. standard lead times).

The *Automated Scheduling Laboratory* is conducting research aimed at providing effective solutions to real world factory scheduling problems. Through extension and integration of research in the fields of Artificial Intelligence (AI) and Operations Research (OR), we intend to formulate the theories and knowledge necessary to overcome the inadequacies of conventional approaches to factory scheduling and enable the creation of computer-based scheduling technologies that provide an integrated basis for factory floor decision-making. To this end, the goals of the Automated Scheduling Laboratory are to:

1. Define an automated scheduling methodology and supporting system architecture that effectively integrates predictive and reactive decision-making processes. We can identify three functional capabilities that the methodology and architecture must encompass:
    - realistic predictive scheduling - Effective coordination of factory production requires an ability to generate schedules that reflect the reality of the production environment. Schedules must be developed with respect to accurate models of the full range of constraints that influence production activities, using knowledge about these constraints to manage the complexity of optimization.
    - incremental reactive scheduling - As unexpected events (both problematic and serendipitous) occur on the factory floor, it is necessary to incrementally revise current schedules to fit current circumstances. Reactive scheduling processes must exploit knowledge of the current production state and characteristics of the current schedule to effectively balance schedule optimization, schedule continuity, and system responsiveness objectives.
    - execution-time schedule interpretation - Knowledge of the sources and degrees of uncertainty in the manufacturing environment should dictate the level of detail of different aspects of the maintained schedule. This typically implies the need for some amount of execution-time decision-making within the confines of the constraints imposed by the schedule, which must necessarily be based on knowledge of the uncertainty assumptions underlying the schedule.

    We are investigating knowledge representations, scheduling methods and system architectures that enable the realization of such integrated scheduling frameworks.
2. Provide a basis for adaptive improvement of production management decision-making. The generation of schedules incorporates predictive assumptions of the performance of the factory, including operation durations, operation failure rates, machine reliability constraints, accuracy, etc. Over time, the system should be able to adapt its model to better reflect the reality of the factory floor. Similarly, the system should be able to adapt and extend its scheduling strategies on the basis of past problem solving experience. We are investigating

frameworks for facilitating and automating these learning processes.

3. Provide a framework for decentralization the overall production management decision-making process. Coordination of production requires planning and reaction at different levels and over different time horizons. Decisions made at higher levels (e.g. decisions regarding manpower requirements and shifts of operation in different areas of the factory) provide constraints on the more detailed decisions that must be made at lower levels (e.g. decisions regarding the short term schedule for a particular area in the factory). Similarly, the results of factory operation necessitate reactive actions that may involve decision-making at several different levels. Given both the complexity of this overall process and the concurrency of manufacturing activities, decentralization of production management responsibility is required to achieve good system responsiveness. We are investigating mechanisms for coordinating the activity within a collection of independent schedulers, organized according to various structural and functional decompositions of the factory.

4. Develop a scheduling system organization that accommodates any division of decision-making responsibility between the user(s) and the system. We recognize that in many manufacturing domains it may not be feasible nor desirable to fully automate all aspects of the coordination process. Furthermore, ultimate transfer of technology resulting from our work into operational environments requires a framework that enables organizations to gracefully change current production management practices. We are investigating scheduling system organizations in which decision-making knowledge and methods are made transparent to the user, can be easily manipulated by the user to influence system decision-making, and can be utilized in decision-support mode to explore decision-making options.

## 6. Intelligent Measurement and Control Laboratory

The premise of our research is that factories designed for flexible automation in the manufacture of narrowly targeted products in small batches will win the battle against hard (fixed) automation only if they are set-up and fine-tuned *quickly*, *correctly*, and *economically*. In small batch operations, set-up and tuning of unit and multiple processes are daily or even hourly activities that in some industries account for more time than production itself. Automation of set-up and tuning -- a sort of second order automation, on top of the process run-time automation that is today's rule -- are becoming increasingly critical to survival.

