

530-63

N 9 8 7 1 8 6 8 9

P-4

JIGSAW: Preference-Directed, Co-operative Scheduling

Theodore A. Linden
David Gaw

Advanced Decision Systems,
a Division of Booz•Allen & Hamilton, Inc.
1500 Plymouth St.
Mountain View, CA 94043
linden@ads.com
415-960-7562

Abstract*

Techniques that enable humans and machines to cooperate in the solution of complex scheduling problems have evolved out of work on the daily allocation and scheduling of Tactical Air Force resources. A generalized, formal model of these applied techniques is being developed. It is called JIGSAW by analogy with the multi-agent, constructive process used when solving jigsaw puzzles. JIGSAW begins from this analogy and extends it by propagating local preferences into global statistics that dynamically influence the value and variable ordering decisions. The statistical projections also apply to abstract resources and time periods—allowing more opportunities to find a successful variable ordering by reserving abstract resources and deferring the choice of a specific resource or time period.

Keywords: Scheduling, constraint propagation, statistical look-ahead, hierarchical planning, resource abstractions, transformational synthesis.

1. Introduction

For many scheduling problems, partial automation is a realistic but difficult goal. Partial automation means that human schedulers can participate in incremental scheduling decisions. Algorithms from operations research and most heuristic search techniques involve humans in the problem set up but not in the generation of schedules. These algorithms work well when the problem is modeled perfectly and is

computationally tractable. Unfortunately, practical scheduling problems occur in very complex environments, it is usually impossible to capture all of the domain complexities in the formal model. In practice, the results of fully automated scheduling algorithms are used primarily to debug the problem set up. This results in a very large debugging loop that is inefficient and does not always converge to an acceptable solution. Furthermore, details about myriads of individual preferences are seldom handled effectively. While a human scheduler may notice that an important task in today's schedule is one on which John Jones performed effectively last week, it is impractical to expect that the knowledge acquisition task can capture all these subtle preferences in advance.

A co-operative approach to schedule generation exploits the strengths of both humans and automation, but co-operation implies that the scheduling software has to work with an incomplete model of the problem domain. Human scheduling decisions should be viewed as dynamic extensions to that model. Furthermore, many scheduling problems are dominated by preferences rather than by hard constraints, and these preferences need to be exploited in the same way that constraints are exploited in constraint-directed scheduling.

2. Background and Overview

JIGSAW generalizes techniques originally developed to partially automate the daily allocation and scheduling of Tactical Air Force resources. The complexity of the knowledge involved in this scheduling problem is such that, when done manually, a team of 8-16 people works over a period of 12 or more hours. An interactive system that solves this problem by allowing humans and the machine to make incremental scheduling decisions was designed three

* This work was partially supported by the Defense Advanced Research Projects Agency (DARPA) under contract DAAH01-90-0080 and partially supported by IR&D funding from Advanced Decision Systems.

years ago, has undergone two years of user evaluations,¹ and is now being hardened for operational use. JIGSAW is a generalization and formalization of the automated reasoning techniques originally designed for this application.

JIGSAW is an open collection of techniques that allow humans to participate as schedules are constructed incrementally. JIGSAW begins with a transformational approach—similar to the transformations commonly used to compile program specifications into programs and to refine design specifications into designs. Correctness-preserving transformations encapsulate knowledge about incremental allocation and scheduling decisions. They separate the definition of these decision rules from the control decisions about when they should be invoked.

JIGSAW extends this transformational approach with statistical look-ahead techniques. Statistical look-ahead uses local constraints and preferences to project the expected contention for resources over time. These statistical projections allow local scheduling decisions to be influenced by statistical knowledge about the global context. Statistical look-ahead enhances both value and variable ordering techniques. Our ongoing work extends these statistical projections to deal with abstract resource groupings. Partially determined time intervals are also handled as abstract resources. An assignment of an abstract resource to a task creates a reservation for an unspecified instance of the abstract resource. These reservations for abstract resources enable incremental commitments that provide more opportunities to find variable orderings that avoid or reduce backtracking.

The name JIGSAW is based on an analogy with jigsaw puzzles where:

- **Many independent agents**—both human and automated—co-operate to construct a solution.
- **The order in which scheduling decisions are made is not predetermined** by the problem.
- **Partial solutions** can (usually) be evaluated as (probably) extensible to an acceptable solution.

