


AD-A245 714



2

Knowledge-Based Production Management: Approaches, Results and Prospects

Stephen F. Smith

CMU-RI-TR-91-21



Center for Integrated Manufacturing Decision Systems
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

December, 1991

This document has been approved
for public release and sale; its
distribution is unlimited.

Copyright © 1992 Carnegie Mellon University

This work was supported in part by a research grant from Intel Corporation and in part by the Robotics Institute.

DEFENSE TECHNICAL INFORMATION CENTER



9203347

92 4 00 099

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 1991		3. REPORT TYPE AND DATES COVERED technical	
4. TITLE AND SUBTITLE Knowledge-Based Production Management: Approaches, Results and Prospects				5. FUNDING NUMBERS	
6. AUTHOR(S) Stephen F. Smith					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Robotics Institute Carnegie Mellon University Pittsburgh, PA 15213				8. PERFORMING ORGANIZATION REPORT NUMBER CMU-RI-TR-91-21	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; Distribution unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In this paper we provide an overview of research in the field of knowledge-based production management. We begin by examining the important sources of decision-making difficulty in practical production management domains, discussing the requirements implied by each with respect to the development of effective production management tools, and identifying the general opportunities in this regard provided by AI-based technology. We then categorize work in the field along several different dimensions, indicating the principal types of manufacturing domains that have received attention, the particular production management and control activities that have been emphasized, and the various perspectives that have emerged with respect to the tradeoff that must be made in practical production management contexts between predictive decision-making to optimize behavior and reactive decision-making to manage executional uncertainty. The bulk of the paper focuses on summarizing the dominant approaches to knowledge-based production management that have emerged. Here, we identify the general concepts, principles, and techniques that distinguish various paradigms, characterize the strengths and weaknesses of each paradigm from the standpoint of different production management requirements, and indicate the results that work within each paradigm has produced to date. Among the paradigms for knowledge-based production management considered are rule-based scheduling, simulation-based scheduling, constraint-based scheduling, fuzzy scheduling, planning and scheduling, iterative scheduling, and interactive scheduling. We also examine work aimed at integrating heterogeneous planning and scheduling methods (both AI and OR based) and the construction of systems for multi-level production management and control. Finally, we survey more					
14. SUBJECT TERMS				15. NUMBER OF PAGES 37 pp	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT unlimited	18. SECURITY CLASSIFICATION OF THIS PAGE unlimited	19. SECURITY CLASSIFICATION OF ABSTRACT unlimited	20. LIMITATION OF ABSTRACT unlimited		

Table of Contents

1. Introduction	1
2. The Production Management Problem	1
2.1. Existing Solutions	3
2.2. Opportunities Provided by AI Technology	4
3. AI in Production Management: Dimensions and Scope	5
4. Approaches to Knowledge-Based Production Management	8
4.1. Rule-Based Approaches	8
4.2. Simulation-Based Scheduling	9
4.3. Constraint Satisfaction Frameworks	10
4.4. Constraint-Directed Search	12
4.5. Multi-Perspective Scheduling	13
4.6. Activity-Based Scheduling	14
4.7. Fuzzy Scheduling	16
4.8. Planning and Scheduling	17
4.9. Iterative Scheduling	18
4.10. Interactive and Mixed-Initiative Scheduling	20
4.11. Heterogeneous, Multi-Level Approaches	21
4.12. Distributed Production Management	24
4.13. Learning Shop Control Policies	26
5. Prospects and Challenges	26
References	29

Approved For	
NUS GRAD	
Date Recd	
By	
Distribution /	
Approved, Date	
Dist	Approved
A-1	

Abstract

Over the past decade, a large (and continually increasing) number of efforts (both research and development) have sought to investigate and exploit the use of Artificial Intelligence (AI) concepts and techniques in production management applications. In some cases, AI-based concepts have provided frameworks for making traditional Operations Research (OR) techniques more accessible and usable in practical production management settings. In others, novel concepts and techniques have been developed that offer new opportunities for more cost-effective factory performance. While this field of "knowledge-based" production management is still fairly young and the literature is still dominated by experimental research systems, results are nonetheless starting to have an impact in actual production environments. In recent years, several systems have made their way into operation, and many have been attributed with substantial manufacturing performance gains.

In this paper, we provide an overview of research in the field of knowledge-based production management. We begin by examining the important sources of decision-making difficulty in practical production management domains, discussing the requirements implied by each with respect to the development of effective production management tools, and identifying the general opportunities in this regard provided by AI-based technology. We then categorize work in the field along several different dimensions, indicating the principal types of manufacturing domains that have received attention, the particular production management and control activities that have been emphasized, and the various perspectives that have emerged with respect to the tradeoff that must be made in practical production management contexts between predictive decision-making to optimize behavior and reactive decision-making to manage executional uncertainty. The bulk of the paper focuses on summarizing the dominant approaches to knowledge-based production management that have emerged. Here, we identify the general concepts, principles, and techniques that distinguish various paradigms, characterize the strengths and weaknesses of each paradigm from the standpoint of different production management requirements, and indicate the results that work within each paradigm has produced to date. Among the paradigms for knowledge-based production management considered are rule-based scheduling, simulation-based scheduling, constraint-based scheduling, fuzzy scheduling, planning and scheduling, iterative scheduling, and interactive scheduling. We also examine work aimed at integrating heterogeneous planning and scheduling methods (both AI and OR based) and the construction of systems for multi-level production management and control. Finally, we survey more recent research in the areas of distributed production management and automated learning of factory floor control policies from experience. We conclude by discussing the current and future prospects of this work. In doing so, we also identify some of the important obstacles and challenges currently facing the field.

1. Introduction

Over the past decade, a large (and continually increasing) number of efforts (both research and development) have sought to investigate and exploit the use of Artificial Intelligence (AI) concepts and techniques in production management applications. This work has attempted to respond to (1) the inadequacies of existing computer-based solutions in this area and the consequent inefficiencies that plague industry today, and (2) the limited impact that results from the fields of Operations Research (OR) and Operations Management (OM) has had over the years in practical factory operations. In contrast to the fields of OR and OM, AI-based approaches to production management and control have emphasized the development of solutions that match the requirements, characteristics and constraints of practical production management problems. In some cases, AI-based concepts have provided frameworks for making traditional OR-based techniques more accessible and usable in practical production management settings. In others, novel concepts and techniques have been developed that offer new opportunities for more cost-effective factory performance. While this field of "knowledge-based" production management is still fairly young and the literature is still dominated by experimental research systems [66], results are nonetheless starting to have an impact in actual production environments. In recent years, several systems have made their way into operation. The ISA system, in use at DEC since 1985, provides support in establishing customer delivery dates and has reduced individual order processing time by an order of magnitude while increasing order processing consistency. The DISPATCHER system [1], developed jointly by DEC and Carnegie Group, has monitored and controlled automated material handling systems at two DEC facilities since 1986, with an estimate increase in production rates of 100%. LMS [149, 43], a real-time, distributed logistics management system developed by IBM, has been operational in its semiconductor manufacturing facilities at Essex Junction, Vermont since 1989, now has over 400 users and has contributed to an overall 10-20% increase in throughput in the first two buildings installed. Seimen's REDS system [50], a multi-level production planning and control system has been operational in a VLSI fabrication plant for over a year now with a reported increase in productivity of at least 10%. [52] report that their JOBCODE scheduler is now operational in 20 different sites. Several other efforts aimed both at installation/testing of field systems and commercialization of research systems are currently underway (e.g., [64, 5, 63, 140, 152]).

In this paper, we summarize work in knowledge-based production and operations management and assess the current state of research and practice. We first consider the important sources of difficulty in practical production management problems, discuss the requirements implied by each with respect to the development of effective production management tools, and identify the general opportunities in this regard provided by AI-based technology. We then characterize the scope of work in knowledge-based production management along several dimensions, including the types of application environments emphasized, the types of production management and control activities emphasized, and the variety of planning and control perspectives that have oriented work in the field. The major portion of the paper is devoted to a review of various approaches to knowledge-based production management and control, in which we identify the important concepts and techniques that underlie various efforts, and examine how this work addresses the requirements previously given. Our intent here is not to exhaustively consider all work and systems that have been reported in the literature (although we include a fairly comprehensive bibliography), but rather to identify, highlight and assess essential aspects of dominant alternative paradigms that have emerged. Other references to work in knowledge-based production management may be found in [147, 65]. We conclude with a discussion of the field's current and future prospects, and the important issues and directions for future research.

2. The Production Management Problem

To provide a context for discussing work in AI and production management, let us first look briefly at the dimensions of and requirements for practical production management, and the current state of affairs of manufacturing practice. Production management is functionally concerned with planning and control of the manufacturing enterprise, the goal being to effect a factory behavior where demands for

manufactured products are produced in a timely and cost effective manner. Several factors contribute to the difficulty of the production management problem:

- *complexity* - The optimal allocation of resources to activities over time is known to be combinatorially complex under idealized problem formulations [48]. In actual domains, the management problem typically involves the synchronization of large numbers of activities and resources subject to a much more complex (and often idiosyncratic) set of time and resource utilization constraints and objectives.
- *constraints* - Production management is invariably driven by a diverse set of constraints originating from many different functional units of the organization. Corporate policy and upper management impose specific organizational objectives on the manufacturing process (e.g., productivity, quality). Demands and demand forecasts provided by marketing must be balanced against manufacturing capabilities and the constraints imposed by purchasing on the acquisition of required materials. Engineering dictates constraints on acceptable and preferred production processes (as well as manufacturing demands relating to ongoing product development). Physical constraints on resource capabilities and utilization requirements, as well as operating preferences, originate from Manufacturing. Status information flowing from the factory floor provides short term information about resource availability and manufacturing progress, as well as longer term trends (e.g., recurring manufacturing problems) that affect production planning decisions. Although a subset of the imposed constraints are non-relaxable, production management is fundamentally a activity that involves negotiation and compromise.
- *uncertainty* - Another fact of life in most manufacturing environments is uncertainty. Decisions must often be made with incomplete and uncertain knowledge of future production circumstances (e.g., raw materials ordered on the basis of forecasted demand), and adjusted as expectations concerning the external marketplace later become invalidated. Within the execution environment, unexpected resource unavailability and unexpected delays and outcomes in the execution of activities continually work against attempts to follow prescriptive plans.
- *access to information* - Responsiveness to changing circumstances requires timely access to relevant information. At the lowest level, this requires data collection capabilities on the factory floor. However, without an information infra-structure for monitoring and analysis of the real time data stream and communication of relevant information to appropriate decision-makers, much of the value of collecting this data is lost.
- *multiple decision-makers* - As alluded to above, production planning and control decisions typically involve multiple decision-makers. Both overall problem complexity and the need for responsiveness to changing circumstances pragmatically necessitate a distribution of decision-making responsibility, raising problems of coordination between decision-makers operating at different levels (e.g., strategic, operational) and from different functional perspectives (e.g., production, marketing, sales). Effective manufacturing performance also requires synergy with the decisions and actions of other organizations (e.g., suppliers and consumers), which may or may not share the same goals and constraints.

Given the different time scales over which production management decisions must be made, there are several levels at which planning must occur (although the emphasis given to different levels will vary across manufacturing disciplines). Longer term, strategic planning/scheduling is required to establish shorter term production requirements that meet global management objectives, effectively calibrate customer delivery dates and provide guidance for raw material acquisition. Shorter term, tactical planning/scheduling is required to organize work flow in the factory in accordance with short term production requirements (e.g., job mix and priorities), managerial objectives (minimize cycle time, maximize resource utilization) and current manufacturing capabilities (e.g., resource, raw material availability). At the lowest time scale, planning/scheduling involves determination of what to do next at individual work centers in the factory. Both overall problem complexity and the inherent uncertainty in decision-making over different time scales dictate the use of coarser models of time and resource

utilization constraints at successively higher levels of planning. On one hand it is a practical necessity to make planning tractable; on the other it provides a hedge against current uncertainty about the future course of events. At the same time, evolving circumstances will inevitably require reactive action and adaptation of plans/schedules over time at all planning levels.

Production planning and control regimes (e.g., problem decomposition assumptions, levels of planning, solution techniques and heuristics) are strongly tied to the characteristics (and particularly the dynamics) of the target production environment. For example, in simple, well-structured environments where demand patterns are stable and appropriate vendor relationships exist (e.g., repetitive manufacturing environments), there are often less stringent requirements placed on global planning and effective performance can be achieved via simple, local coordination patterns (e.g., KANBAN). Alternatively, in small batch, job-shop manufacturing environments good performance may require stronger predictive guidance. Generally speaking, the types and extent of uncertainties present in the environment (e.g., frequency of machine breakdowns, yield characteristics, demand fluctuation) will dictate the level of detail at which planning/scheduling can feasibly be performed. What we should expect, therefore, is that different classes of production management problems will require different types of solutions.

2.1. Existing Solutions

Most existing production scheduling systems do not adequately address the solution requirements implied by these characteristics of the production management problem. The complexity of the problem and its constraints demand the advance development of schedules. It is only through anticipation of potential constraint interactions (e.g., resource contention) that the harmful effects of these conflicts can be minimized and performance can be optimized. However, the utility of this activity depends on the "match" between the schedule and the production environment, i.e., the extent to which the schedule reflects the actual constraints and objectives of the target domain. Traditional production scheduling research (e.g., [49]) has concentrated on methods for obtaining optimal solutions to idealized problems that have little to do with actual production environments. Similarly, most existing scheduling systems adopt oversimplified models of the domain and its constraints. MRP systems, for example, typically operate with standard lead time and infinite capacity assumptions, and consequently produce schedules that often bear little relationship to the actual problem that must be solved. As a result, the schedules must be constantly corrected to account for actual production constraints. Better solutions require an ability to represent and schedule with respect to the actual constraints and objectives of the manufacturing environment.

At the same time, production management is not a static optimization problem, but an ongoing activity in an unpredictable environment. An ability to generate schedules that accurately reflect the actual constraints and objectives of the target domain is of limited use without a companion ability to responsively adapt the current schedule to the specifics of the current execution state over time. In fact, in many production scheduling domains, the management problem is more accurately described as a rescheduling problem rather than a scheduling problem. In this regard, as has been recently pointed out by several researchers (e.g., [85]), work in scheduling theory has been solving the wrong problem altogether. For the most part, OR research relevant to managing executional uncertainty has focused on the development of local dispatch heuristics for dynamic decision-making (e.g., [111]). These approaches offer decision-making robustness, but make only marginal use, if any, of the results of advance scheduling (e.g., release and due dates) and provide no direct basis for reacting to changing circumstances in a globally coherent manner. Most existing scheduling systems provide no support for reactive management of schedules over time. Moreover, support for simply exploring the consequences of manual schedule changes is often beyond the capabilities of current systems. MRP systems, as well as existing finite capacity schedulers, are rigid software programs that sometimes take days to run. The result is ad hoc "fire fighting" by human decision-makers when unexpected situations demand change and rapid deterioration of the overall coherence of organizational activity.

Finally, effective solutions to the production management problem should recognize that decision making is typically a collective activity involving decision-makers with varying perspectives, expertise and levels of authority. This implies the need for an organizational structure that facilitates decentralized management and control, and mechanisms for communication and collaborative decision-making within this organizational structure. In most existing manufacturing environments, there is inadequate computer support for timely access to and communication of relevant decision-making information, and for direct interaction and negotiation among distinct decision-making agents.