The goal of the laboratory is threefold: first, to understand the acquisition and application of set-up, tuning, and control knowledge for both the unit process and multiple interacting processes; second, to understand how this practical knowledge contrasts with formal scientific and engineering knowledge; and third, to use this understanding to automate the acquisition and application of practical knowledge by the process control computer. Ironically, the task becomes increasingly feasible as the manufacturing process machinery becomes increasingly isolated from the human operator by the interposition of the computer between them. As the man becomes more and more a feedback element in the process control loop, interacting with the process only by watching its sensors on the computer screen and affecting its actuators only through the computer keyboard or joystick, we become more and more confident that the system, with its sensors and actuators, is complete and closed. Since the process can be made to work without the intervention of human sensors or manual adjustments, a sufficient set of sensors and actuators must be in place. Our task then becomes a matter of synchronizing, interpreting, and connecting the computer's internal (but until now unconnected) records of process state evolution and operator control action.

But in reality it is not necessarily this easy. There are underlying research issues, even given the complete record of machine state evolution and operator activity:

- is there really enough information on record, in principle and in practice, to set-up, tune, and control the process in configurations that replicate previous configurations?
- if so, is there also enough information on record, in principle and in practice, to set-up, tune, and control the process in new configurations?
- if so, then how do we describe the new configurations to the computer?
- and, if in either case there is not enough information, what is missing? how did the operator acquire it? and, how can the computer acquire it?

Our research is pragmatic: its context is industry and the measure of its success is the acceptance of our systems in industrial working environments. We build prototypes in our laboratory, then work hand-in-hand with industrial researchers, engineers, and designers to implement the results on the factory floor.

## 7. Manufacturing Logistics Laboratory

Consider the situation where a customer in a far off region needs a part. Questions arise such as where can the part be found? Is it in a distribution center relatively close? If so, what is the fastest way of transporting it to the customer? If the part is not in inventory, perhaps it should be manufactured? If so, where and when? Does manufacturing have the materials and the responsiveness? Can we communicate our needs to them? Can they change their schedule? If the plans for the part no longer exist, can it be made? Reverse engineered? In the global marketplace, providing sufficient products in the right place at the right time, i.e., parts on demand, is an expensive and difficult task. Optimizing

- the order entry process,
- where and how much to manufacture,
- distribution center locations and inventories,
- routes, transportation modes and policies, and
- how to administer an order including sourcing and routing

is beyond the current state of the art. Yet, with the enormous costs involved in manufacturing logistics, even a savings of 1% is enormous.

The first goal of the laboratory is to develop a decision support system for manufacturing logistics design and analysis. The system will enable the user to design a manufacturing and distribution logistical system, simulate and analyze it. The system will have embedded expertise which will guide the further optimization of the design so that products may be provided to the customer rapidly while at the same time reducing inventory and transportation costs.

Secondly, a realtime logistics control system will be built which provides for the distributed monitoring of production, administration of orders, and sourcing and routing of orders. Its architecture will support the logistical integration of design, planning, production, distribution and field service.

Lastly, the systems will be connected so that feedback from the logistics control system will be used to further analyze system performance so that alterations may be performed to provide further optimization.

In summary, the objective is to achieve the coordinated control of multi-plant complexes to meet the overall business goals of the enterprise. This includes the following sub-tasks : linking plant level and enterprise level decision systems; evaluation of the impact of coordination and uncertainty on Just-In-Time systems; and the choice of vendors and order decisions in multi-plant systems.

## 8. Manufacturing System Architecture Laboratory

Optimization of manufacturing can only be achieved by greater integration of activities throughout the product manufacturing life cycle. Integration must not only address the issues of shared information and communication, but how to coordinate decisions and activities throughout the firm. To achieve this level of integration it is necessary to develop a manufacturing systems architecture which provides a shared representation of manufacturing information and knowledge, and protocols which enable cooperative decision making by supporting queries and responses, feedback, information distribution, task distribution and negotiation.

The goals of this laboratory include:

***Distributed Architecture.*** Historically, manufacturing systems have tended to be islands of automation focusing on a small set of problems. While integration of these islands has been a goal, technology has only recently made this practical (i.e., networks of low cost workstations). The goal of the *Manufacturing System Architecture* project is to integrate decision making across the entire manufacturing product life cycle. The question is, what does it mean to integrate activities within production and across the manufacturing product life cycle. Even more interestingly, what does it mean to integrate all decision making within the organization? What is the appropriate architecture for a manufacturing system which integrates all of the activities throughout the manufacturing product life cycle? How can this architecture support the parallel performance of activities yet allow them to constrain each other's decision making. In addition, how can it also provide for feedback in any relevant direction?

