AD-A158 721 THE USE OF MULTIPLE PROBLEM DECOMPOSITIONS IN TIME CONSTRAINED PLANNING T..(U) CARNEGIE INST OF TECH PITTSBURGH PA GRADUATE SCHOOL OF INDUST.. 1/1

UNCLASSIFIED S F SMITH ET AL. 06 MAY 85 AFOSR-TR-85-0637 F/G 6/4 NL

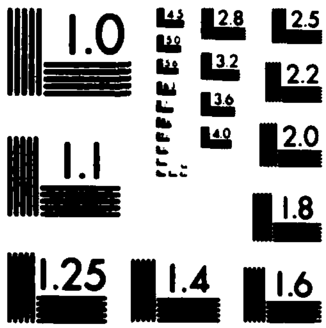

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AFOSR-TR- 85-0637

The Use of Multiple Problem Decompositions in Time Constrained Planning Tasks

Stephen F. Smith

Intelligent Systems Laboratory
The Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213 USA

Peng Si Ow

Grad. School of Ind. Admin.
Carnegie-Mellon University
Pittsburgh, PA 15213 USA

6 May 1985

To appear in *Proceedings 9th International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August, 1985.

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR- 85-0637
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b. OFFICE SYMBOL (If applicable)
7a. NAME OF MONITORING ORGANIZATION Air Force Office of Scientific Research	
6c. ADDRESS (City, State and ZIP Code) Pittsburgh, PA 15213	7b. ADDRESS (City, State and ZIP Code) Directorate of Mathematical & Information Sciences, Bolling AFB DC 20332
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (If applicable) NM
9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F49620-82-K-0017	
8c. ADDRESS (City, State and ZIP Code) Bolling AFB DC 20332	10. SOURCE OF FUNDING NOS.

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.
61102F	2304	A7	

11. TITLE (Include Security Classification) The Use of Multiple Problem Decompositions in Time Constrained Planning Tasks

12. PERSONAL AUTHOR(S) Stephen F. Smith

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Yr., Mo., Day)	15. PAGE COUNT
Reprint	FROM ______ TO ______	6 May 1985	10

16. SUPPLEMENTARY NOTATION
9th International Joint Conference on Artificial Intelligence, Los Angeles, CA Aug, 1985

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR.	multiple problem decompositions resource-based and agent-based perspectives
XXXXXXXXXXXXXXXX			

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Problems requiring the synthesis of a collection of plans accomplishing distinct (but possibly related) goals has received increasing attention within AI. Such problems are typically formulated as multi-agent planning problems, emphasizing a problem decomposition wherein individual agents assume responsibility for the generation of individual plans while taking into account the goals and beliefs of other agents in the system. One consequence of such a problem decomposition is a simplified view of resource allocation that assumes avoidance of conflicts to be the sole concern. The validity of this assumption comes into question in time constrained problem domains requiring the allocation of multiple, shared resources. In job shop scheduling, for example, where sequences of manufacturing operations must be determined and scheduled for multiple orders, it is necessary to consider much more than availability to efficiently allocate resources over time. We argue that in such domains, an ability to reason from both resource-based and agent-based perspectives is essential to appropriate consideration of all domain constraints.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS ☐	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL Robert N. Buchal	22b. TELEPHONE NUMBER (Include Area Code) (202) 767-4939	22c. OFFICE SYMBOL NM

DD FORM 1473, 83 APR   EDITION OF 1 JAN 73 IS OBSOLETE   UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

Abstract

Problems requiring the synthesis of a collection of plans accomplishing distinct (but possibly related) goals has received increasing attention within AI. Such problems are typically formulated as multi-agent planning problems, emphasizing a problem decomposition wherein individual agents assume responsibility for the generation of individual plans while taking into account the goals and beliefs of other agents in the system. One consequence of such a problem decomposition is a simplified view of resource allocation that assumes avoidance of conflicts to be the sole concern. The validity of this assumption comes into question in time constrained problem domains requiring the allocation of multiple, shared resources. In job shop scheduling, for example, where sequences of manufacturing operations must be determined and scheduled for multiple orders, it is necessary to consider much more than availability to efficiently allocate resources over time. We argue that in such domains, an ability to reason from both resource-based and agent-based perspectives is essential to appropriate consideration of all domain constraints.