2.2. Opportunities Provided by AI Technology

In relation to the requirements of practical production management problems identified above, AI-based technology offers several broad opportunities for improved effectiveness:

- Representation of production management constraints - Declarative knowledge representation schemes that accommodate both symbolic and numeric constraints (stemming principally from AI research in automated planning) provide a basis for constructing realistic models of the production environment and its operating constraints (in contrast to the restrictive modeling assumptions made in the mathematical models of OR and in the conventional production management systems used in industry today). Such models, by virtue of their declarative nature, are accessible, interpretable and extensible. Frame-based and object-oriented representation systems (originating from AI research in knowledge representation) offer structured modeling frameworks that enhance these inherent model properties and enable economical encoding of large declarative models.
- Planning/scheduling under diverse constraints - With respect to incorporation of constraints into the scheduling process, decomposable "heuristic search" procedures provide modular, extensible frameworks for introducing heuristic knowledge to focus the scheduling process and manage the complexity of the problem. Similarly, "constraint satisfaction" frameworks combine deductive constraint propagation techniques (which serve to limit the search) with the flexibility to inject appropriate heuristic search knowledge. The emphasis on decomposable, "open" search techniques within AI is a significant distinction that can be drawn relative to the nondecomposable, "closed" procedures typically emphasized in OR.
- Representation and incorporation of production management expertise - Rule-based representations and reasoning provide a straightforward basis for encoding and applying managerial heuristics, either in the context of heuristic search (the typical case in AI-based scheduling research) or as a direct means for problem solution.
- Management of uncertainty - Techniques for representing and reasoning with uncertainty coupled with temporal constraint management techniques provide a basis for generating production schedules that (1) pro-actively take into account current uncertainties and (2) avoid overcommitment whenever possible. Techniques for constraint management coupled with decomposable decision (search) procedures enable isolation of problems in the schedule resulting from unexpected events and incremental schedule revision.
- Integration of heuristic and analytical procedures - In many respects, the heuristic problem solving frameworks that have emerged from the field of AI can be seen as complementary to the analytic techniques produced by OR. These frameworks can provide a basis for exploiting knowledge of model assumptions, parameters, setup, and applicability to (1) make existing OR techniques more accessible and usable to an end user, or (2) opportunistically exploit a collection of analytic/heuristic procedures as appropriate during the planning/scheduling process.
- Flexible decision support and schedule manipulation - Interactive, two-way graphical interfaces, coupled with constraint management techniques and relevant production management knowledge can provide information access, schedule editing, constraint checking, explanation and advisory capabilities that greatly enhance a human decision maker's productivity. Knowledge-based techniques also provide a basis for explicitly

modeling the user and providing customized system support commensurate with user goals, expertise, and preferences. Similar opportunities exist for the construction of intelligent "front-ends" to existing production management tools.

- **Monitoring, and analysis of status information** - Knowledge-based monitoring and diagnostic techniques are quite mature and can be applied to problem of detecting exceptional situations (e.g., consistent production bottlenecks) in the stream of status information received over time. Knowledge of the implications of such situations can be exploited to either recommend or initiate corrective actions.
- **Adaptable production management decision-making over time** - Much of the research in AI-based scheduling has focused on techniques that are driven by characteristics of the current problem state. Such techniques can be seen as adapting their scheduling strategies to fit current circumstances (e.g., the emergence of a new bottleneck resource). At another level, research in the field of machine learning offers the potential of improving system scheduling strategies over time on the basis of prior experience (e.g., refining through experience the applicability of different dispatch heuristics in various production circumstances). Some preliminary work in this area has been reported.
- **Support for coordination of independent decision-makers** - Although also relatively recent, research in the area of distributed problem solving (generally) and distributed production management (specifically) offers the promise of improved coordination among different participants in the production planning and control activity.

In the following sections we consider the nature of these opportunities in more detail in the context of recent work in AI and production management. We first summarize the scope of work in the field along several different dimensions. This is followed by a more detailed examination of the techniques that have been proposed and applied, and an assessment of their strengths and weaknesses.

3. AI in Production Management: Dimensions and Scope

We can categorize the work in AI and Production Management along several dimensions. One important dimension concerns the manufacturing discipline of focus. As indicated above, the dominating characteristics of a given manufacturing environment have a significant influence the utility of various approaches and heuristics. AI-based production management research has addressed a variety of application domains:

- **Job shop scheduling** - As has been noted in previous surveys [147, 65], job shop manufacturing disciplines have historically received the most attention and much work continues in this area (e.g., [5, 9, 18, 29, 33, 37, 44, 92, 52, 53, 58, 63, 68, 71, 76, 82, 84], [86, 106, 125, 140, 157]).
- **Flow shop environments** - Work has also considered management and control of a variety of flow shop environments, ranging from steel making (e.g., [21, 100], to beer production [140, 158] to windshield manufacturing [54] to computer board assembly and test [1, 28, 57, 140]) and other high volume manufacturing environments [26].
- **Wafer fabrication** - A problem domain that has gained increasing interest in recent years is wafer fabrication, which exhibits problem characteristics common to both job shop and flow shop environments [27, 40, 50, 149, 141, 159].
- **Flexible manufacturing systems** - Another substantial focus of research has been intelligent control of flexible manufacturing systems [8, 19, 23, 42, 72, 62, 73, 74, 101, 113, 120, 133], [134].

Orthogonally to manufacturing discipline, we can categorize research according to the level (or levels) at which the production management and control problem has been addressed:

- **higher-level, strategic planning** - Work in AI-based production management in the area of longer term, strategic planning has been fairly sparse (although admittedly the boundary

between strategic planning and shorter term tactical planning is somewhat ill-defined). While there have been isolated successes relative to providing better support for specific higher-level production management activities (e.g., order processing based on aggregate but finite capacity constraints and domain-specific allocation heuristics [102]), it appears that the most promising contributions to be made by AI at higher levels production planning might lie in the development of knowledge-based frameworks that provide linkages between strategic and operational decision-making and facilitate various production management activities within the organization (e.g., [153]). In point of fact, the quantitative models of OR provide the most leverage at aggregate longer range levels of planning (where detailed operational constraints are less influential), and it is not clear in many cases that AI-based techniques can offer better analysis and planning methods. At the same time, AI can augment such methods by providing supporting frameworks for determining which methods to apply when (and how) to better coordinate overall production management activities.

- shop/factory level scheduling - By far, the most dominant focus for research in AI-based production management has been short-term, detailed scheduling of factory production (e.g., [4, 19, 21, 22, 92, 40, 52, 58, 63, 64, 71, 76, 79, 80, 106, 121, 125, 140, 156]), either seeking to fill the gap between coarser level master scheduling (or MRP planning) and shop floor control or, in other cases, advocating finite-capacity scheduling alternatives to MRP systems. In many respects, this focus is natural. First, the lack of support for short term, detailed scheduling and control is perhaps the most identifiable problem in current production management practice and a primary source of production inefficiencies. Second, this is the aspect of production management for which AI has the most to offer in terms of more effective solution techniques. This is where the limiting modeling assumptions of conventional and OR-based techniques become intolerable and the ability to incorporate the full range of constraints on production becomes imperative.
- work cell/work center control - Another substantial body of work has focused on more localized problems of work cell or work center scheduling and control [27, 32, 149, 54, 75, 77, 83, 101, 141, 148] and material movement [1, 99, 114, 155]. The motivations here are the same as indicated above at the detailed factory scheduling level: work has emphasized the power of incorporating more knowledge of the actual execution state and constraints in the management of actual factory operations.
- multi-level production planning and control - A smaller (but growing) number of efforts have focused on multi-level frameworks for production planning and control (e.g., [21, 94, 50, 61, 66, 96, 115, 145, 153]). For the most part, this work has concentrated on integrating operational levels of the overall production management function (in contrast to the longer range strategic levels of planning mentioned above). Some frameworks have assumed the use of common techniques at different levels, relying on hierarchical models of domain constraints as a basis for determining scope, temporal horizon, and precision of decision-making at different levels (e.g., [61, 115, 145]). Others have ascribed specific functional roles to different levels (e.g., determination of production quantities, tactical predictive scheduling, real-time reactive schedule repair) and have emphasized the integration of disparate methods (e.g., [50, 66, 96]).

A third dimension along which work in the field can be distinguished is in the position taken relative to the complexity/uncertainty tradeoff:

- pro-active planning orientation - Due to both the influence of classical approaches to production scheduling problems and the influence of AI research in generative planning, a large body of research in the field has advocated a strong predictive view of the scheduling problem (e.g., [22, 92, 50, 76, 80, 121, 125, 135, 140, 159]). This work starts from the assumption that better factory performance requires better solutions to the complexity problem (i.e., generation of schedules that more accurately reflect actual production constraints and objectives), and has emphasized the development of techniques capable of providing better predictive guidance. Under such approaches, the problem of executional uncertainty is commonly seen as one of effectively managing the current predictive schedule

when mismatches are recognized relative to the current execution state (e.g., [18, 26, 29, 119, 138, 155]). It is assumed that such unexpected circumstances lead to either incremental schedule revision or complete schedule regeneration relative to the constraints known to be true about the current factory state, depending on the flexibility/decomposability of the scheduling techniques employed.

To address the potential computational burden of managing schedules over time, work in predictive scheduling has also investigated techniques for pro-actively taking uncertainty into account during schedule generation. Hierarchical models have been used in conjunction with uncertainty related criteria (e.g., time to execution) to control the level of precision at which various scheduling decisions are made over time (e.g., [50, 138]). Similarly, probabilistic and fuzzy set frameworks for managing time and capacity constraints (e.g., [37, 71, 97]) have been proposed to support least commitment decision-making and allow generation of schedules that contain appropriate degrees of freedom with respect to execution (e.g., remaining imprecise with respect to timing details, resource assignments).

- **reactive control orientation** - Sensitivity to the uncertainty inherent in operational control of many actual factory environments has motivated a second body of research to de-emphasize detailed predictive scheduling and focus instead on dynamic management of production activities at execution time (e.g., [1, 25, 14, 39, 83, 134, 149]). This work is dominated by more of a reactive control perspective of the scheduling problem and has its roots instead in classical research in dispatch heuristics and adaptive control. It generally assumes a much looser connection between activities of planning and execution, where whatever guidance that can be reasonably provided by advance planning/scheduling (e.g., job release and due dates, job priorities) is exploited in an "open loop" fashion in conjunction with characteristics of the current execution state. Predictive guidance provides an additional bias in making control decisions and interaction with the predictive scheduler is limited to periodic intervals (e.g., at the beginning of each shift, daily).

The downside of such a reactive orientation to production management is the tendency for myopic decision-making, due to inadequate treatment of the complexity aspect of the problem. To combat this tendency, a variety of approaches have been investigated. In some approaches (e.g., [1, 134, 149]), advance analysis of the behavior of the operating environment (often based on past experience) is used to determine a set of state-dependent control heuristics and operating policies that cover a wide range of manufacturing circumstances. Other work has sought to inject stronger predictive guidance into operational decision-making, either through postprocessing of rigid schedules in light of known uncertainties (e.g., [25]) or association of additional preference information that serve to highlight critical execution constraints and provide a looser basis for event-based interaction with the predictive scheduler (e.g., [83]).

- **interactive support** - A third perspective on the complexity/uncertainty tradeoff that complements the above two is inspired by the practical reality that no system can be expected to contain the knowledge necessary to reasonably account for all decision-making circumstances (either predictively or reactively), and that the human scheduler must be an integral component of factory floor planning and control [86]. This perspective has led to the development of a variety of frameworks for support of interactive scheduling by a human decision-maker (e.g., [13, 59, 86, 88, 100, 63, 129]). This work has emphasized establishment of productive user and system roles, access to and use of system knowledge of production management constraints, and effective user integration with system scheduling capabilities.

Speaking more generally, the appropriateness of different production management orientations (pro-active, reactive, interactive) depends ultimately on the nature of the complexities and uncertainties in the target production environment. In high volume environments with random yields and thus considerable variance in product cycle times, there may be little useful operational guidance to be gained from detailed predictive scheduling. Alternatively, in small batch manufacturing contexts, specific production units traveling through the factory are typically associated with specific customer orders and it is important to

operate with much more detailed predictive guidance. What has been evident from the applications addressed by research rooted in any one perspective is that typically a mix of predictive, reactive and interactive scheduling capabilities is required for effective production management. Work in each camp has tended to push in the direction of the other. [50] provides a good example of an operational system that integrates pro-active and reactive perspectives based on the specific uncertainties of the target environment.

A final dimension of comparison, obviously influenced by each of the dimensions identified above, are the problem solving techniques that have been developed for addressing the production planning and control problem. A diversity of approaches have been explored, and it is along this dimension that we choose to organize our more detailed summary of work in AI and Production Management.

4. Approaches to Knowledge-Based Production Management

Given the complexity of production planning and control problems, heuristic knowledge must play an integral role in any practical solution technique. This leads to two basic (and related) issues: one concerning how to represent the problem to be solved and the knowledge required to solve it, and the other concerning how to apply this knowledge to effectively search the problem space for a solution. The various approaches discussed below can all be seen in relation to these two basic types of choices.

Our overview of approaches to knowledge-based production management is structured as follows. We first summarize alternative heuristic problem solving paradigms have been investigated as a basis for operational production management and control, considering in turn rule-based scheduling, simulation-based scheduling, constraint-based scheduling, fuzzy scheduling, planning and scheduling, iterative improvement scheduling, and interactive scheduling. We then examine work aimed at integrating heterogeneous planning and scheduling methods (both AI and OR based) and the construction of systems for multi-level production management and control. Finally, we survey more recent research in the areas of distributed production management and automated acquisition of production management decision-making expertise from experience.

4.1. Rule-Based Approaches

One basic approach to production scheduling and control problems that has been pursued is the so called "expert systems" or rule-based approach [9, 1, 19, 20, 27, 70, 72, 74, 82, 102, 149], [156], which produced the initial real-world successes of AI in other application domains. The objective under this approach classically is to mimic the decision-making of a human expert. Experiential knowledge consisting primarily of empirical associations between decision-making conditions and problem solving actions are extracted and encoded as a set of IF THEN rules. For example [19],

IF a lot is available and could be scheduled but it is known that an
essential machine must undergo maintenance in a short time
THEN delay this lot and try to schedule another

Within production planning and scheduling applications, such rules are typically applied in a forward-chaining (constructive) manner, moving forward from an initial state (e.g., an empty schedule or an initial schedule with problems) to a satisfactory result. In terms of search, consideration of alternative schedules is typically severely limited (if not eliminated). But this need not be the case. In some cases, human expertise has been profitably augmented by injection of heuristic search procedures at appropriate points (e.g., [1]).

There are several arguments as to why a pure expert systems approach may generally not be an effective approach to production scheduling problems. First, it is not clear that a human scheduler does a particularly good job at his task. As has been observed [92], schedulers are often more accurately characterized as jugglers of constraints and are overwhelmed by problem complexity. Scheduler's rules

are often myopic in nature, aimed at solving small subproblems or putting out immediate fires, and it is difficult to extract useful global scheduling strategies. Finally, as stated in [5], expert systems approaches have been most successful in problems where the number of possibilities at each step is fairly small (e.g., diagnosis problems); this is not the case in production scheduling and rule-based approaches can result in fairly arbitrary pruning of scheduling possibilities.

At the same time, there do appear to be production management circumstances where rule-based approaches, either employed in conjunction with other techniques or applied under different base assumptions, can produce significant leverage. Much of the work in rule-based scheduling has focused on lower level, relatively local production control problems (e.g., FMS control and shop floor dispatching). Rule-based approaches, in fact, provide a natural and efficient framework for dynamic, state-dependent decision-making and decision-support. Furthermore, decision-making heuristics need not be limited to those of a single human operator but can be expanded to encompass the collective wisdom of the organization [149] and/or the results of problem analysis [82]. Simply combining timely status reporting and information access capabilities with coordination and decision-making rules can have a significant impact on overall factory performance, which can lead in turn to shorter higher-level production planning cycles. Moreover, in highly unpredictable environments, such approaches might offer the most practical solution to short term production control.