JIGSAW extends this analogy with a combination of techniques for reasoning about preferences, abstraction levels, variable ordering, and uncertainty. Unlike jigsaw puzzles, JIGSAW seeks a globally good solution by making a series of local decisions that are

informed by statistical knowledge about how the local decision is likely to impact global optimality.

The overall JIGSAW approach involves associating a transformation with each incremental, atomic allocation and scheduling decision. The users can commit some scheduling decisions, and the automated JIGSAW techniques accept and work with partial schedules developed by users. The users control which transformations will be candidates for execution. The control software invokes the transformations that produce the most promising extensions of the current partial schedule.

3. Exploiting Value and Variable Ordering Opportunities

To fully exploit value and variable ordering opportunities when constructing a schedule incrementally, individual transformations of partial assignments should be kept as atomic as possible. Most job shop scheduling techniques exploit variable ordering opportunities only at the level of complete orders or resources; that is, they make assignments to all the tasks involved in an order or they completely schedule a single resource. Like Cortes [Fox & Sycara 90, Sadeh 91], JIGSAW enables separate decisions for each individual task or activity.² JIGSAW allows a task to be assigned a resource or scheduled into a time period without simultaneously committing to decisions about other tasks or times. Furthermore, by introducing resource abstraction hierarchies, JIGSAW can reserve an abstract resource for a task while deferring the assignment of a specific resource or time interval. These assignments of abstract resources allow more opportunities for variable ordering heuristics to be effective.

When allowing very small incremental transformations that may be made in almost any order, one has a problem preserving the property that any partial assignment that is generated can be extended to a nearly optimal solution. In particular, related tasks must all eventually receive consistent assignments, tasks that are assigned an abstract resource must eventually receive specific resources, tasks that receive resources must eventually be scheduled, and tasks assigned a flexible time period must eventually be scheduled into a specific time interval. These problems are largely avoided in earlier scheduling systems where all of the decisions associated with an order or resource are made simultaneously; however,

¹The realities of a large implementation have led to an early focus on machine assistance for human decision-making; implementation of the automated decision-making techniques on which JIGSAW is based is quite recent.

²JIGSAW's tasks are equivalent to Cortes' activities. The terms "operation" and "variable" are also used in the literature with an equivalent meaning.

this grouping of decisions limits the opportunities to fully exploit value and variable orderings.

JIGSAW includes substantial bookkeeping functions and statistics that summarize the state of current assignments and project the probable effects of future assignments. This information is used to inhibit transformations that are likely to interfere with the completion of existing partial assignments. Projections about the expected demand on resources allow the incremental transformations to achieve a balance between greedy local optimization and altruistic minimization of resource conflicts [Sycara et al. 90]. The bookkeeping functions and statistics apply to abstract as well as specific resources and time periods. Reservation for abstract resources are guaranteed in that transformations making assignments to other tasks will preserve enough instances of the abstract resource to fulfill all prior reservations. Significantly, many of these same bookkeeping functions and statistics are also useful to the human experts who co-operate in the problem solving process.

The bookkeeping functions and statistics are also used to dynamically select the order in which the transformations are executed. The goal is to defer each transformation until there is enough information available to predict that its decision is a step toward achieving a nearly optimal assignment. Note that if all transformations meet this goal, then whenever a specific task-resource or task-time-period pairing is required to achieve an optimal assignment, other transformations will not use up the last instance of the resource or time period that is needed by this task. Of course, with invocation criteria as stringent as this, the problem is whether there will always be a transformation that does not need to be deferred. An experimental hypothesis being evaluated is: for many large problems that are characterized by many preferences and that can be solved adequately by teams of human experts, there will usually be some "obvious" transformation that does not need to be deferred. When there is no such transformation, then either human intervention or a branch and bound search strategy can be used effectively.

In summary, the automated portion of JIGSAW starts from any consistent partial assignment (initially from the empty assignment unless human experts make some initial decisions), finds a transformation that is statistically the least in need of being deferred, executes that transformation, and iterates. Humans control the overall process and can interleave their own decisions between transaction invocations.

4. Statistical Projections

In the Tactical Air Force application, statistical look-ahead was used to give more sophistication to what is

basically a greedy algorithm augmented with plan repair techniques. However, the statistical look-ahead techniques together with reservations for abstract resources also work in the context of backtracking or breadth-first search strategies. The choice of the search strategy is controlled by the size of the problem and the need to interact with human schedulers, not by the statistical look-ahead. Human schedulers appear to be most comfortable with divide-and-conquer, greedy algorithms, and plan repair strategies—together with a very limited amount of breadth-first search and backtracking.