Table of Contents

1. Multi-Agent Planning and Resource Allocation

Problems requiring the synthesis of a collection of plans accomplishing distinct (but possibly related) goals has received increasing attention within AI. Systems that address this problem have been called multi-agent planning systems [5], so termed because of the emphasis on a system of loosely-coupled, cooperative planning agents, each responsible for the generation of a single plan but cognizant of, and taking into account, the goals and beliefs of other agents in the system. Work in this area has focused primarily on the issue of goal protection, i.e. the planning of activities to achieve desired goals despite the dynamic nature of the surrounding environment. The issue of allocating resources to such activities has typically been given secondary importance, the assumption being that the avoidance of resource conflicts is the sole concern. The validity of this assumption comes into question in time constrained domains requiring the allocation of multiple, shared resources. In job shop scheduling, for example, where sequences of operations must be determined and scheduled for multiple orders, resource (e.g. machine) assignments to support these operations are influenced by much more than the mere availability of the resource during the time period in question. Other constraints, such as capacity limitations, sequencing preferences, and order splitting preferences, must also be considered. Efficient allocation of resources under such constraints is difficult within the problem decomposition ascribed to multi-agent planning systems above, given the local and incomplete view of resources held by each individual agent. Our experience with the ISIS job shop scheduling system [4], which adopts such a problem decomposition, confirms this claim. What is needed to directly exploit these types of constraints is an ability to reason from a resource based perspective, suggesting the use of multiple problem decompositions. This paper describes an initial approach to providing such a reasoning capability.

The remainder of the paper is organized as follows. In Section 2 we briefly review the approach taken by ISIS in reasoning with the large and conflicting set of constraints encountered in the job shop scheduling domain. This is followed in Section 3 by a closer examination of the limitations of decomposing the problem solely from a order (or agent) based perspective, and issues surrounding the integration of a resource based reasoning capability are explored. An initial system architecture possessing an ability to reason from both perspectives is presented. Finally, in Section 4, some research directions are identified.

2. Constraint-Directed Reasoning in ISIS

The scheduling domain of ISIS is realistically complex, requiring the consideration of such diverse and conflicting factors as due date requirements, cost restrictions, production levels, machine capabilities, alternative production processes, order characteristics, resource characteristics and resource availability. To address this complexity, the ISIS design advocates two key ideas:

- an *explicit formalization* of the various scheduling influences as constraints in the system's knowledge base, and

- the formulation of schedule construction as a *constraint-directed* heuristic search.

The first point above presumes a fairly broad view of constraints, and it is important to note that the ISIS constraint representation encompasses scheduling objectives, goals, and preferences as well as the range of necessary conditions that delineate the space of admissible schedules. Constraints of the former variety provide a basis for optimization during the evaluation of alternative solutions, by assigning utilities indicative of the degree to which they have been satisfied. The representation also captures other knowledge necessary to effectively reason with the constraints, including constraint importance, constraint relevance, and constraint interdependencies.

The generation of a shop schedule is accomplished in an incremental fashion. For each order to be scheduled, the system proceeds through multiple levels of analysis, principal of which is the heuristic search procedure employed to make detailed selections of operations, resources, and time intervals for production of the order. Working from a set of allowable routings for the order (i.e. a directed graph of operations capturing operation precedence constraints, alternative manufacturing processes and resource substitutability), the search proceeds either forward from the order's requested start date or backward from the requested ship date. The search space is composed of states that represent alternative partial schedules, and the application of the search operators serves to generate new states that further specify the partial schedules under development (e.g. add another operation to a partial schedule for the order, bind a particular machine to an operation, allocate a particular time interval for the order on a particular machine). Using a beam search, only the best *n* search states are extended at each iteration of the search, and, as indicated above, the quality of a given state is estimated on the basis of how well it satisfies the objectives, goals, and preferences that are relevant to the scheduling decisions it embodies. The outcome of this search is a particular routing for the order along with an assignment of time bounds to the resources required to produce it. Once refined into the order's final schedule, these commitments serve to additionally constrain any subsequent scheduling that must be performed.[1]

3. Reasoning from Multiple Perspectives

The ISIS heuristic search paradigm outlined above attempts to provide a framework for incorporating the full range of constraints that typically influence human scheduler's decisions in the automatic construction of job shop schedules. Unfortunately, its commitment to a particular decomposition of the scheduling problem places undue emphasis on the exploitation of a certain class of constraints to the effective exclusion of others. The "focus on one order at a time" approach employed, while useful in reducing the overall complexity of the problem, does not provide an adequate basis for attending to *inter-order constraints*, i.e. constraints that influence the allocation of resources over a number of orders. One example of how this weakness manifests itself is in the consideration of order sequencing preferences, which arise due to machine setup costs[2]. These preferences relate to the total set of orders requiring a given resource in the shop. Yet, within the ISIS framework, they can be considered only in the context of the partially constructed shop schedule that exists at the time each order is selected for scheduling. As such, the extent of their influence is somewhat coincidental.