LMS [149] provides an example of the benefits of such broader perspectives regarding rule-based knowledge. LMS is a distributed logistics management system developed for real-time control of operations in IBM's wafer fabrication facilities in Essex Junction, Vermont. LMS can be seen at several levels (which parallel the actual development phases of the system). At the lowest level, LMS provides a transaction-based framework that integrates real-time data collected by various control systems. Superimposed on this framework are analysis methods which turn this data into useful information and presentation systems that make this information (and relevant part data) available to operators at individual work stations on the floor. At the next level, rule-based heuristics for pro-active intervention are introduced, which monitor the incoming data stream and communicate status alerts (e.g., machine down) to appropriate decision-makers (e.g., upstream work stations, maintenance personnel). At the highest level, rule-based heuristics that combine state information, part information, and organizational objectives can be utilized, either automatically or in decision-support mode, to make actual control decisions. In the cases of both alerting relevant decision makers and dispatch decision-making, rules were developed not on the basis of a particular expert, but instead based on analysis by a team of people with different organizational and decision-making specialties. As is stated by the system developers [149], "there is no single expert who can do what LMS does". As of 1989, LMS was operational in two buildings at the Essex Junction facility, and had over 400 users. Although the overall project also involved an analysis of fab bottlenecks and consequent changes other than the introduction of LMS, it has nonetheless contributed substantially to the observed 10-15% increase in throughput in both applications.

RBD [27] is a centralized, rule-based dispatching shell based on similar assumptions (e.g., integration with a factory control and status report system, capabilities for incorporating state-dependent rules, provision for work station dependent control heuristics), which has also been applied to control production flow in wafer fabrication facilities.

4.2. Simulation-Based Scheduling

A second basic approach to production scheduling that has been explored is scheduling via discrete-event simulation (e.g., [19, 58, 62], the long range scheduling component of [159]). Under this approach, a schedule is created by simply simulating the execution of the factory with an appropriate dispatch heuristic (or set of heuristics) and taking the recorded execution history as the schedule. Factory entities (e.g., resources, jobs, operations) are typically modeled as objects, and object-oriented programming and knowledge representation techniques provide a flexible basis for enforcing many complex scheduling constraints (e.g., conditional routings, material transport). In [19], extensions are described to basic

discrete-event simulation frameworks to enable treatment of activity durations that vary depending on the specific product mix that is allocated to given resources.

Despite this ability to straightforwardly accommodate complex allocation constraints, there are identifiable disadvantages to simulation-based scheduling approaches:

- The quality of the resulting schedule is questionable, due to the myopic nature of the decision-making. The overall problem is decomposed principally in an event based fashion (i.e., decisions are made in chronological order), and secondarily into a set of local resource scheduling problems. Search is thus confined to the operations that are eligible to acquire a free resource at a given point in (simulated) time. A decision made at any point may unwittingly restrict alternatives for critical future decisions which eventually leads to otherwise avoidable problems (e.g., unnecessary downstream congestion). [19] employs a closed queuing network algorithm as a "load evaluator" for look-ahead purposes, but proposes a broader search capability as a necessary extension.
- Simulation-based approaches also provide no reactive capabilities, except to rerun the simulation from the point of the problem onward. One consequence here is little control over the degree of continuity in the schedules produced over time. There is also difficulty in directly accommodating desired user changes to the schedule. In [58], an iterative scheduling framework involving a simulate and adjust cycle is described, where the human scheduler can "adjust" the results by imposing new constraints (e.g., new earliest start times to avoid blockage of rush jobs) or move jobs between two broad priority classes. In the latter case, a 2-pass simulation is performed, first considering only the high priority jobs and then restarting the the simulation and adding the lower priority jobs. Thus, the user must anticipate how specific adjustments will affect the simulator's output. This system is in operational use in a job shop. However, the effects on factory performance are not clear and it appears that the benefits may have more to do with prior lack of any decision support capability.

Generally speaking, it would seem that simulation is more suited for analysis of factory behavior (e.g., as a basis for determining appropriate shop floor operating policies) than as a basis for producing actual detailed constraints on execution. However, depending on the match of the simulator to actual floor control policies, it could also prove to be a viable basis for establishment predictable production requirements in highly uncertain environments. In either role, the modeling flexibility provided by knowledge-based approaches to simulation provides clear advantages over traditional simulation frameworks.

4.3. Constraint Satisfaction Frameworks

Observed inadequacies in expert scheduling heuristics and the inability of simulation-based techniques to produce globally satisfactory solutions have led researchers to seek other scheduling frameworks that (1) provide a less arbitrary basis for restricting search and (2) allow search to be more profitably integrated into the scheduling process. In this regard, the perspective of scheduling as fundamentally a constraint-driven process [92] has become a dominant organizing principle in the field. This perspective is not uncommon to work in constrained optimization within the field of Operations Research, and like this work, research in AI-based approaches has sought techniques for exploiting constraints to decompose the overall problem, and guide the search for a solution. However, there are several essential differences between these approaches and OR-based counterparts:

- approaches have been based on much richer representational frameworks, enabling incorporation of the full range of constraints that influence production (in contrast to mathematical models).
- work has emphasized use of heuristic knowledge about constraints as a basis for guiding the search for a solution and frameworks for flexibly integrating heuristics into the search process (in contrast to analytical, optimization techniques). Aside from providing tractable solution techniques to practical problems, this emphasis also better reflects the nature of actual

production planning and scheduling, where the notion of "optimal schedule" is ill-defined and there are typically a range of conflicting objectives and preferences that must be balanced.

- the techniques developed, by and large, have been incremental and decomposable in nature (in contrast to a "black box" solution technique perspective). This provides a much more flexible basis for both integration of users into the decision-making process and reactive management of schedules in response to unexpected execution circumstances.

A broad commonality observed in many approaches to constraint-based scheduling is a basic problem solving organization that distinguishes two components: a decision-making component and a constraint management component (e.g., [5, 6, 29, 40, 68, 88, 119, 125, 138]). Under this organization, the decision-making component is responsible for making choices among alternative scheduling decisions and conducting the search for an acceptable schedule. The decision-making component is typically heuristic in nature owing to the complexity of the underlying search space. The role of the constraint management subsystem, alternatively, is to incrementally maintain a representation of the current set of feasible solutions as scheduling choices are made by the heuristic decision-making component. This is accomplished by deductive constraint propagation techniques, which combine constraints specified in the underlying factory model (e.g., activity durations, precedence relations between activities, capacity and availability constraints on required resources) with external requirements (e.g., earliest release dates and due dates) and any new constraints introduced by the decision-making component and determine new constraints on the decisions that remain to be made. Inconsistent solutions are detected whenever there remain no feasible alternatives for one or more decisions that remain to be made. In such situations it is necessary for the decision-maker to either backtrack over one or more previous decisions or relax one or more specified problem constraints (e.g., loosen due dates, increase the number of shifts worked).

A variety of approaches to constraint management in support of scheduling have been defined. Some approaches have emphasized general-purpose frameworks for temporal reasoning [3, 36, 30, 38, 122]. These frameworks generally support least commitment styles of planning and scheduling, and many allow the generation of constraint satisfying solutions that retain degrees of freedom where constraints allow (e.g., schedules that dictate ordering constraints between activities but associate feasible execution intervals with activities as opposed to crisp start times). Other work has relied on more specialized propagation techniques, that are more closely aligned with the specific decision-making techniques employed, exploit characteristics of the problem at hand and/or compromise with respect to general modeling flexibility for efficiency of operation (e.g., requiring assignment of crisp activity start and end times, assuming fixed activity durations) [7, 21, 40, 137, 78]. Some approaches (e.g., [5, 21, 88]) additionally record "dependencies" between constraints as new constraints are inferred (to enable more selective retraction of constraints in situations of constraint conflict), record sets of decisions that have resulted in conflict for future avoidance (referred to as "shallow" or "deep" learning, depending on the level of analysis used to isolate the conflict set) and provide a basis for managing of alternative solutions (although this last capability typically comes at considerable computational expense). Others (e.g., [137, 78]) have defined techniques that combine management of both time and resource capacity constraints and support scheduling at different levels of abstraction.

One simple paradigm for scheduling based on the use of such constraint management techniques involves a decision-making component that simply conducts a depth first, backtracking search for a solution. Scheduling decisions to be made are treated as sets of variables requiring assignments (e.g., start time, resource), and the constraint management component is used to incrementally maintain the set of possible assignments for each variable as the search proceeds. At any point, a value is selected for a given unassigned variable from those that are currently feasible and the constraint management subsystem is invoked to determine the consequences (i.e., prune away possible values from the sets of other currently unassigned variables). Whenever it becomes no longer possible to make an assignment to a given variable, the search backtracks to a previous choice point and another assignment is selected.

Such basic constraint satisfaction search techniques, employing dependency-directed backtracking and shallow learning techniques, have been applied in several contexts. In [39, 40], such techniques are used to provide support for managing short-term production schedules in a semi-conductor manufacturing environment. A set of feasible schedules is incrementally maintained as execution proceeds. Guidance is provided to the user with respect to feasible scheduling alternatives in a given decision-making context, and in understanding the implications of various schedule changes. The DAS scheduler [21, 119] employs similar techniques as a basis for generating and reactively maintaining machine schedules in a steel manufacturing application. In this case, the problem is solved by a hierarchy of communicating schedulers, each responsible for a distinct portion of the overall problem. Schedulers at the lowest level maintain schedules for an individual machine, schedulers at the intermediate level manage loads on a specific group of substitutable machines, and a global scheduler coordinates the flow of jobs from machine group to machine group so as to respect job release and due dates.

Generally speaking, constraint propagation and management techniques provide a strong foundation for generating feasible schedules and reactively maintaining them as executional circumstances warrant. At the same time, basic constraint satisfaction search techniques, in and of themselves, are limiting in two respects:

- While they can be effective in locally contained or user-focused scheduling contexts, the combinatorics of undirected, systematic exploration of alternatives can rapidly become debilitating in larger, more complex scheduling domains.
- practical production scheduling problems are more often than not over-constrained (e.g., all due dates cannot be met), and the crux of the problem is determining a good compromise (or which constraints to relax) in the presence of conflicting objectives and preferences. Moreover, even within the space of feasible solutions, some solutions are better than others. A constraint satisfying solution that implies long queue times is less valued than one that doesn't. Pure constraint satisfaction techniques provide no leverage addressing either of these issues.

In the next several sections, a variety of heuristic approaches that extend this basic paradigm in different ways to address these problems are summarized.

4.4. Constraint-Directed Search

One of the earliest and most influential systems in the field of AI and production management was the ISIS job shop scheduling system [47, 89, 92, 93, 90, 143]. ISIS represented the first attempt to formulate and operationalize the view of scheduling as a heuristic, constraint-directed activity. It emphasized a broad view of constraints that recognized the conflicting and negotiable nature of many specified requirements and objectives, and focused on representation and use of knowledge about constraints to support effective conflict resolution and constraint relaxation. A representation of preference (i.e., relaxable) constraints was defined to model various factory objectives (e.g., meeting deadlines, minimizing work-in-process time) and operational preferences (e.g., reliability preferences among substitutable machines, sequencing preferences to minimize machine setups), and embedded within a larger relational model of the entities and physical constraints of the production environment. The representation encoded knowledge relating to possible relaxations of the preferred choice, the utility of each alternative, the decision-making relevance of the preference and the relative importance of satisfying different types of preferences. ISIS introduced several heuristic techniques for using constraints to guide the scheduling process:

- generative constraint relaxation - at the core of the ISIS approach was a beam search strategy for adding a new order (job) into the shop schedule. Within this search, physical constraints (i.e., operation precedence constraints, resource requirements and availability) were used as a basis for generating alternative sets of decisions. Relevant preference constraints provided the basis for evaluation and pruning of alternatives at each step of the search, thus implementing a generative approach to constraint relaxation.

- rule-based analysis - domain specific rules were applied both prior to the search to fix specific relaxable constraints (e.g., specify backward scheduling to ensure satisfaction of due date) and after the search to assess results and propose constraint relaxations if appropriate.
- staged consideration of constraints - The beam search was embedded in an iterative, hierarchical search framework where additional types of constraints were considered at each successive level. After selecting the next order to schedule on each iteration, a high level analysis of remaining resource availability was used to emphasize specific allocation intervals (through the introduction of additional preferences). After detailed forward or backward beam search scheduling, final local optimization to minimize WIP time was performed.

ISIS was applied to a Westinghouse Turbine Components Factory in Winston Salem, NC. The system was never actually fielded. At the same time, ISIS provided perhaps the primary trigger of the field of knowledge-based production management, and the system has heavily influenced much subsequent work in the field. A Westinghouse perspective of the ISIS project can be found in [112].

4.5. Multi-Perspective Scheduling

Experiences with the ISIS scheduler led to subsequent development of the OPIS scheduling system [108, 105, 109, 110, 107, 144, 138, 139, 140, 142]. OPIS incorporated and extended many of the concepts introduced in ISIS, but also differed in fundamental ways. From the standpoint of scheduling methodology, OPIS first introduced and demonstrated the power of reasoning from multiple local problem solving perspectives in producing a schedule that effectively compromises among conflicting objectives [146]. The concept is illustrated by consideration of two common scheduling objectives: minimizing work-in-process (WIP) time and maximizing resource utilization. A job-oriented problem decomposition provides an opportunity to minimize WIP time, since subproblems are solved that contain all constraints involved in optimizing this objective. At the same time, a job-oriented decomposition works against the objective of optimizing resource usage as the constraints relevant to this tradeoff are spread across subproblems. A resource-oriented decomposition strategy, alternatively, provides the complementary opportunity to optimize resource usage (at the expense of leverage in minimizing WIP). The subproblems solved in this case contain multiple requests for a given resource, enabling resource setups to be minimized. Whereas the ISIS system employed a strict job-oriented decomposition of the overall problem, OPIS advocated the cooperative use of both job-centered and resource-centered scheduling techniques.

Multi-perspective scheduling in OPIS introduced several novel techniques for constraint-based scheduling:

- opportunistic problem decomposition and subproblem formulation - Decisions as to what part of the overall schedule to concentrate on next and which local scheduling method to apply are made dynamically through repeated analysis of the characteristics of current solution constraints. Schedule generation was focused by analysis of projected resource contention, using the basic heuristic: "schedule the bottleneck resources first" [103, 104].
- constraint-based schedule repair - In situations of constraint conflicts, heuristics relating characteristics of conflicts and properties of existing time and capacity constraints (e.g., relative constrainedness, opportunities for nondisruptive repair, important optimization concerns) to the relative strengths and weaknesses of alternative repair methods, were used to opportunistically direct the schedule revision process. In addition to basic job-oriented and resource-oriented scheduling methods, more specialized repair methods (schedule shifting, demand swapping) were also incorporated.
- blackboard-based scheduling architectures - opportunistic, multi-perspective scheduling in OPIS was supported by a specialized "blackboard-style" architecture [41] emphasizing centralized constraint management, modular integration of analysis and scheduling methods and a rich underlying representational structure for modeling the manufacturing environment.

The architecture itself constitutes a flexible test bed or shell for the development and customization of schedulers in particular domains [118].

OPIS has been extensively tested in several domains, including the Westinghouse Turbine Components Plant mentioned above and an IBM Computer Board Assembly and Test Line at the Poughkeepsie facility. Comparative studies have quantified the utility of multi-perspective scheduling relative to both ISIS and other less flexible, single-perspective scheduling techniques in both generative and reactive contexts. OPIS remains an active research system and a version of the scheduler is currently being customized by Horizon Research Inc. for operational use in a beer manufacturing facility.

A system influenced by OPIS and similarly based on the opportunistic use of a variety of analysis and scheduling methods is the SONIA job shop scheduler [30, 29, 31]. In this case, the problem is to select and schedule a subset of jobs for the next shift/day from a pool of eligible candidates, and a different scheduling strategy is employed. Three basic methods are defined: one for job selection (descended from the earlier developed SOJA scheduler [76, 131]), one for job rejection and one for job-oriented scheduling. These three methods are opportunistically applied according to dynamic analysis of the current state of the schedule (e.g., resource contention levels, existing conflicts, decisions remaining to be made). A fourth shifting method is used exclusively in reactive scheduling contexts. SONIA's principal contributions, however, lie in the further development of system architectures for building constraint-based scheduling systems.