The critical part of JIGSAW is the inner loop where statistics about expected resource availability are projected and a transformation that does not need to be deferred is found. This section summarizes the steps used in the Tactical Air Force scheduling problem from which JIGSAW evolved. A more formal, general treatment can be found in [Linden 91], and we are currently trying to formalize these ideas more directly in terms of Bayesian networks and decision theory.

This description of the inner loop in JIGSAW is a step toward generalizing the computations, not optimizing them. The Air Force application where these techniques were applied deals with the optimization issues; many optimizations are available by reusing previous computations.

The steps of the inner loop are:

1. **Local rating:** Use constraints to identify the alternative resources and time periods that can be assigned to each task, and use preferences to order or rate these possible values. This local rating is based on the easily-processed constraints and preferences directly associated with the task; initially, it does not deal with global issues like resource availability.
2. **Global statistics:** Translate the local ratings for each alternative value assignment into a subjective probability that this assignment will be made, and sum these probabilities across all the tasks to project global statistics about the expected demand for each resource. Comparison of the expected demand for resources with the available resources identifies probable bottlenecks.
3. **Trade off:** Re-evaluate the alternative value assignments in terms of which choice is most likely to be part of a globally optimal assignment. This re-evaluation uses the statistics about resource contention and makes a trade off between local utility and global resource contention.
4. **Commit:** For one or more tasks, "commit" to a transformation that is projected to lead toward a good complete assignment. Choose to make this

commitment for tasks where the decision is "obvious" and/or "influential":

a. **Obvious decisions** are those where one can project a very high confidence level that a decision made now will be "right." This confidence is evaluated in terms of:

- Strength of the local preference for the proposed commitment relative to alternative possible values. This may be computed as the delta between the rating of the value to be committed and the rating of the next best value.
- The commitment's use of low contention resources based on the statistical projections of the expected demand for each resource at various times.
- The quality of the current understanding about how interactions with other tasks might affect this task.

b. **Influential decisions** are decisions which clarify many other decisions; for example, a decision to commit bottleneck resources is influential because it narrows the choices that remain open for all others decisions.

5) **Plan repair:** Plan critics are available as a way of undoing a previous decision—along with the decisions that directly depend on it. Plan critics resolve conflicts that arise from imperfect look-ahead or from changing conditions in the external environment. Plan critics have been included in the design of JIGSAW applications, but they have not yet been added to the formal JIGSAW model.

5. Conclusions

JIGSAW evolved from work on large scheduling applications that must be solved co-operatively and are dominated by preferences rather than by hard constraints. JIGSAW exploits those preferences to project statistical characteristics of the global situation which are then used to enhance local value and variable ordering decisions. JIGSAW extends these statistical projections to abstract groupings of resources and allows partial schedules to include reservations for abstract resources. These reservations for abstract resources open more opportunities for value and variable ordering techniques to be effective.

JIGSAW is proposed as one of a range of scheduling techniques. It is appropriate for large resource allocation and scheduling applications that are currently solved by teams of human experts. It is especially appropriate for problems where the evaluation criteria are complex, changing, and not fully formalized--problems for which human schedulers

need to be involved to help evaluate the feasibility and effectiveness of the evolving schedules.

6. References

- [Fox & Sycara 90] Mark S. Fox and Katia P. Sycara, "Overview of CORTES: A Constraint Based Approach to Production Planning, Scheduling, and Control," Proc. of the Fourth Inter. Conf. on Expert Systems in Production and Operations Management, May, 1990.
- [Linden 91] Theodore A. Linden, "Preference-directed, Co-operative Resource Allocation and Scheduling." Final Technical Report, DARPA Order No. 6685, Advanced Decision Systems Report TR-1270-3, Sept. 1991.
- [Sadeh 91] Norman Sadeh, "Look-ahead Techniques for Micro-opportunistic Job Shop Scheduling." PhD Thesis, School of Computer Science, Carnegie Mellon University, 1991.
- [Sycara et al. 90] Katia P. Sycara, S. Roth, N. Sadeh, and M. Fox, "Managing Resource Allocation in Multi-Agent Time-constrained Domains." DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control., Morgan Kaufman Publ., San Mateo, CA, Nov. 1990, pp. 240-250.