We next address the problem of how to cope with these inter-order constraints while preserving the ability to adequately exploit intra-order constraints. Specifically, we explore the use of a resource-

---

[1] The above description is necessarily brief and omits several important issues that are not directly relevant to the discussion below. The reader is referred to [2, 3, 4, 7] for more detailed accounts of this work.

[2] It is sometimes necessary to prepare or "set up" a machine before an operation can be performed on it. This is typically caused by a change in the type of operation to be performed on the machine and the amount of machine time consumed for setups is dependent on the type of change. A proper sequencing of orders on a machine can greatly increase its throughput.

based decomposition in conjunction with an order-based decomposition to provide this added ability.

3.1. Partitioning the Problem Solving Effort

The primary source of difficulty in constructing good job shop schedules stems from the conflicting nature of the domain's constraints. Constraints are said to be in conflict when a scheduling decision made with respect to satisfying any one affects the extent to which the others may be satisfied. The optimal resolution of a given conflict necessarily requires a problem decomposition in which all constraints involved in the conflict are grouped within the same subproblem. The order-based decomposition utilized by ISIS groups together the constraints surrounding a particular order and, consequently, provides an opportunity for effectively resolving order-centered conflicts (e.g. conflicts involving precedence constraints). A resource-based problem decomposition, in contrast, produces a grouping of constraints that promotes the resolution of a different set of conflicts. Here the strategy becomes one of scheduling on a resource by resource basis, and the grouping of constraints contained within a given subproblem includes a cross-section of the constraints associated with a number of orders. Conflicts brought to the foreground under this decomposition center around the resource allocation decisions that must be made at a particular resource (e.g. conflicts involving various setup preferences). It is obvious that there are many conflicts that cannot be effectively isolated within either decomposition strategy, and, consequently, conflicts that cannot be optimally resolved from either problem solving perspective. Nonetheless, it is felt that a broadening of the range of constraints that can be meaningfully addressed through the use of multiple perspectives will lead to a more equitable consideration of the domain's constraints.

Given the decision to employ multiple problem decompositions, the task becomes one of how to best partition the problem solving effort between distinct perspectives so as to maximize the number of conflicts that can be directly addressed. Since the formation of specific subproblems (e.g. schedule operations on the milling machine from a resource-based perspective, and schedule the other operations on an order by order basis) will determine which conflicts can be directly resolved, a partitioning that associates essential resource-based conflicts with the resource-based reasoning component and, likewise, essential order-based conflicts with the order-based reasoning component, is highly desirable. Fortunately, the majority of resource-based conflicts can be identified through the detection of *bottleneck* resources, so called because they are scarce resources of the shop. Accordingly, a division of effort in which the resource-based reasoning component is employed to make resource allocation decisions at the bottleneck resources and non-bottleneck resources are scheduled from an order-based perspective is seen as most appropriate.

Despite an ability to derive a fairly useful problem decomposition, interactions amongst subproblems remain an important concern. Resource allocation decisions made with respect to a particular bottleneck, for example, might quite likely limit the extent to which we can effectively resolve conflicts (or satisfy constraints) in subsequent order-based subproblems. It is felt that the harmful effects of these interactions can be minimized somewhat by a judicious ordering the subproblems identified. Specifically, the relative importance of satisfying various constraints is seen as a useful criterion for coordinating the overall effort.

3.2. A Specific System Architecture

To gain a better understanding of these issues, a specific system architecture possessing both resource-based and order-based reasoning components has been implemented. Adopting the problem decomposition strategy discussed in Section 3.1 (i.e. that resource-based reasoning is most critical with respect to bottleneck resources), a simple scheme for coordinating the overall effort has been imposed. Specifically, the system first employs its resource-based reasoning component to establish resource reservations at the bottleneck resources. These resource allocation decisions, which are guaranteed to be feasible, then serve as "islands of certainty" [1] for subsequent exploitation by the order-based reasoning component in developing the remainder of the schedule. The resource based reasoning strategy currently employed is based on a particular OR (Operations Research) heuristic developed in [6]. The order-based reasoning component is a derivative of the strategy employed by ISIS that has been generalized to operate on arbitrary portions of the set of routings associated with a particular order.