- flexible constraint propagation - a decomposable framework for temporal constraint management was defined, enabling explicit control over the extent and type of constraint propagation performed in different problem solving circumstances, and consequently more efficient scheduler performance.
- control architectures for constraint-based scheduling - A blackboard-based control framework was defined to provide a flexible, declarative structure for programming opportunistic scheduling strategies (including control over the behavior of the constraint propagation system).

A different approach to multi-perspective scheduling, referred to as reinforcement scheduling, is implemented in the RESS job shop scheduling system [80, 81]. Here, the basic notion (common to the OR community) of solving a simpler problem to gain insight for solving the actual problem is exploited. A "pre-planning" component is used to construct a rough utilization plan for the identified critical resource and provide a global view of future events. This rough plan is then used as guidance in constructing the actual schedule in conjunction with details of local constraints. Similar use of abstraction through omission of constraints can be seen in [148].

Despite the observed limitations of a fixed, single-perspective scheduling strategy, such techniques have been successful and made their way into operational use. The JOBCODE system [52] is reported operational in 20 different sites. JOBCODE schedules in a job-oriented fashion that appears remarkably similar, both terms of the search performed and the governing heuristics, to the approach implemented in ISIS. Much of the advantage of the JOBCODE scheduler could lie in the flexibility it provides for manipulating individual job schedules and its ability to deal effectively with more complex resource allocation constraints (e.g., job splitting, operation overlapping).

4.6. Activity-Based Scheduling

Motivated by the experimental success of the OPIS scheduler in combining dynamic analysis of expected resource contention and bottleneck-based focusing heuristics, as well as other recent work with bottleneck-based approaches to scheduling (e.g., [104, 87]), one line of more recent research has focused on the development of finer-grained approaches to opportunistic schedule generation [10, 68, 125]. This work is based on the observation that a coarse-grained approach to problem decomposition (e.g., one where it is assumed that bottleneck resources will be completely scheduled

before re-examining estimates of projected resource contention) only roughly approximates the underlying philosophy of "bottleneck-based" scheduling: that the scheduler should be consistently focused on the most critical decisions that remain to be made. It is often the case, for example, that resources are not highly contended for (i.e., bottlenecks) over the entire scheduling horizon, but only over some (possibly discontinuous) portion of the horizon. Furthermore, in many complex scheduling environments (e.g., job shops), contention for various resources is affected by allocation decisions made at other resources, and bottlenecks can shift as the schedule generation process proceeds. In both of these situations, many of the decisions made during generation of an entire bottleneck schedule may be artificially constraining (i.e., there exist other more-critical decisions from the standpoint of resource contention that should be considered first) and thus a coarse-grained approach to subproblem formulation can prematurely limit overall possibilities for schedule optimization. By reducing the granularity of subproblem formulation to the level of individual activities, i.e., operating with a control cycle that repeatedly determines and schedules the most critical unscheduled activity, the flexibility to better exploit possibilities for overall optimization is provided.

Such activity-based approaches to scheduling closely parallel the basic constraint satisfaction problem solving (CSP) paradigm discussed above in Section 4.3. The important distinction here is the introduction of so-called "variable ordering" and "value ordering" heuristics for *dynamically* reprioritizing (1) the set of scheduling decisions that remain to be made at each problem solving step and (2) the set of commitments that can be made relative to the selected (highest priority) scheduling decision within a given problem solving step. As indicated above, this dynamic reprioritization of possible choices is based on repeated analysis of current resource contention. Within the basic CSP approaches described in Section 4.3, in contrast, variable selection (i.e., which scheduling decision to consider next) and value assignment (i.e., which scheduling alternative to commit to) proceeds with respect to fixed (and often randomly established) orderings of alternatives (with these orderings being pruned whenever possible as a result of dependency-directed backtracking and shallow learning).

The first reported investigation of such opportunistic, activity-based approaches to scheduling was that of [68, 69], referred to as "interaction sensitive planning". This work introduced and explored the principle of coupling variable selection based on activity criticality from the standpoint of resource contention with value assignment based on minimizing impact with respect to resource contention. A more sophisticated approach to micro-opportunistic, activity-based scheduling, based on the use of a probabilistic model of the underlying search space, was developed and extensively analyzed within the MICRO-BOSS scheduler [91, 125, 126, 127, 128]. In fact, two variants of this probabilistic model were developed: the first based on the idea of counting feasible schedules at each decision step (viewing the scheduling problem strictly as a constraint satisfaction problem), and the second incorporating a cost model in the probabilistic analysis (thus interpreting the scheduling problem as one of minimizing overall costs). Variable and value ordering heuristics based on similar concepts of criticality and minimal impact were developed for both "constraint satisfaction (CSP)" and "constraint optimization (COP)" problem formulations. Comparative experimental studies conducted with the CSP model demonstrated the superiority of these heuristics vis a vis the "interaction sensitive" heuristics of [69] in terms of computational efficiency. Other computational studies relative to the COP formulation (which more accurately reflects the production scheduling problem) demonstrated superiority of the MICRO-BOSS approach in solution quality (i.e., minimizing wip, tardiness, earliness costs) over a collection of dispatch scheduling heuristics and several macro-opportunistic variants of the MICRO-BOSS algorithm. A third approach to micro-opportunistic, activity-based scheduling, developed independently but bearing considerable similarity to the MICRO-BOSS framework, is reported in [10, 11, 12] and used as a basis for strategic level decision-making within the DAS distributed scheduler (see Section 4.12 below). One key difference here is the use of global utility functions (as opposed to a cost model), which enables the scheduler to be focused according to specific management objectives.

Pragmatically, the issue of granularity in opportunistic approaches to scheduling is a tradeoff that must be struck between the better quality solution that can be produced by more flexible, finer-grained control

and the added computational expense associated with more frequent analysis of the current problem solving state. There are several factors that contribute to this tradeoff:

- nature of resource demand patterns - If patterns of resource demands in the target domain are quite sensitive to individual resource allocation decisions (e.g., relatively even loads across all machines in a flow line), then finer-grained control is more likely to yield substantial benefits. On the other hand, if resource demand patterns in the target domain are not very dynamic (i.e., there are clear bottleneck resources and clear non-bottleneck resources), then more frequent analysis may provide little additional information, and the optimization advantages gained at additional computational expense may be negligible (although experimental results in [125] point out that some seemingly static resource demand patterns are in fact quite dynamic).
- nature of uncertainty - Depending on the level of uncertainty present in the execution environment, there may be little operational leverage to be gained by finer-grained decomposition. If status in the execution environment is rapidly changing in an unpredictable fashion, a higher premium might be placed on computational efficiency (to keep up with the environment) than on producing the most optimized result. In contrast, if the execution environment is fairly stable, then the extra computational expense may be well worth the effort.
- complete vs incomplete search framework - Much of the micro-opportunistic scheduling research to date has emphasized the use of backtracking search through previous decisions when infeasible solutions are encountered. Under these assumptions, [125] has shown that a micro-opportunistic approach actually results in less backtracking than more coarser-level decomposition counterparts (and thus is more efficient). However, it remains an open question how approaches that rely on backtracking search over an "expanded" search space that includes various constraint relaxations (to guarantee the existence of a solution) compare in performance to approaches (such as those adopted by the coarser-level approaches discussed in Section 4.5 or for that matter the micro-opportunistic approach described in [12]) where the original (typically over-constrained) space is searched and constraints are selectively relaxed as necessary. Under such "constraint relaxation" approaches, the added computational cost of finer-grained decomposition becomes directly proportional to the cost of problem space analysis and the other factors raised above become the parameters of the efficiency/effectiveness tradeoff.

We do not yet have a good understanding of these tradeoffs. At the same time, the results that have been obtained with micro-opportunistic schedulers make it clear that such approaches offer new opportunities for optimizing production. Micro-opportunistic, activity-based scheduling is an important area for future research.

4.7. Fuzzy Scheduling

Given the uncertainty inherent in most manufacturing environments, convincing arguments can be made against approaches to detailed scheduling that yield "crisp" assignments to activity start and end times. The average time until portions of such a schedule are invalidated by events on the factory floor and require revision will typically be small (often minutes). Moreover, as argued in [71], significant computational effort may be unnecessarily expended when scheduling over longer horizons to avoid scheduling conflicts (e.g., minor temporal overlaps) that are really insignificant. This perspective has led some to investigate approaches to scheduling which explicitly factor uncertainty into the generated schedule [25, 37, 38, 71, 97]. The most developed of these efforts have been based on the use of fuzzy sets.

Under fuzzy set approaches, temporal constraints (e.g., due dates, activity start and end times, activity durations) are represented as fuzzy numbers, which represent sets of numbers (actually ranges) with associated levels of confidence as to membership in the set. Resource capacity constraints are defined

similarly, in this case relating to the amount of overlap in resource requests. Fuzzy arithmetic operations are used as a basis for combining fuzzy numbers and constraint propagation (e.g., through operation precedence constraints). FSS [71] uses fuzzy constraint management techniques in conjunction with a rule-based decision-making framework to produce job shop schedules in which fuzzy constraints (e.g., due dates) have a specified minimum threshold degree of satisfaction. The heuristic scheduling strategy implemented here focuses first on recognizing the primary bottleneck resource (from propagated fuzzy activity start and end times) and ordering the operations that require it, and then completing the schedules for other resources in a similar manner. Domain specific managerial and technological preferences are incorporated in situations where time constraints are not severe.

Although FSS has only been applied to a simplified job shop domain, [71] claim several potential advantages of fuzzy scheduling:

- generated schedules exhibit a graceful degradation of precision as one moves forward along the time line,
- scheduling time is not wasted producing consistent decisions further out on the scheduling horizon. Experimental results have indicated that their fuzzy approach is much more efficient than a deterministic counterpart, despite the increased expense of fuzzy constraint propagation.
- If the timings of specific future events are known (e.g., scheduled maintenance periods), they naturally provide "islands of certainty" from which to plan.

Another approach to fuzzy job shop scheduling, similarly based on fuzzy temporal constraint management techniques, is described in [37]. In [25], fuzzy set concepts are explored as a basis for post-processing an existing crisp schedule to hedge against machine failures in the context of a single machine problem.

A different application of fuzzy set concepts to production scheduling is exploited in OPAL [5, 6], a job shop scheduling system which constructs sequences of jobs for individual machines subject to operation precedence constraints and imposed release and due dates. The issue here is mediation among a set of expert scheduling heuristics, which range in origin from work shop managers to scheduling experts, potentially provide antagonistic advice, and are typically only relevant in a specific but generally uncertain range of situations. Within OPAL, such rules are used to provide advice as to the next sequencing decision to make during schedule development. Remaining choices (or conflicts) are identified by a supporting constraint management system. Fuzzy set theory is used as a basis for a weighted voting scheme, in which advice for alternative decisions is accumulated from different rules. A commercial version of the OPAL system is reportedly currently under development [5].

4.8. Planning and Scheduling

Much of the work in constraint-based approaches to production scheduling has relied on "pre-compiled" representations of part production processes as partially (or totally) ordered networks of potentially schedulable activities, which are known at the outset of any attempt to solve the scheduling problem. It is this representational assumption, for example, that allows propagation and analysis of time and capacity constraints for purposes of identifying bottleneck regions. The use of a pre-compiled representation by no means prohibits the possibility of production alternatives (e.g., alternative resources and processes), but constraint propagation complexity results regarding disjunctive constraints do place practical limits on the extent of disjunction that can be reasonably tolerated. In many manufacturing contexts, there is sufficient structure in the manufacturing process that such representational assumptions are reasonable.

However, in other production management environments (e.g., Flexible Manufacturing Systems) there is not only flexibility in routings of individual jobs, but a large set of supporting activities (e.g., part movement, machine reconfiguration, part buffering, etc) that must additionally be managed to

successfully accomplish target production activities. The presence and character of many of these supporting activities in the schedule is state-dependent. The need for a part buffering operation (along with transport operations to and from the store) depends on the state of the manufacturing system at the point when the part in question is initially ready to move. Likewise, the need for machine setup activities depends on the prior state of the machine.

The problem of effectively attending to such state-dependent production constraints has led to approaches that depend more directly on AI generative planning representations and techniques (e.g., [35, 98, 133, 135, 158]). These approaches operate with explicit models of the dynamics of actions (e.g., descriptions of the required enabling states of actions - or pre-conditions - and the resulting effects of executing actions - or post-conditions). Activity networks are constructed by a search that either works forward from an initial state description (e.g., all jobs ready to load into the FMS) to a goal description (e.g., all jobs completed and unloaded from the FMS) or backward from the goal description to the initial state, repeatedly matching activity pre and post conditions to the description of the current state. In some cases (e.g., [35, 98]), search flexibility is provided by using temporal constraint management techniques to maintain an underlying description of the evolution of the state of the manufacturing system over time.

The problem of managing the complexity of the search while producing a good schedule remains. In many cases, the issue of efficient resource allocation has been de-emphasized in favor of simply generating a feasible plan. In applying the SIPE planning system to the problem of scheduling beer production, an overall plan/schedule for production (including necessary plumbing reconfiguration activities, etc.) is incrementally generated by repeatedly selecting the next highest priority order and adding the activities required for that particular order into the current schedule (similar in overall decomposition assumptions to the ISIS approach). The resulting schedule is feasible, but opportunities to optimize production through direct consideration of resource utilization constraints are not exploited. In [35], a hierarchical, least-commitment temporal planning framework is employed, using a simulator to guarantee the existence of a feasible schedule each time a decision is made to add or connect a new activity to the plan. Aside from the obvious computational drawbacks of determining whether a feasible schedule exists from scratch at each planning step, this framework also places no emphasis on efficient usage of resources.

Other work has attempted to integrate resource allocation concerns into the plan generation process. In addressing the problem of FMS scheduling, [133] also assumes a job-centered decomposition of the problem, where individual plans for the production of each job are first developed independently using a heuristic search. At this point, however, plans are examined for resource interactions, and plan/schedule repair actions are taken to better synchronize resource usage. In the HSTS scheduler [98], an abstract level of representation where state-dependent constraints are modeled implicitly as adjustments to activity durations is employed to provide a basis for integrating dynamic constraint analysis and resource contention-based focus of attention heuristics with detailed generative planning. [55] has also recently described a Genetic Algorithm based approach to simultaneously optimizing process planning and resource allocation objectives.

4.9. Iterative Scheduling

The constraint-based approaches discussed above from Section 4.3 onward (with a few exceptions) have emphasized techniques for incremental convergence to a final schedule (or set of schedules if some degrees of freedom are retained in the final solution). Search proceeds by maintaining and incrementally extending one or more partial schedules and terminates when an acceptable complete solution has been found. We can view this type of search generally as one which incrementally winnows away alternatives as it proceeds (backing up when necessary) and thus restricts attention to increasingly smaller sets of candidate schedules (or regions of the underlying search space) over time. An alternative class of approaches, referred to here generally as iterative scheduling approaches, operate with respect to different ground assumptions as to the structure of the search. Within these approaches, search

proceeds (in the simplest case) by repeated movement from one complete schedule (or set of schedules) to another (or more generally from one set of points in the underlying space to another). Search terminates when either an acceptable schedule is found or a pre-specified amount of computational effort has been expended.

The use of iterative scheduling techniques in support of production management and control can be argued from several perspectives:

1. The iterative scheduling paradigm enables use of a variety of probabilistic search algorithms motivated by analogies to natural systems, many of which are provably efficient in searching non-linear, combinatorial spaces (such as those encountered in most production scheduling contexts). These algorithms, which include simulated annealing, genetic algorithms (GAs), and neural network techniques, provide general-purpose heuristics for finding schedules that globally optimize with respect to conflicting constraints and objectives. One conceptual advantage of using such techniques is that all domain specific knowledge can be localized in the specification of the schedule evaluation function.
2. The extensive sampling of points in the space of possible schedules that is performed under many iterative schemes can provide an inexpensive alternative to problem space analysis methods for building an understanding of the structure of the current scheduling problem. In [16], for example, iterative random reassignment of resources to activities is used to identify regions of high contention on resource time lines, which subsequently become foci for more extensive search effort.
3. Iterative scheduling approaches emphasize generation and evaluation of many different schedules. Thus, from a reactive scheduling perspective, a schedule can always be made available for use very quickly if needed, with the search continuing to find better schedules as decision-making time permits.
4. The iterative scheduling paradigm provides a natural framework for repair-based approaches to managing schedules over time (e.g., [106, 160]), where minimizing disruption to the current schedule is an important additional concern to be balanced with optimization and responsiveness objectives.