In order to answer these questions, one must analyze the tasks performed in an organization. For example, the task of a product manager requires the ability to allocate and negotiate tasks and resources across many functional units (e.g., design, production, field service). On the other hand, design groups need to access warranty repair information from field service in order to focus their redesign efforts.

In these cases, the "organization information systems" must support many modes of interaction, including:

- intelligent networking where the system is able to route the information to individuals and groups who need it.
- feedback to appropriate portions of the organization.
- the definition and support of control/authority relationships.
- negotiation of changes to product and activities.
- task allocation and integration
- the constraining of each other. For example, design constrains fabrication and vice versa.
- allocation and reallocation of resources. For example, computers and personnel.
- learning from each other. For example, problems with products in the field may result in changes in the design.

In other words, application level protocols must exist to support these interactions.

***Manufacturing Modeling.*** A major goal is the establishment of representational standards for all types of information, including: product, resources, operations, organization structure, orders, personnel, etc. Without standardization, at least at the level of communication, integration cannot occur. The architecture will provide a clear and concise semantics for the representation of manufacturing knowledge, including:

- Activities: states, acts, time, causality, goals, milestones, ...
- Product: components, versions, revisions, ECOs., ...
- Organizations: departments, personnel, resources, functions, goals, ...

***Testbed.*** The project will also provide a testbed (simulator) in which all the other projects can "plug in" their modules in order to validate their operating behavior.

## 9. Bibliography


Akella, R., (1989), "Dispatch in Flexible Assembly, Test, and Manufacturing Systems," *Operations Research Models in Flexible Manufacturing Systems*, F. Archetti, M. Lucertini, and P. Serafini (Eds), CISM Courses and Lectures No. 306, International Center for Mechanical Sciences, Springer Verlag, 1989.

Akella, R., (1989), "Multi-plant Coordination under Material Flow Constraints," *Proceedings of the 1989 IFAC Workshop on Decisional Structures in Automated Manufacturing Systems*.

Baykan C., and M.S. Fox, (1987), "An Investigation of Opportunistic Constraint Satisfaction in Space Planning", *Proceedings of the International Joint Conference on Artificial Intelligence*, Milano, Italy.

Bourne D.A., and Wright, P.A., (1988), "Automation and Craftsmanship", *Computer Integrated Manufacturing Systems*, Vol. 1, No. 3, pp. 131-142.

Finger, S., Fox, M.S., Navinchandra, D., Prinz, F.B., and Rinderle, J., (1988) "Design Fusion: A Product Life-Cycle View for Engineering Designs," *Proceedings of the Second IFIP WG 5.2 Workshop on Intelligent CAD*, 19-22 September 1988, Cambridge, UK, edited by H. Yoshikawa and T. Holden.

Fox, M.S., Y.V. Reddy, N. Husain, M. Roberts, (1989), "Knowledge Based Simulation: An Artificial Intelligence Approach to System Modeling and Automating the Simulation Life Cycle", *Artificial Intelligence, Simulation and Modeling*, L.E. Widman (Ed.), New York: John Wiley & Sons.

Smith, S., Fox, M.S., and Ow, P.S., (1986), "Constructing and Maintaining Detailed Production Plans: Investigations into the Development of Knowledge-Based Factory Scheduling Systems", *AI Magazine*, pp. 45-61.

Smith, S.F., and Ow, P.S., (1985), "The Use of Multiple Problem Decompositions in Time Constrained Planning Tasks", *Proceedings of the International Joint Conference on Artificial Intelligence*, Los Altos, CA: William Kaufmann Inc. pp. 1013-1015.

Reddy Y.V., and M.S. Fox, (1983), "INET: A Knowledge-Based Simulation Approach to Distribution Analysis", *Proceedings of the IEEE Computer Society Trends and Applications*, National Bureau of Standards, Washington DC.