Testing of the system is currently proceeding, and preliminary results using simulated plant data[3] appear promising. Moreover, the reconfiguration of the ISIS scheduling system to treat previously imposed resource reservations as fixed points or islands from which to expand the search has also provided an opportunity to employ the user as the resource-based reasoning component of the system. In experiments where we have manually scheduled bottleneck resources before invoking the order-based reasoning component, the benefits of providing an ability to reason directly about the critical resource allocation decisions that have to be made can be immediately seen, and give considerable credence to the approach we have adopted.

4. Discussion

In this paper we have pointed up the inadequacy of the problem decomposition typically embodied by multi-agent planning systems for certain classes of problems involving the synthesis of multiple plans. In particular, we argue that problems requiring the efficient allocation of a collection of shared resources over time cannot be effectively addressed by relying solely on the incomplete and local knowledge possessed by individual agents attending to the construction of individual plans. The specific focus of our work, the generation of schedules to govern production in a job shop, exemplifies this type of problem wherein efficient allocation of resources requires direct consideration of inter-order constraints, and, hence, an ability to reason globally from a resource-based perspective. Recognizing this, we have proposed a problem solving framework that employs both resource-based and order-based decompositions of the scheduling problem. We have suggested that the division of effort between these distinct perspectives can be usefully guided by an attempt the maximize the range of constraint conflicts that can be meaningfully resolved, and that the harmful effects of the interactions amongst the resulting subproblems can be minimized by an appropriate prioritization of the domain's constraints. A initial system architecture was presented to demonstrate the feasibility of this approach.

The work reported here has only begun to address the larger issues associated with the use of

---

[3] The specific job shop model employed is based on the Westinghouse Turbine Components Plant located in Winston-Salem, North Carolina. The reader is referred to [4] for a discussion of the characteristics and complexities of this specific scheduling environment.

multiple problem decompositions in balancing a large set of conflicting objectives. We have made specific assumptions about the relative importance of satisfying various constraints which have led to a static partitioning of the problem solving effort (*i.e. first schedule bottlenecks from a resource based* perspective and then schedule non-bottlenecks using an order-based decomposition). While this, in general, might constitute a reasonable guideline for prioritizing the domain's constraints, there are obviously situations in which important constraints will not be appropriately attended to. The improvement of matters requires a more dynamic interplay between the reasoning strategies associated with distinct problem solving perspectives in which decisions made while reasoning from a particular perspective can be questioned and undone in light of the constraints that become relevant as different perspectives are employed. This is an issue that we are currently pursuing.

A second important issue concerns the level of sophistication of the resource-based reasoning strategy. In adopting an OR heuristic as the sole basis for decision making, we have limited the system's attention to the constraint knowledge that is implicitly captured by the heuristic. While the specific heuristic employed does, in fact, attend to several important inter-order constraints, it nonetheless operates with a restrictive model of the scheduling environment. Ultimately, the reasoning strategy must be capable of exploiting any constraint found to be relevant to the resource allocation decisions under consideration. This implies a strategy that reasons with an explicit characterization of constraint knowledge, analogous to the heuristic strategy currently employed by the system when reasoning from an order-based perspective. The development of such a strategy is also currently under investigation.

Acknowledgements

The ideas presented in this paper have benefited greatly from many discussions with Mark Fox. This research was supported in part by the Air Force Office of Scientific Research under contract F49620-82-K-0017, Westinghouse Electric Corporation, and the Robotics Institute.

References

1. Erman, L.D., Hayes-Roth, F., Lesser, V.R., and Reddy, D.R. "The HEARSAY-II Speech Understanding System: Integrating Knowledge to Resolve Uncertainty." *Computing Surveys 12*, 2 (June 1980).

2. 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.

3. Fox, M.S. *Constraint-Directed Search: A Case Study of Job Shop Scheduling*. Ph.D. Th., Carnegie-Mellon University, 1983.

4. Fox, M.S. and Smith, S.F. "ISIS: A knowledge-Based System for Factory Scheduling." *Expert Systems 1*, 1 (July 1984), 25-49.

5. Konolidge, K., N.J. Nilsson. Multi-Agent Planning Systems. Proceedings 1st Conference of the American Association for Artificial Intelligence, August, 1980.

6. Ow, P.S. *Heuristic Knowledge and Search for Scheduling*. Ph.D. Th., Graduate School of Industrial Administration, Carnegie-Mellon University, 1984.

7. Smith, S.F. Exploiting Temporal Knowledge to Organize Constraints. Tech. Rept. CMU-RI-TR-83-12, Robotics Institute, Carnegie-Mellon University, July, 1983.

END

FILMED

10-85

DTIC