Several efforts have investigated the application of general-purpose probabilistic search procedures as a basis for schedule generation [28, 33, 45, 46, 44, 53, 60, 152, 157, 160]. One difficulty in applying such techniques to scheduling problems is that of specifying a suitable problem representation. The effectiveness of general-purpose algorithms like GAs, simulated annealing, and neural networks depends heavily on an ability to easily and rapidly generate new candidate schedules for evaluation. This requirement is typically at odds with the need to account for the idiosyncratic constraints of the target production environment, which complicate the representation of the search space. Reconciliation of these antagonistic requirements pragmatically entails either (1) the use of domain specific heuristics in the "candidate solution generators" (to insure the feasibility of each new schedule that is generated) or (2) the introduction of penalties into the schedule evaluation function to bias the search away from infeasible schedules when they are generated. Consequently, much of the work reported in using such techniques has the drawback that fairly simplified representations of problem constraints are employed. One interesting exception, representing perhaps the most promising approach to this representation problem, is the work of [152]. In this case, a GA is used in conjunction with an efficient domain-specific simulator (called the "schedule builder") to generate new candidate schedules for evaluation. The GA is focused strictly on the simpler, more-manageable representation space of possible sequences, and each sequence generated is then transformed into a schedule that satisfies all domain constraints by the schedule builder. This GA-based system has been successfully applied to the problem of scheduling training exercises at a Naval firing range and is currently being placed into operational use. The approach taken to problem decomposition is equally applicable to production scheduling domains.

A second broad focus of work in iterative scheduling techniques (as suggested above) has been

schedule repair [16, 106, 60, 95, 130, 140, 160]. Repair-based scheduling methods proceed from the assumption that the current schedule contains some number of constraint conflicts (i.e., the schedule is infeasible in that resource capacity might be over-allocated, activity precedence constraints might be violated, etc), and the goal (minimally) is to produce a feasible schedule. Current conflicts may be a consequence of new information about the status of the factory (e.g., a heavily allocated machine has failed and will now be down for some amount of time) or the result of prior scheduling actions. Emphasis is placed on the use of local repair actions aimed at resolving specific conflicts. Approaches vary according to objectives considered (e.g., achieving a feasible schedule, maintaining a schedule that continues to reflect optimization objectives, minimizing disruption to the starting current schedule, speed of the overall reactive process), granularity of the repair actions taken, and positions taken as to where search energy should be expended. The NEGOPRO [130, 129], SPIKE [60, 95], and KEDS [160] schedulers, for example, operate with simple, small-granularity repair actions (e.g., reschedule most-constrained conflicting activity at its least disruptive feasible alternative) and a new solution is produced through extensive iterative application and evaluation of the results of these simple actions. The OMP scheduler [15, 16] shares this philosophy of simple repair actions and extensive iteration, but operates with larger granularity repair actions (in this case, randomly rescheduling activities that contribute to current bottleneck regions of the timeline). Schedule repair in OPIS [106], alternatively, is representative of more deliberate, knowledge-directed approaches, relying less on extensive iteration and more on extensive analysis of current solution constraints at each step and the use of larger-grained, search-based repair actions that are matched to the recognized optimization needs and opportunities for non-disruptive schedule revision produced by conflict analysis.

4.10. Interactive and Mixed-Initiative Scheduling

Most of the work described thus far has emphasized the development of knowledge-based scheduling techniques that more effectively address practical production management and control problems in some respect than do conventional techniques (e.g., produce solutions that more accurately reflect domain constraints and objectives, produce solutions that account for executional uncertainty, incrementally patch schedules when changing circumstances dictate). Research in interactive and mixed-initiative scheduling, alternatively, has been grounded by objectives that center around effective user support and involvement. This research is rooted in the pragmatic assumption that many aspects of production management decision-making (even at operational levels) can never be totally automated, and attention should instead be focused on the development of tools that effectively support the human scheduler's task. In some cases, interactive scheduling goals have been considered complementary to those that have driven the development of automated techniques, and work has focused on frameworks for interactively exploiting these techniques (or their constituent components) as decision-making aids [40, 59, 63, 88, 118, 129]. However, other work in interactive and mixed-initiative scheduling has argued that issues relating to appropriate user and system roles in the production management process impose specific system requirements, which should actually focus the development of supporting techniques [13, 86, 100].

One commonly recognized basis for effective user support in developing and managing production schedules is constraint checking. Several interactive scheduling tools (e.g., SEMIMAN [40], LOGOS [88]) have been built directly around the basic constraint management and temporal reasoning frameworks described in Section 4.3. Such frameworks provide convenient specification of complex time and resource capacity constraints on manufacturing activities to be scheduled, allow the user to see the immediate consequences of specific scheduling decisions, and keep the user continually aware of both feasible scheduling alternatives and critical constraint interactions. More recent work with probabilistic and utility-based representations of constraint interactions (e.g., [12, 97, 125]) raise interesting possibilities with respect to more sophisticated constraint checking capabilities. In fact, early work with the ISIS scheduler demonstrated representational techniques that enabled constraint checking to be extended to include consideration of preferences, providing assessments of the utility as well as the feasibility of user decisions [144].

Frameworks for interacting with such constraint checking and management components, and for shifting decision-making responsibility from the user to the system have adopted distinct philosophies. At the most basic level, there is general agreement in the need for schedule building and editing facilities such as selection, deletion, insertion, substitution, and relocation of activities. However, perspectives diverge with respect to higher-level user interaction. Some approaches (e.g., [88, 129] have advocated user focusing of system scheduling actions within backtrack search frameworks. In [91], it is argued instead that scheduling activity should be based on a model that allows schedules to be directly manipulated without reverting to a previous state and rebuilding from that point. Still other approaches (e.g., [86, 100]) have developed models of interaction based on perceived spheres of expertise of human and machine decision-making. Likewise, there are varying viewpoints on the appropriateness of various approaches to user specification of aggregate scheduling actions (e.g., reschedule job *x* to start tomorrow) and system treatment of constraint conflicts (e.g., default resolution/relaxation strategies vs disallowance of conflict generating actions). At the user interface level, recent research has advocated the concept of a "electronic leitstand" [63]: a production control "command center" that combines graphical displays with underlying constraint checking and scheduling capabilities to provide a variety of editable views of the current production schedule. Leitstand tools are having an increasing operational impact (over 100 installed systems by conservative estimates) as a means for integrating existing MRP planning systems with real-time shop floor control.

One somewhat novel framework for interactive scheduling, rooted in the assumption that the source of human schedulers' unique expertise is their knowledge of uncertainty, is proposed in [86]. Based on the results of an extensive study of human schedulers in a large number of job shop environments, this work argues that schedulers become experts over time by (1) developing better and better categorizations of various sources of uncertainty in the production environment (which in turn provide a basis for distinguishing between important problems that require attention and regions of stability in the schedule that can be ignored), and (2) getting better and better at predicting the consequences of identified uncertainties. Two conclusions are drawn from this viewpoint with respect to the design of automated scheduling tools. First, the human scheduler's knowledge of uncertainty will always outstrip the knowledge contained in the system, and thus the human scheduler must always be ultimately in charge of decision-making. Second, this process of evolving expertise in classifying and managing uncertainty provides a good model for scheduling system design. Under the proposed scheme, system scheduling, at any point in time, is driven by rules which map categorizations of certainties and uncertainties in the current production environment to appropriate scheduling procedures, and results are always then passed to the human scheduler for final revisions. The system monitors user manipulation of the schedule, requesting the reasons for each revision that is made. This information is then used to augment/refine the system's knowledge of uncertainty categories as well as its rules for responding to recognized uncertainties. While this approach remains a proposal at this point, techniques that have emerged from the field of machine learning give the approach credibility and provide a basis for investigating its viability.

4.11. Heterogeneous, Multi-Level Approaches

The preceding sections have outlined several AI-based techniques, frameworks and methodologies relevant to various aspects of practical production management and control problems. For the most part, the specific research and development efforts described have emphasized the uniform use of a particular method or methodology, making particular (sometimes unstated) assumptions about the type of the production environment of interest, its dominant characteristics, and the level at which planning and control decisions are being made. In this section, we consider work which has proceeded from a somewhat broader perspective: that practical production management problems require an integration of disparate techniques. This work has concentrated both on approaches to integrating the strengths of various analytic and heuristic decision-making models (e.g., [32, 121, 154]), and on frameworks for heterogeneous, multi-level decision-making over different time scales. In some cases (e.g.,

[32, 50, 84, 159]) problem decomposition assumptions, system organization and decision-making heuristics have been heavily influenced by the particular characteristics of the target domain. Others (e.g., [153, 83, 96, 121]) have sought more generally applicable frameworks for integration. We summarize a few successful and representative systems in the following paragraphs.

REDS [34, 51, 50] is a multi-level production planning and control system developed by Siemens for use in managing VLSI production. The dynamic nature of VLSI manufacturing environments is directly reflected in the design of REDS, which operates according to the general principle of least-commitment planning. Production plans are developed that retain maximal flexibility to local constraint and requirements changes and allow some of the problem's disruptive dynamics to be absorbed. This is accomplished through two principal devices. First, the production planning process is decomposed into a set of heterogeneous decision-making levels that operate over different time scales. Second, various abstractions of domain constraints (in particular resource capacity) are utilized as a basis for decision-making at different levels.

More specifically, the architecture of REDS identifies three types of decision-making that must be integrated at different stages in the VLSI production management process:

1. feasibility analysis - Feasibility analysis is concerned with the general questions of whether and (ultimately) when new jobs can be processed. This is accomplished using a job-release heuristic [51] designed to globally balance to management specifications of the ratio between jobs completed on time and machine utilization. This strategy is augmented with heuristics for relaxing constraints if necessary (e.g., extending due dates, adding work shifts, and pre-empting existing orders). At this level of decision-making, only the capacity of known critical, bottleneck resources is taken into account (referred to as the "essence" of an order's requirements).
2. detailed scheduling - Detailed scheduling is concerned with distributing the production requirements of a given order over particular production intervals (i.e., determining what quantities will be produced when within the overall production interval set during feasibility analysis). To this end, a hierarchy of production intervals (or time buckets) is defined (e.g., ranging from a year interval at the root, to constituent quarters at the next highest level to constituent months at the next level and so on down to day intervals), and associated with each node is a "constraint pool" defining local constraints (e.g., resource capacity, management objectives, etc.). This abstraction hierarchy can be exploited in a breadth-first manner for purposes of long term production planning or in a depth-first manner to establish short-term, daily production requirements. A load balancing heuristic is used as the basis for decision-making at a given level of granularity.
3. real-time scheduling - Real-time scheduling is concerned with the sequencing decisions that must be made to keep production moving and achieve the production requirements established for the current day by detailed scheduling. Decisions at this level are made through application of a simulation-based scheduler built around the Dynamic Sequencing Rule [132]. This rule is sensitive to important aspects of the current production state, including setup time, availability of downstream resources, buffer capacity, and management objectives such as order priority (hot or not), higher machine utilization, lower cycle time, or minimum tardiness. All factors are assigned weights to provide a method of evaluation (and the final priority).

These three decision-making modules are applied in a top-down manner over increasingly smaller scheduling horizons. A fourth module, called the statistician, monitors the real-time data stream to provide feedback to detailed scheduling and feasibility analysis levels about longer term problems (e.g., identification of developing bottleneck work centers or production quality problems). REDS has been in operational use at a West German VLSI fabrication facility for over a year now, and its use is said to have led to a 10% increase in productivity.

A second framework for multi-level production planning and control is illustrated by the AISIS family of

schedulers [159, 65, 64, 66, 140], developed by INTEL Corporation and recently installed in a wafer fabrication facility in Albuquerque, New Mexico. AISIS similarly structures the production planning and control process into decision-making levels, in this case providing a collection of schedulers appropriate to specific short term production management activities. In more detail, AISIS is comprised of the following components:

- V1 - The V1 scheduler is operates with respect to a planning horizon comparable to the cycle time of products through the fab, supports analysis of fab capabilities to meet demands, and provides a basis for establishing job release dates. It is designed to provide flexible interactive support to a human planner and schedules are developed using a simulation-based approach.
- V2 - Operating with respect to the constraints and objectives set by V1, the V2 scheduler is responsible for developing a detailed production schedule (i.e., resource assignments and time table) for the next work shift in each of the four principal areas of the fab. Decision-making at this level is based on a more detailed model of fab operating constraints. The simulation-based approach of the V1 scheduler is abandoned in favor of a more constraint-directed framework based on the combined use of a collection of scheduling methods and heuristics [64]. Reminiscent of the concept of multi-perspective scheduling employed in OPIS, schedules are developed incrementally through repeated recognition of the most dominant scheduling concerns (e.g., hot jobs, bottleneck resources) and selection of the method or heuristic best suited to the characteristics of the identified scheduling subproblem.
- V3 - The V3 scheduler is responsible for reactively managing the V2 schedule as execution proceeds. It is incrementally delegated responsibility for the executing frontier of the schedule, issues control decisions to the factory floor control system, and monitors fab status as the results of executing job steps are received. When mismatches are detected between the expected and actual fab state (signaling problems), a repair-based scheduling strategy [140, 77] is invoked which attempts to exploit existing job slack and machine idle time in the predictive V2 schedule to locally rearrange execution. If execution problems cannot be locally resolved, they are passed back to the V2 level for more global consideration.

To ensure accuracy of the information driving the scheduling process at each level, an underlying data consistency checking module interprets the real-time data stream produced by the factory control system. Here an expectation-based model of status is exploited to infer missing information (perhaps due to operator error) and detect inconsistencies.

Both REDS and AISIS provide frameworks for integrating various operational levels of production management and control decision-making. Other such multi-level frameworks are described in [83, 84, 96]. What these frameworks don't address, however, are the larger issues of integrating strategic and operational planning, and integrating of production management activities with other functional activities of the manufacturing enterprise. One of the few research efforts to date to investigate the potential of knowledge-based technology in this larger arena is the SIMENS system [153]. SIMENS is an intelligent support system aimed at integrating strategic "market positioning" decision-making with operational production planning and control activities in various functional areas of the enterprise. The approach is targeted at multi-stage production and assumes that the key decisions linking strategic and operational planning concern product buffer sizes at intermediate production stages. The key underlying concept is that buffer policy determines the firm's position in the marketplace. For example, larger amounts of product buffering advocates faster product delivery at the expense of higher inventory costs. Buffer policy decisions are supported through generation of a so called "frontier curve" (FC), which indicates the optimum tradeoff between average inventory cost and average order processing delay that can be realized by some buffer strategy (assuming the operational production scheduling strategy in use). By contrasting the firm's FC with those of its competitors, points on the curve where a competitive advantage is held can be identified. Analytic models for generating the FC and determining the buffer strategy once a point on the curve has been selected provide the tools for strategic decision making within SIMENS. The determined buffer strategy together with projected demands provide production

requirements that drive the development of production plans at operational decision-making levels.

At the operational level, the concept of "controlled reasoning" is used as a basis for integrating the activities of various functional units and for responding to unexpected exogenous events. In summary, problems are detected from monitoring of real time events and classified as either "minor deviations from plans", "familiar deviations" or "unfamiliar deviations". In the first case, no reaction is warranted, and problems of the second type are solved by tried and true procedures (expressed within SIMENS as a collection of rules and algorithms). In the case of unfamiliar deviations, a set of rules is consulted to first determine which functional unit(s) should be called upon (e.g., marketing, production, purchasing, distribution) and which problem(s) it should solve. Associated with each functional unit modeled in the system is a set of "analysis templates" and a set of rules for determining the relevance of alternative templates to the current problem. Analysis templates may correspond to mathematical models, heuristic procedures, or report generators and may variably suggest corrective planning actions, indicate required involvement by other functional units, or simply provide the human decision-maker with information upon which to base reactive planning decisions.

Aside from the approach to linking strategic and operational decision-making, the important contribution of SIMENS lies not in the specific knowledge, models and heuristics employed for analysis and decision-support (in fact many off-the-shelf components such as MRP modules are actually used) but rather in the framework developed for structuring and applying knowledge and methods particular to different functional activities and knowledge relating to their integrated use in specific planning circumstances. While aspects of the controlled reasoning framework and their underlying assumptions can be criticized and certainly warrant further research, SIMENS, nonetheless suggests the potential of knowledge-based technology in integrating decision-making at the enterprise level.

4.12. Distributed Production Management

Another important requirement for integration in practical production management domains is support for distributed decision-making. Factories are, in fact, decentralized systems with many distinct activities proceeding concurrently, and the organizational structures that drive them involve multiple decision-makers. Decentralization and distribution of production management and control responsibility are a practical necessity to insure adequate responsiveness to unfolding manufacturing circumstances over time. Accordingly, the application of knowledge-based techniques to the problem of distributing production management and control has been the focus of a growing number of recent research efforts [17, 21, 56, 61, 106, 113, 145, 150].

Within the subfield of distributed AI, a variety of frameworks have been developed for cooperative problem solving by multiple decision-making "agents". The "contract net" formalism [136] provides a structured framework for distribution of subtasks among agents through protocols for task announcement, collection of bids and awarding of tasks to bidding problem solving agents. Blackboard-based systems [41] (discussed previously in Section 4.5), provide an alternative framework for integrating the differential expertise of a group of problem solvers, in this case based on the use of a globally accessible data base (the blackboard). Other message-passing formalisms (e.g., [2]) provide mechanisms for more focused coordination between specific problem solving agents. Collectively, these frameworks have provided a base for investigating approaches to distributed production management. At the same time, the characteristics of scheduling and resource allocation problems diverge significantly from those of the types of problems historically considered by field of distributed AI, and research in distributed production management has reciprocally led both to extensions/specializations of the above frameworks and the development of alternative, more appropriate coordination structures and protocols (e.g., [123, 124]).

Research to date in distributed production management has largely focused attention on mechanisms and frameworks for shorter term production scheduling and control. Approaches can be differentiated and contrasted with respect to problem decomposition assumptions made along several dimensions.

One fundamental issue that arises in distributing problem solving responsibility among multiple scheduling agents concerns how the overall problem should be partitioned across distinct agents, and how this set of agents should be organized. Most work has exploited some form of structural decomposition of the manufacturing system (e.g., the entire factory at the highest level, constituent work areas at the next level, smaller work areas at the next level and so on) as a basis specifying an organization of scheduling agents and partitioning scheduling responsibility. In many cases [21, 61, 113, 145, 56], hierarchical organizational structures have been super-imposed over this structural decomposition (motivated by many of the same issues that motivated the multi-level systems described above in Section 4.11). Scheduling responsibility is ascribed to individual agents at each level by varying (1) the portion of the overall manufacturing process considered, (2) the granularity of the constraints used to drive decision-making, and/or (3) the types of decisions made. Other work has emphasized heterarchical frameworks for cooperative scheduling, focusing instead on mechanisms for locally coordinating decision-making in various areas in the factory at some level [17, 106, 116, 150].

A second dimension of problem decomposition along which approaches can be distinguished concerns the capabilities of individual scheduling agents: whether agents are assumed to have homogeneous capabilities and thus responsibility for specific components of the overall solution (e.g., responsible for maintaining schedules for particular areas of the factory) or whether agents have differential capabilities which define their role in developing the overall solution. The distributed CORTES project [151, 150] provides one example of a homogeneous agent approach, where investigation centers on decentralization of the MICRO-BOSS approach to scheduling (Section 4.6). The current order set is partitioned and distributed to a set of agents, each operating with a copy of the MICRO-BOSS scheduler, and resource demand profiles and reservations are communicated as a means of coordinating local resource allocation decisions. A second homogeneous agent architecture, aimed at lower-level work cell control, is defined in the NUTS system [17, 101].

The CSS system [106], in contrast, adopts a heterogeneous agent approach. In this case, a problem decomposition scheme was investigated that was motivated more by the existing approach production scheduling in the target domain (a Pratt and Whitney jet engine factory), where schedules were developed and maintained through compromise between agents with different production management objectives. Using a contract net framework, "resource broker" agents, biased toward resource-based scheduling objectives, were defined for different areas of the shop, and a "work order manager" agent, motivated by order-centered scheduling objectives, was defined to bid for resource capacity and award job contracts to resource brokers. The DAS scheduler [21, 119, 12] implements a second heterogeneous agent approach, in this case also exploiting a three-level hierarchical decision-making framework. At the lowest level, agents are responsible for maintaining the schedule for each individual machine in the factory. At the second level, tactical agents preside over specific groups of substitutable machines and are responsible for distributing jobs across these machine groups. A single strategic agent sits at the top level, responsible for selection among alternative process plans, establishing job release sequences and relaxing due dates if necessary. Coordination is accomplished through both message-passing and a centrally accessible description of the current schedule.

A third dimension of problem decomposition along which approaches have varied is the granularity of agent decision-making. At one extreme along this dimension is the CASCADE system [115, 114, 117, 116], designed for real-time management of material flow in an FMS. Agents associated with each work station in the FMS are each provided with simple, parameterized rules for directing and receiving material, that are applied within a contract net coordination framework. Specific parameter settings give rise to different identifiable work flow behaviors (e.g., just in time, just in case, etc.). System developers have since observed that the system's fine-grained organizational structure maps directly to a neural network, suggesting the use of neural network propagation techniques as a means of tuning system parameters to the specific production environment. The OPIS-based BOSS system for decentralized production management [145, 56, 57], alternatively, operates with a much coarser partitioning of scheduling responsibility. Here, a hierarchical decision-making framework similar to that of

the DAS scheduler is employed. Within BOSS, however, lowest level agents have responsibility for larger segments of factory operation, and agents at different levels in the organization are concerned with managing schedules at different levels of abstraction over different time horizons. An underlying hierarchical model of the domain and its constraints provides a framework for coordination: constraints are imposed top down with respect to abstractions of time and capacity constraints; lateral message passing enables cooperative management of more detailed schedules within these imposed constraints by lower level agents over shorter time horizons.

Research in distributed production management is still at too formative a stage to draw general conclusions. The diversity of approaches explored in large part reflects differing emphases with respect to type of manufacturing environment and/or level of production management decision-making, and it is likely that eventual, overall solutions in specific production management contexts will combine elements of several approaches. It is, however, clear that frameworks for decentralization of production management and control activities are essential to more flexible organizational behavior, and consequently an important area for future research.

4.13. Learning Shop Control Policies

In the discussion of interactive scheduling frameworks in Section 4.10, the use of mixed-initiative learning techniques was identified as a means for incrementally building system knowledge about uncertainty in the execution environment. Other more recent work has considered the use of induction techniques for learning from examples (developed within the field of Machine Learning) as a basis for developing effective factory floor control policies [23, 134]. This work adopts a dynamic, state-dependent approach to the control problem (similar, for example, to the LMS system described in Section 4.1 and the classical dispatch rule literature [111]), and has focused on the issue of inferring appropriate control policies from examples of the behavior of specific dispatch decisions or dispatch strategies under different problem circumstances. So called "training examples" are obtained through repeated simulation of the target manufacturing system, each identifying a relevant characteristics of the manufacturing state (e.g., balanced/unbalanced load, contention metrics, buffer sizes, etc.) and resulting performance. In [134], for example, such training data is used to infer a decision tree which expresses the conditions under which each of a set of predefined dispatch rules should be used as a basis for control decisions in an FMS environment. Experimental results in this case show that the more complex, learned control policy significantly outperforms any one of the individual base rules.

Speaking more generally, it would seem that such machine learning approaches hold the promise of broader applicability to other, pro-active aspects of production management as well. Many of the approaches described above advocate the cooperative use of several strategies and heuristics both in generating and revising production schedules (e.g., [29, 66, 88, 140]), and it would seem that a similar approach to generating and/or validating appropriate conditions of applicability (in this case additionally relating to the state of the schedule and scheduling process) could be taken. Moreover, many inductive machine learning techniques are incremental in nature, raising the possibility of using examples of actual manufacturing behavior in addition to simulation-driven approaches as a means of adaptively improving decision-making performance over time. This is one obvious direction for future research.

5. Prospects and Challenges

In this paper, we have provided an overview of research in the field of knowledge-based production management. We have categorized this work along several different dimensions, identifying the principal types of manufacturing domains that have received attention, the particular production management and control activities that have been emphasized, and the various perspectives that have emerged with respect to the tradeoff that must be made in practical production management contexts between predictive decision-making to optimize behavior and reactive decision-making to manage executional

uncertainty. The bulk of the paper has focused on summarizing the dominant approaches to knowledge-based production management that have emerged. Here, our goal has been to identify the general concepts, principles, and techniques that distinguish various paradigms, characterize the strengths and weaknesses of each paradigm from the standpoint of different production management requirements, and indicate the results that work within each paradigm has produced to date. Among the paradigms for knowledge-based production management considered were rule-based scheduling, simulation-based scheduling, constraint-based scheduling, fuzzy scheduling, planning and scheduling, iterative scheduling, and interactive scheduling. We also examined work aimed at integrating heterogeneous planning and scheduling methods (both AI and OR based) and the construction of systems for multi-level production management and control. Finally, we surveyed more recent research in the areas of distributed production management and automated learning of factory floor control policies from experience. We conclude by commenting briefly on the practical prospects of this work. In doing so, we also identify some of the important obstacles and challenges currently facing the field.

It is no longer a question of whether research in knowledge-based production management will impact manufacturing practice, but rather to what ultimate extent and in what time frame. Several examples of operational systems employing knowledge-based techniques have been identified, and many have been attributed with substantial manufacturing performance gains. Examination of these successful applications gives insight into the current obstacles to technology transfer as well as a basis for projecting the likely future progression of research results into operational production management environments. With very few exceptions, successful applications have emphasized knowledge-based tools for short-term production scheduling and control. Insertion of this technology has followed a lengthy, incremental development path involving (1) development of the data collection and information access capabilities necessary to support computerized decision-making and decision-support (including necessary linkages to existing systems), (2) introduction of flexible interactive decision-support capabilities (e.g., interactive schedule manipulation with constraint checking), and (3) subsequent expansion of production management knowledge and system decision-support functionality to encompass a larger and larger decision-making role. Thus, we can expect more generally that research concerned with more sophisticated constraint checking and interactive scheduling frameworks will continue to have the most significant short-term impact on actual manufacturing operations. It is important to note, given the current state of computerized production management support in many environments, that steps 1 and 2 above will often have a much larger impact on factory operations than will the introduction of more sophisticated techniques for planning, scheduling and control. Nonetheless, as interactive tools and supporting data collection and access capabilities become integrated into production management operations, the utility for more sophisticated techniques for production management and control become more apparent, and the results of research in more sophisticated approaches production scheduling will begin to have a larger impact. In the longer term, focus will shift to integration of production management activities at different levels and involving different functional areas of the organization, again moving from interactive frameworks that support flexible coordination to more sophisticated frameworks that automate more of the coordination process.

One important gap with the respect to the current state of research in knowledge-based scheduling is the lack of results quantifying the relative advantages and disadvantages of various approaches that have been and are being pursued. In contrast to the field of OR, where specific classes of scheduling problems are precisely defined and alternative algorithms and heuristics must compete with current benchmark results, work in knowledge-based scheduling has tended to proceed more from the specifics of a particular scheduling domain and the validity of approaches has been argued solely on the basis of their success in addressing the problem at hand. This is, of course, more than adequate if the result is an operational system and actual benefits to manufacturing organization can be measured. But much of the research in the field is concerned, as it should be, with the development of approaches and techniques that enable more effective general solutions (using specific production management domains as research drivers), and, in this context, calibration of results obtained relative to other work in the field is an

important issue. Comparative experimental analyses that have appeared tend to be focused more on performance comparisons to OR-based alternatives (e.g., dispatch heuristics) that have not necessarily been designed with respect to the problem at hand. The field of knowledge-based production management does not currently have a strong set of guiding principles for operational system builders to draw on in mapping various approaches and techniques to the specific production management problems they face. This is evidenced by the "roll your own" philosophy that has been adopted in most system design and development projects to date, and this trend will continue until there is better understanding of the comparative strengths and weaknesses, and complementary relationships of alternative approaches.

There is, of course, some justification for this state of affairs. Unlike research in OR, the field of knowledge-based scheduling has concentrated on solution techniques for practical problems that are not easily formalizable, and the ultimate utility of a given approach is typically a direct function of the approach's ability to deal effectively with the specific operational constraints, objectives and characteristics of the target domain. The argued strength of knowledge-based scheduling has always been its ability to conveniently represent and incorporate knowledge about complex, domain-specific constraints and operating procedures. This emphasis, coupled with the diversity of different manufacturing production environments, justifies the broad range of approaches and techniques that have been explored. It is also not the case that there has been no interaction in the field and building on previous results. Quite to the contrary, one can see clear progressions in the approaches that have been reported over time (e.g., the work in constraint-based scheduling discussed in Sections 4.3 through 4.6) as well as convergence with respect to basic scheduling system frameworks and components (e.g., constraint management subsystems). One major goal of this paper has been to identify such trends and commonalities among work in the field.

At the same time, the field of knowledge-based scheduling has reached a point where there now is experience in many classes of manufacturing environments and production management problems. Many so called "generic" techniques and methodologies have emerged from this experience, and it is important at this stage for research to emphasize results aimed at comparison and characterization of strengths and limitations along this landscape of investigated applications and existing techniques. It should be possible, for example, to assemble a common set of benchmark problems, reflective of different types of production environment domains, different sets of assumptions with respect to dominant constraints, etc., to provide an standard, accepted suite of problems for comparative analysis [67]. A strong start in this direction could be made simply by making more experimental data accessible (e.g., [24]). The simple fact that most published experimental studies do not contain sufficient detail of the domain to enable reconstruction of the problem is a serious problem that the field should no longer be willing to tolerate.

With respect to overall direction of the field, this paper has identified several promising avenues of research aimed at improved functionality in specific production planning, scheduling and control activities. Continued research in areas such as use of problem state analysis and look-ahead methods, representation and use of knowledge about uncertainty, repair-based scheduling frameworks and methods, integration of resource allocation and plan synthesis processes, and cooperative use of AI and OR based techniques promises significant new opportunities for optimizing manufacturing system performance. This paper has also described more recent research activity in integrating production management and control decision-making at different levels and coordinating the activities of multiple decision-makers. Further research into methodologies for problem decomposition, representation of constraints at different levels of abstraction, protocols and strategies for negotiation, cooperative problem solving among heterogeneous agents, and integration of strategic and operational decision-making holds the potential for much more profound advances in organizational performance, and the areas of integrated and distributed production management can be expected to receive increasingly more attention within the field over the next few years. Perhaps the largest long-term payoffs of knowledge-based production management lie in the application of machine learning techniques. Practical production management is fundamentally a dynamic activity that draws on evolving knowledge of both the production environment and the external marketplace. Accordingly, techniques for incremental acquisition,

refinement and adaptation of production management constraints, objectives, operating procedures, and decision-making strategies are central to realizing the full potential of knowledge-based production management technologies. We have described some initial research activities and proposals in this direction, but the possibilities and technical issues here remain largely unexplored at this point.

References

1. Acock, M. and R. Zemel. DISPATCHER: AI Software for Automated Material Handling Systems. Proceedings ULTRATECH - AI in Manufacturing Conference, Society of Manufacturing Engineers, Long Beach, CA, September, 1986.
2. Agha, G. *ACTORS: A Model of Concurrent Computation in Distributed Systems*. The MIT Press, Cambridge, MA, 1986.
3. Allen, J.E. "Maintaining Knowledge About Temporal Intervals". *Communications of the ACM* 26, 11 (November 1983), 832-843.
4. Arff, S. and G. Hasle. AIPLAN: A Knowledge-Based System for Production Planning in Aluminum Foundries. Proceedings 3rd International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1989, pp. 603-617.
5. Badie, C., G. Bel, E. Bensana, and G. Verfaillie. Operations Research and Artificial Intelligence Cooperation to Solve Scheduling Problems: the OPAL and OSCAR systems. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 1-5.
6. Bel, G., E. Bensana, and D. Dubois. "OPAL: A Multi-Knowledge Based System for Industrial Job Shop Scheduling". *International Journal of Production Research* 26, 5 (May 1988).
7. Bell, C. and A. Tate. Using Temporal Constraints to Restrict Search in a Planner. University of Iowa, December, 1984.
8. Ben-Arieh, D. and C.L. Moodie. "Knowledge-Based Routing and Sequencing for Discrete Part Production". *Journal of Manufacturing Systems* 6, 4 (1985), 287-297. simulation-based scheduling, design issues.
9. Bensana, E., M. Correge, G. Bel, and D. Dubois. An Expert System Approach to Industrial Job Shop Scheduling. Proceedings International Conference on Robotics and Automation, San Francisco, CA, 1986, pp. 1645-1650.
10. Berry, P.M. Satisfying Conflicting Objectives in Factory Scheduling. Proceedings 6th Conference on AI Applications, Santa Barbara, CA, March, 1990, pp. 101-106.
11. Berry, P.M. Resolving Conflicting Objectives in Factory Scheduling. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 16-21.
12. Berry, P.M. *A Predictive Model for Satisfying Conflicting Objectives in Scheduling Problems*. Ph.D. Th., University of Strathclyde, Glasgow, 1991.
13. Fox, B.R. Mixed Initiative Scheduling. Proceedings AAAI Spring Symposium Series: AI in Manufacturing, Stanford University, March, 1989, pp. 25-28.
14. Fox B.R. & Kempf K.G. Complexity, Uncertainty and Opportunistic Scheduling. Proceedings of the Second IEEE Conference on AI Applications, IEEE-85, 1985, pp. 487-492.
15. Biefeld, E. and L. Cooper. Operations Mission Planner: Final Report. Publication 90-16, Jet Propulsion Laboratory, March, 1990.

16. Biefeld, E., and L. Cooper. Bottleneck Identification through Chronology-Directed Search. Proceedings 12th International Conference on Artificial Intelligence, Sydney, Australia, August, 1991.
17. Bowen, J., P. O'Grady, H. Nuttle and M. Terribile. "An Artificial Intelligence Approach to Loading Workstation Resources in a Distributed Job Shop Controller". *Computer Integrated Manufacturing Systems* (February 1989).
18. Brown, M.C. The Dynamic Rescheduler: Conquering the Changing Production Environment. Proceedings 4th IEEE Conference on AI Applications, San Diego, March, 1988.
19. Bruno, G., Elia, A., and P. Laface. "A Rule-Based System to Schedule Production". *IEEE Computer* 19, 6 (July 1986).
20. Bruno, G. and M. Morisio. The Role of Rule-Based Programming for Production Scheduling. Proceedings IEEE International Conference on Robotics and Automation, Raleigh, 1987, pp. 545-550.
21. Burke, P., and P. Prosser. A Distributed Asynchronous System for Predictive and Reactive Scheduling. Tech. Rept. AISL42, Dept. of Computer Science, University of Strathclyde, October, 1989.
22. Canzi, U., G. Guida, W. Poloni, . CRONOS-II: A Knowledge-Based System for Scheduling in a Complex Manufacturing Environment. Politecnico di Milano, Dipartimento di Elettronica, 1988.
23. Chaturvedi, A.R. and G.K. Hutchinson, and D.L. Nazareth. FMS Scheduling Using Goal-Directed Conceptual Aggregation. Proceedings 7th IEEE Conference on AI Applications, Miami Beach, FL, February, 1991, pp. 315-321.
24. Chiang, W.Y., M.S. Fox and P.S. Ow. The OPIS Test Experiments. Tech. Rept. CMU-RI-TR-89-??, The Robotics Institute, Carnegie Mellon University, ??, 1989.
25. Chiang, W.Y., and M.S. Fox. Protection Against Uncertainty in a Deterministic Schedule. Proceedings 4th International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head, May, 1990.
26. Chiodini, V. SCORE: An Integrated System for Dynamic Scheduling and Control of High Volume Manufacturing. Proceedings 5th IEEE Conference on AI Applications, Miami, FL, March, 1989, pp. 272-278.
27. Clancy, D.P. and S. Mohan. RBD: Rule-Based Job Dispatching Software for Implementing Production Plans. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 100-103.
28. Cleveland, G. and S.F. Smith. Using Genetic Algorithms to Schedule Flow Shop Releases. Proceedings 3rd International Conference on Genetic Algorithms, Washington D.C., June, 1989.
29. Collinot, A., C. LePape, and G. Pinoteau. "SONIA: A Knowledge-Based Scheduling System". *Artificial Intelligence in Engineering* 2, 2 (1988).
30. Collinot A., and C. LePape. Controlling Constraint Propagation. Proceedings 10th International Joint Conference on AI, Milano, Italy, August, 1987, pp. 1032-1034.
31. Collinot, A. and C. LePape. Testing and Comparing Reactive Scheduling Strategies. Proceedings 1989 AAAI-SIGMAN Workshop on Manufacturing Production Scheduling, AAAI-SIGMAN, Detroit, MI, August, 1989.
32. Czigler, M.H. and C.R. Whitaker. A Hybrid Algorithmic and Knowledge-Based Implementation for Workcenter-Based Production Scheduling. Proceedings 4th International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1990, pp. 145-155.
33. Davis, L.D. Job Shop Scheduling with Genetic Algorithms. Proceedings 1st International Conference on Genetic Algorithms and their Applications, Pittsburgh, PA, June, 1985.

34. De Waele, W. REDS: A Planning and Scheduling Environment for Manufacturing. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 11-15.
35. Dean, T., Firby, J., and D. Miller. The FORBIN Paper. Yale University, Computer Science Department, 1989.
36. Dean, T. and D. McDermott. "Temporal Data Base Management". *Artificial Intelligence* 32 (1987), 1-55.
37. Dubois, D. Fuzzy Knowledge in an Artificial Intelligence System for Job-Shop Scheduling. In *Applications of Fuzzy Set Methodologies in Industrial Engineering*, Elsevier Science Publishers, 1989, pp. 73-79.
38. Dubois, D. and H. Prade. "Processing Fuzzy Temporal Knowledge". *IEEE Transactions on Man, Systems, and Cybernetics* 19, 4 (1989), 729-744.
39. Elleby, P., H.E. Fargher, and T.R. Addis. Reactive Constraint-Based Job-Shop Scheduling. In *Expert Systems and Intelligent Manufacturing*, M.D. Oliff, Ed., Elsevier Science Publishing Co., 1988, pp. 1-10.
40. Elleby, P., H.E. Fargher, and T.R. Addis. A Constraint-Based Scheduling System for VLSI Wafer Fabrication. In *Knowledge-Based Production Management Systems*, J. Browne, Ed., Elsevier Science Publishers B.V. (North-Holland), 1989, pp. 107-114.
41. Erman, L.D., F. Hayes-Roth, V.R. Lesser and D.R. Reddy. "The Hearsay-II Speech Understanding System: Integrating Knowledge to Resolve Uncertainty". *Computing Surveys* 12 (1980), 213-253.
42. Farah, B.N. "Expert Systems: An Application in Flexible Manufacturing". *Journal of Intelligent and Robotic Systems* (1988), 73-88.
43. Feigenbaum, E., P. McCorduck, and H.P. Nii. *The Rise of the Expert Company*. Random House Inc., New York, 1988.
44. Foo, Y-P.S. and Y. Takefuji. Integer Linear Programming Neural Networks for Job-Shop Scheduling. University of South Carolina, Dept. of Electrical and Computer Engineering, August, 1988.
45. Foo, Y-P.S. and Y. Takefuji. Stochastic Neural Networks for Solving Job-Shop Scheduling: Part 1. Problem Representation. University of South Carolina, Dept. of Electrical and Computer Engineering, August, 1988.
46. Foo, Y-P.S. and Y. Takefuji. Stochastic Neural Networks for Solving Job-Shop Scheduling: Part 2. Architecture and Simulations. University of South Carolina, Dept. of Electrical and Computer Engineering, August, 1988.
47. Fox, M. S., B.P. Allen and G.A. Strohm. Job Shop Scheduling: An Investigation in Constraint-Directed Reasoning. Proceedings of the 2nd National Conference on Artificial Intelligence, American Association of Artificial Intelligence, August, 1982, pp. 155-158.
48. Garey, M.R. and D.S. Johnson. *Computers and Intractability*. W.H. Freeman and Company, San Francisco, 1979.
49. Graves, S.C. "A Review of Production Scheduling". *Operations Research* 29, 4 (July-August 1981), 646-675.
50. Hadavi, K., M.S. Shahraray, and K. Voigt. "REDS - A Dynamic Planning, Scheduling and Control System for Manufacturing". *Journal of Manufacturing Systems* 9, 4 (1990), 332-344.
51. Hadvai, K. and M.S. Shahraray. Release No Job Before Its Time. Proceedings 3rd International Conference on Expert Systems and the Leading Edge in Production and Operations Management, May, 1989.

52. Hastings, N.A.J. and C.-H. Yeh. "Job Oriented Production Scheduling". *European Journal of Operational Research* 47 (1990), 35-48. JOBCODE system.
53. Hilliard, M.R. and G.E. Liepins. Machine Learning Applications to Job Shop Scheduling. Proceedings SIGMAN Workshop on Production Planning and Scheduling, AAAI-SIGMAN, St. Paul, MN, August, 1988.
54. Hsu, W.L., M. Prietula, D. Steier. MERL-SOAR: Scheduling Within a General Architecture for Intelligence. Proceedings 3rd International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1989, pp. 467-481.
55. Husbands, P. and F. Mills. Simulated Co-Evolution as the Mechanism for Emergent Planning and Scheduling. Proceedings 4th International Conference on Genetic Algorithms, San Diego, CA, July, 1991.
56. Hynynen, J. *A Framework for Coordination in Distributed Production Management*. Ph.D. Th., Helsinki University of Technology, Laboratory of Information Processing Science, October 1988.
57. Hynynen, J. BOSS: An Artificially Intelligent System for Distributed Factory Scheduling. Proceedings 3rd International IFIP Conference on Computer Applications in Production and Engineering (CAPE-89), Tokyo, JAPAN, October, 1989.
58. Inoue, I., M. Fuyuki. Job-Shop Scheduling Expert System Based On Know-How Activated Simulation Concept - System Development and Case Study. Proceedings 1989 IIE Integrated Systems Conference & Society for Integrated Manufacturing Conference, Institute of Industrial Engineers, 1989, pp. 483-488.
59. Jacobs, H., J. Kolts, S. Barrett, and M. DeLucia. A Knowledge-Based Manufacturing Scheduling Shell. Proceedings 3rd International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1989, pp. 401-413. simulation-based scheduling.
60. Johnston, M. Spike: AI Scheduling for NASA's Hubble Space Telescope. Proceedings 6th Conference on AI Applications, Santa Barbara, CA, March, 1990, pp. 184-190.
61. Jones, A.T. and C.R. McLean. "A Proposed Hierarchical Control Model for Automated Manufacturing Systems". *Journal of Manufacturing Systems* 5, 1 (1986).
62. Kachitvichyanukul, V., S.-W.J. Cheng, and D.R. Denzler. KBSIM: A Knowledge-Based Simulator for FMS Scheduling and Dispatching. Proceedings Summer Computer Simulation Conference, San Diego, CA, 1986, pp. 696-698.
63. Kanet, J.J. and V. Sridaran. "The Electronic Leitstand: A New Tool for Shop Scheduling". *Manufacturing Review* 3, 3 (Sept. 1990), 161-169.
64. Kempf, K.G. Manufacturing Scheduling: Intelligently Combining Existing Methods. Proceedings AAAI Spring Symposium Series: AI in Manufacturing, Stanford University, March, 1989, pp. 51-55.
65. Kempf, K.G. Manufacturing Planning and Scheduling: Where We Are and Where We Need to Be. Proceedings 5th IEEE Conference on AI Applications, Miami, FL, 1989.
66. Kempf, K., B. Russell, S. Sidhu, and S. Barrett. "AI-Based Schedulers in Manufacturing Practice". *AI Magazine* 11, 5 (January 1991), 46-56.
67. Kempf, K., C. LePape, S.F. Smith, and B. Fox. "Issues in The Design of AI-Based Schedulers". *AI Magazine* 11, 5 (January 1991), 37-46.
68. Keng, N.P., D.Y.Y. Yun, and M. Rossi. Interaction-Sensitive Planning System for Job-Shop Scheduling. In *Expert Systems and Intelligent Manufacturing*, M.D. Oliff, Ed., Elsevier Science Publishing Co., 1988, pp. 57-69.
69. Keng, N.P. and D.Y.Y. Yun. A Planning/Scheduling Methodology for the Constrained Resource Problem. Proceedings AAAI-89, Detroit, Michigan, August, 1989, pp. 998-1003.

70. Kerr, R.M. and R.V. Ebsary. "Implementation of an Expert System for Production Scheduling". *European Journal of Operational Research* 33 (1988), 17-29.
71. Kerr, R.M. and R.N Walker. A Job Shop Scheduling System Based on Fuzzy Arithmetic. Proceedings 3rd International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1989, pp. 433-450.
72. Kim, J., K.H. Funk, and E.F. Fichter. Towards an Expert System for FMS Scheduling: A Knowledge Acquisition Environment. In *Expert Systems and Intelligent Manufacturing*, M.D. Oliff, Ed., Elsevier Science Publishing Co., 1988, pp. 215-234.
73. Kusiak, A. "Artificial Intelligence and Operations Research in Flexible Manufacturing Systems". *Information Systems and Operations Research* 25, 1 (1987), 2-12.
74. Lecocq, P., and T. Guiot. An Expert Systems Application to Increase the Flexibility and the Efficacy of Real-Time FMS Controllers. *Expert Systems and Intelligent Manufacturing*, 1988, pp. 248-264.
75. Lee, J.K. and M.S. Suh. "PAMS: A Domain-Specific Knowledge-Based Parallel Machine Scheduling System". *Expert Systems* 5, 3 (August 1988), 198-213. domain-specific shell for parallel machine scheduling.
76. LePape, C. SOJA: A Daily Workshop Scheduling System. *Expert Systems 85: Proceedings BCS Specialist Group on Expert Systems Conference*, University of Warwick, 1985, pp. 195-212.
77. LePape, C. Constraint Propagation in Planning and Scheduling. Stanford University, Dept. of Computer Science, Dec, 1990.
78. LePape, C. and S.F. Smith. Management of Temporal Constraints for Factory Scheduling. Proc. IFIP TC 8/WG 8.1 Working Conference on Temporal Aspects in Information Systems (TAIS 87), May, 1987, pp. 165-176.
79. Lim, B.S., B.C. Nieuw, and N.C. Ho. Knowledge-Based Master Production Scheduler. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 88-93.
80. Liu, Bing. "Scheduling via Reinforcement". *Artificial Intelligence in Engineering* 3, 2 (1988), 76-85.
81. Liu, B. A Reinforcement Approach to Scheduling. Proceedings 8th European Conference on AI, 1988, pp. 580-585.
82. Lou, X.C., and G. Van Ryzin. "Optimal Control Rules for Scheduling Job Shops". *Annals of Operations Research* 17 (1989), 233-248.
83. Martinez, J., P.R. Muro, M. Silva, S.F. Smith, and J.L. Villarroel. *IMACS Transactions on Scientific Computing-88*. Volume : Merging AI Techniques and Petri Nets for Real Time Scheduling and Control of Production Systems. In *AI and Expert Systems in Scientific Computing*, J.C. Baltzer Publishing, Switzerland, 1989.
84. May, J.H., L.G. Vargas, R. De', S.A. Slotnick, T.E. Morton, D.W. Pentico, and G.J. Tatar. Multex: An Integrated AI/OR System for Factory Floor Scheduling. Tech. Rept. PITT-AIM-22, J.M. Katz Graduate School of Business, University of Pittsburgh, 1990.
85. McKay, K., F. Safayeni, and J. Buzacott. "Job Shop Scheduling Theory: What is Relevant?". *Interfaces* 18, 4 (1988), 84-90.
86. McKay, K., J. Buzacott, and F. Safayeni. The Scheduler's Knowledge of Uncertainty: The Missing Link. Proceedings IFIP Working Conference on Knowledge-Based Production Management Systems, Galway, Ireland, 1988.
87. Meleton, M.P. "OPT: Fantasy or Breakthrough". *Production and Inventory Management* (Second Quarter 1986), 13-21.

88. Meng, C.-C., and M. Sullivan. "LOGOS: A Constraint-Directed Reasoning Shell for Operations Management". *IEEE Expert* 6, 1 (Feb 1991), 20-28.
89. Fox, M.S. *Constraint-Directed Search: A Case Study in Job Shop Scheduling*. Ph.D. Th., Carnegie Mellon University, Computer Science Dept., 1983.
90. Fox, M.S. Observations on the Role of Constraints in Problem Solving. Proceedings 6th Annual Conference of the Canadian Society for Computational Studies of Intelligence, Montreal, Canada, July, 1986.
91. Fox, M.S., N. Sadeh, and C. Baycan. Constrained Heuristic Search. Proceedings IJCAI-89, Detroit, MI, August, 1989, pp. 309-315.
92. Fox, M.S., and S.F. Smith. "ISIS: A Knowledge-Based System for Factory Scheduling". *Expert Systems* 1, 1 (July 1984), 25-49.
93. Fox, M.S. and S.F. Smith. The Role of Intelligent Reactive Processing in Production Management. Proceedings 13th Annual CAM-I Technical Conference, Clearwater Beach, November, 1984.
94. Fox, M.S. and K. Sycara. The CORTES Project: Integrated Production Planning, Scheduling and Control. Proceedings 4th International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head, May, 1990.
95. Minton, S., M.D. Johnston, A.B. Phillips and P. Laird. Solving Large-Scale Constraint Satisfaction and Scheduling Problems Using a Heuristic Repair Method. Proceedings 8th National Conference on Artificial Intelligence (AAAI-88, Boston, MA, July, 1990, pp. 17-24.
96. Morton, T.E., et.al. PATRIARCH: A Multi-Level System for Production Management. Carnegie Mellon University, Graduate School of Industrial Administration, 1987.
97. Muscettola, N. and S.F. Smith. A Probabilistic Framework for Resource-Constrained Multi-Agent Planning. Proceedings 10th International Joint Conference on Artificial Intelligence, Milano, Italy, August, 1987.
98. Muscettola, N. and S.F. Smith. Integrating Planning and Scheduling to Solve Space Mission Scheduling Problems. Proceedings 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control, San Diego, CA, November, 1990.
99. Newman, P.A. and K.G. Kempf. Opportunistic Scheduling for Robotic Machine Tending. Proceedings 2nd IEEE Conference on AI Applications, Miami Beach, FL, 1985, pp. 168-175.
100. Numao, M. and S. Morishita. A Scheduling Environment for Steel-Making Processes. Proceedings 5th IEEE Conference on AI Applications, Miami, FL, March, 1989, pp. 279-286. cooperative scheduling.
101. O'Grady, P. and K.H. Lee. "An Intelligent Cell Control System for Automated Manufacturing". *International Journal of Production Research* 26, 5 (1988), 845-861.
102. Orciuch, E. and J. Frost. ISA: Intelligent Scheduling Assistant. Proceedings 1st IEEE Conference on AI Applications, Denver, CO, December, 1984, pp. 314-320.
103. Ow, P.S. *Heuristic Knowledge and Search in Scheduling*. Ph.D. Th., Carnegie Mellon University, Department of Industrial Administration, 1984.
104. Ow, P.S. "Focused Scheduling in Proportionate Flowshops". *Management Science* 31, 7 (July 1985), 852-869.
105. Ow, P.S. Experiments in Knowledge-Based Scheduling. CMU Robotics Institute Technical Report, forthcoming, 1987.
106. Ow, P.S., S.F. Smith, and R. Howie. CSS: A Cooperative Scheduling System. In *Expert Systems and Intelligent Manufacturing*, M.D. Oliff, Ed., Elsevier Science Publishing Co., 1988.

107. Ow, P.S., S.F. Smith, and A. Thiriez. Reactive Plan Revision. Proceedings AAAI-88, St. Paul, Minn., August, 1988.
108. Ow, P.S. and S.F. Smith. Toward an Opportunistic Scheduling System. Proceedings 19th Hawaii International Conference on System Sciences, January, 1986.
109. Ow, P.S. and S.F. Smith. "Two Design Principles for Knowledge-Based Systems". *Decision Sciences* 18, 3 (Summer 1987).
110. Ow, P.S. and S.F. Smith. "Viewing Scheduling as an Opportunistic Problem Solving Process". *Annals of Operations Research* 12 (1988), 85-108.
111. Panwalker, S.S. and W. Iskander. "A Survey of Scheduling Rules". *Operations Research* 25 (1977), 45-61.
112. Papas, P. The ISIS Project in Review. Proceedings Symposium on Real Time Optimization in Automated Manufacturing Facilities, Gathersburg, MD, 1986.
113. Parunak, H.V.D., B.W. Irish, J. Kindrick and P.W. Lozo. Fractal Actors for Distributed Manufacturing Control. Proceedings 2nd IEEE Conference on AI Applications, December, 1985.
114. Parunak, H.V.D., P.W. Lozo, R. Judd, B.W. Irish, and J. Kindrick. A Distributed Heuristic Strategy for Material Transportation. Proceedings 1986 Conference on Intelligent Systems and Machines, Rochester, Michigan, 1986.
115. Parunak, H.V.D., J.F. White, P.W. Lozo, R. Judd, B.W. Irish, and J. Kindrick. An Architecture for Heuristic Factory Control. Proceedings 1986 American Control Conference, Seattle, WA, 1986.
116. Parunak, H.V.D. A Connectionist Model for Material Handling. Tech. Rept. DFSG 86-22, University of Michigan Industrial Technology Institute, October, 1986.
117. Parunak, H.V.D. Manufacturing Experience with the Contract Net. Tech. Rept. DFSG 86-14, University of Michigan Industrial Technology Institute, Ann Arbor, Michigan, May, 1986.
118. Potvin, J. and S.F. Smith. Flexible Systems for the Design of Heuristic Algorithms in complex OR Domains. Proceedings ORSA Conference on Impact of Recent Computer Advances On Operations Research, Williamsburg, VA, January, 1989.
119. Prosser, P. A Reactive Scheduling Agent. Proceedings AAAI-89, Detroit, Michigan, August, 1989.
120. Prowner, E.T. and D.H. Walburn. A Knowledge-Based Scheduler. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 82-87.
121. Reinschmidt, K.F., J.H. Slater, and G.A. Finn. Expert Systems for Plant Scheduling Using Linear Programming. Proceedings 4th International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head Island, SC, May, 1990, pp. 198-211.
122. Rit, J.F. Propagating Temporal Constraints for Scheduling. Proceedings AAAI-86, Philadelphia, PA, August, 1986, pp. 383-388.
123. Roboam, M., M.S. Fox, and K.P. Sycara. Enterprise Management Network Architecture: Distributed Knowledge Base Support. Tech. Rept. CMU-RI-TR-90-21, The Robotics Institute, Carnegie Mellon University, November, 1990.
124. Roboam, M., M.S. Fox, and K.P. Sycara. Enterprise Management Network Architecture: The Organization Layer. Tech. Rept. CMU-RI-TR-90-22, The Robotics Institute, Carnegie Mellon University, November, 1990.
125. Sadeh, N. *Look-Ahead Techniques for Micro-Opportunistic Job Shop Scheduling*. Ph.D. Th., Carnegie Mellon University, School of Computer Science, May 1991.

126. Sadeh, N. Micro-BOSS: A Micro-opportunistic Factory Scheduler. School of Computer Science, Carnegie Mellon University, January, 1991.
127. Sadeh, N. and M.S. Fox. Preference Propagation in Temporal/Capacity Constraint Graphs. Tech. Rept. CMU-CS-88-193, Carnegie Mellon University, Computer Science Dept., November, 1988.
128. Sadeh, N. and M.S. Fox. Variable and Value Ordering Heuristics for Job Shop Scheduling. Proceedings 4th International Conference on Expert Systems and the Leading Edge in Production and Operations Management, Hilton Head, May, 1990.
129. Safavi, A. Planning and Scheduling of Software Manufacturing Projects. Tech. Rept. CMU-RI-TR-91-04, The Robotics Institute, Carnegie Mellon University, March, 1991.
130. Safavi, A. and S.F. Smith. A Heuristic Search Approach to Scheduling Software Manufacturing Projects. In *Advances in Artificial Intelligence, Natural Language and Knowledge-Based Systems*, M. Golumbic, Ed., Springer-Verlag Publishers, 1990.
131. Sauve, B., and A. Collinot. "An Expert System for Scheduling in a Flexible Manufacturing System". *Robotics and Computer Integrated Manufacturing* 3, 2 (1987), 229-233.
132. Shahraray, M.S. Dynamic Sequencing Rule Applied to Job-Shop Scheduling. Proceedings Simulation and Artificial Intelligence in Manufacturing Conference, SME, Dearborn, MI, 1987.
133. Shaw, M.J. "Knowledge-Based Scheduling in Flexible Manufacturing Systems: An Integration of Pattern-Directed Inference and Heuristic Search". *International Journal of Production Research* 26, 5 (1988), 821-844.
134. Shaw, M.J. "A Pattern-Directed Approach to Flexible Manufacturing: A Framework for Intelligent Scheduling, Learning and Control". *International Journal of Flexible Manufacturing Systems* 2 (1989), 121-144.
135. Shaw, M. and A.B. Whinston. Automated Planning and Flexible Scheduling: A Knowledge-based Approach. Proceedings IEEE International Conference on Robotics and Automation, St. Louis, MO, 1985, pp. 890-894.
136. Smith, R.G. "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver". *IEEE Transactions on Computers* C-29, 12 (December 1980), 1104-1113.
137. Smith, S.F. Exploiting Temporal Knowledge to Organize Constraints. Tech. Rept. CMU-RI-TR-83-12, CMU Robotics Institute, July, 1983.
138. Smith, S.F. A Constraint-Based Framework for Reactive Management of Factory Schedules. in *Intelligent Manufacturing*, M.D. Oliff, Ed., Benjamin Cummings Publishers, 1987.
139. Smith, S.F. The OPIS Framework for Modeling Manufacturing Systems. Tech. Rept. CMU-RI-TR-89-30, CMU Robotics Institute, December, 1989.
140. Smith, S.F., P.S. Ow, J.Y. Potvin, N. Muscettola, and D. Matthys. "An Integrated Framework for Generating and Revising Factory Schedules". *Journal of the Operational Research Society* 41, 6 (June 1990), 539-552.
141. Smith, S.F., N. Keng, and K. Kempf. Exploiting Local Flexibility During Execution of Pre-Computed Schedules. Tech. Rept. CMU-TR-RI-90-13, The Robotics Institute, Carnegie Mellon University, June, 1990.
142. Smith, S.F., P.S. Ow, N. Muscettola, J.Y. Potvin and D. Matthys. OPIS: An Opportunistic Factory Scheduling System. Proceedings Third International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE-90), Knoxville, TE, May, 1990.
143. Smith, S.F., P.S. Ow, C. LePape, B. McLaren, and N. Muscettola. Integrating Multiple Scheduling Perspectives to Generate Detailed Production Plans. Proceedings SME Conference on AI in Manufacturing, Long Beach, CA, September, 1986.

144. Smith, S.F., M.S. Fox, and P.S. Ow. "Constructing and Maintaining Detailed Production Plans: Investigations into the Development of Knowledge-Based Factory Scheduling Systems". *AI Magazine* 7, 4 (Fall 1986).
145. Smith, S.F. and J.E. Hynynen. Integrated Decentralization of Production Management: An Approach for Factory Scheduling. Proceedings ASME Annual Winter Conference: Symposium on Integrated and Intelligent Manufacturing, Boston, MA, December, 1987.
146. Smith, S.F. and P.S. Ow. The Use of Multiple Problem Decompositions in Time-Constrained Planning Tasks. Proceedings 9th International Joint Conference on Artificial Intelligence, Los Angeles, California, August, 1985.
147. Steffen, M.S. A Survey of Artificial Intelligence-Based Scheduling Systems. Proceedings Fall Industrial Engineering Conference, Institute of Industrial Engineers, 1986, pp. 395-405.
148. Steffen, M.S. and T.J. Greene. An Application of Hierarchical Planning and Constraint-Directed Search to Scheduling Parallel Processors. Proceedings IEEE International Conference on Robotics and Automation, San Francisco, CA, 1986.
149. Sullivan, G. and K. Fordyce. Problem Solving Technologies for Manufacturing Logistics. Proceedings AAAI Spring Symposium Series: AI in Manufacturing, Stanford University, March, 1989, pp. 98-101.
150. Sycara, K.P., S.F. Roth, N. Sadeh, and M.S. Fox. "Distributed Constained Heuristic Search". *IEEE Transactions on Man, Systems, and Cybernetics* 21, 6 (December 1991).
151. Sycara, K.P., S.F. Roth, N. Sadeh, and M.S. Fox. "Resource Allocation in Distributed Factory Scheduling". *IEEE Expert* 6, 1 (Feb. 1991), 29-40.
152. Syswerda, G. Schedule Optimization Using Genetic Algorithms. In *Handbook of Genetic Algorithms*, L.D. Davis, Ed., Addison Wesley, 1990, ch. 10.
153. Tse, E., M. Zhu and S-Y. Juang. SIMENS: An Intelligent Support System for Manufacturing Enterprise. Proceedings 1990 IEEE Conference, 1990, pp. 282-289.
154. Van Hentenryck, P. Application of CHIP to Industrial Scheduling Problems. Proceedings 1989 AAAI-SIGMAN Workshop on Manufacturing Production Scheduling, AAAI-SIGMAN, Detroit, MI, August, 1989.
155. Verbraeck, A. Adaptive Production Planning Environments Using Decision Rules and Simulation. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 147-151.
156. Warwick, A.M., and H.M.J. Walters. A Rule-Based Planning and Scheduling System for Manufacturing Industries. Proceedings International Conference on Expert Planning Systems, Institution of Electrical Engineers, Brighton, United Kingdom, June, 1990, pp. 104-109.
157. Whitley, D., T. Starkweather, and D. Fuquay. Scheduling Problems and the Traveling Salesman: The Genetic Edge Recombination Operator. Proceedings 3rd International Conference on Genetic Algorithms, Washington D.C., June, 1989.
158. Wilkins, D. Can AI Planners Solve Practical Problems? SRI, 1990.
159. Yu C., Scott G & Kempf K.G. Artificial Intelligence and the Scheduling of Semiconductor Wafer Fabrication Facilities. Proceedings of the Intel Technology Conference, Intel Corporation, 1988, pp. 135-138.
160. Zweben, M., M. Deale, and R. Gargan. Anytime Rescheduling. Proceedings 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control, San Diego, CA, November, 1990, pp. 251-259.