

227-63  
137253  
N93-18686

# IOPS Advisor: Research in Progress on Knowledge-Intensive Methods for Irregular Operations Airline Scheduling\*

John E. Borse and Christopher C. Owens  
The University of Chicago  
Artificial Intelligence Laboratory, Department of Computer Science  
1100 East 58th Street, Chicago, IL 60637  
borse@cs.uchicago.edu

## Abstract

Our research focuses on the problem of recovering from perturbations in large-scale schedules, specifically on the ability of a human-machine partnership to dynamically modify an airline schedule in response to unanticipated disruptions. This task is characterized by massive interdependencies and a large space of possible actions. Our approach is to apply both qualitative, knowledge-intensive techniques relying on a memory of stereotypical failures and appropriate recoveries, and quantitative techniques drawn from the Operations Research community's work on scheduling. Our main scientific challenge is to represent schedules, failures and repairs so as to make both sets of techniques applicable to the same data.

This paper outlines ongoing research in which we are cooperating with United Airlines to develop our understanding of the scientific issues underlying the practicalities of dynamic, real-time schedule repair.

## Irregular Operations Scheduling (IOPS)

Airline schedules are highly complex, structured objects, with large numbers of internal interdependencies. Airlines must confront the consequences of uncertainty in the execution of their daily schedules — uncertainty stemming from inclement weather, sick calls from crew members, mechanical problems with aircraft, constraints on airport resources, and other problems. A

\*This work is supported in part by the Air Force office of Scientific Research under contract AFOSR-91-0112, and in part by the Defense Advanced Research Projects Agency and Rome Laboratory under contract F30602-91-C-0028. The authors gratefully acknowledge the assistance of United Airlines in providing data and observer access to live operations. Nothing in this paper represents any policy, position, or opinion of United Airlines.

snowstorm at a key airport, for example, can have devastating consequences on the operations of an airline, effects from which it may take days to recover. The interdependencies among factors like crew scheduling, maintenance routing, and congestion at airports add further complication to the daily planning problem. Because of these interdependencies, even a single disruption and the consequent attempts at recovery typically involve widespread and long-lasting downstream effects. The search space of possible recoveries to a schedule disruption is enormous.

Airlines employ schedule planners who attempt to mitigate the effects of schedule disruptions. Their main goals are to minimize both passenger inconvenience and the cost of implementing the repair, while accounting for crew work rules, aircraft maintenance schedules, and other factors. An additional goal is to minimize the overall complexity of a repair.

Controllers attempt to balance the trade-offs and uncertainties of irregular events, typically using information provided by various decision support systems such as real-time scheduling displays and passenger booking data. However, very few, if any, of these systems provide the planner with decision-making advice in the form of strategies or specific recommendations to counteract the adversity of a particular event. The goal of our research is to develop the scientific foundations for a new class of decision support tool to address this problem.

From the viewpoint of Artificial Intelligence planning and decision support, the key features of the irregular operations planning task are:

- Airline schedules are large, complex, and highly interdependent.
- Solving schedule problems by exhaustive search is generally infeasible.
- Current situations typically share more with past situations than they differ from them.

- While they may be similar, no two situations are ever entirely identical. This means that simply storing and reusing a "library" of solutions will not suffice.

The size of the search space, together with the recurring nature of typical problems, suggests a solution based on the re-use of plans. But re-using plans means more than just retrieving and replaying old solutions. Because the details of situations change over time, the system will need to be able to notice that a retrieved plan does not exactly fit the current situation, therefore it will need to modify its retrieved plans to fit new situations.

Our approach to plan repair is to provide qualitative expertise in the form of a case library linking descriptions of stereotypical problems with appropriate recovery strategies, and quantitative expertise in the form of optimization techniques drawn from the Operations Research (OR) community. The goal of our research is to develop the scientific foundations for a new class of decision support tool. The IOPS Advisor, currently under development, couples the experiential knowledge of schedulers, which is essential in generating strategies for solving a schedule problem, with the quantitative power of operations research techniques, which are effective in comparing the costs and effectiveness of the potential solutions generated by those strategies. Furthermore, the quantitative models may be responsible for optimizing the details missing from a sketchy solution suggested by a qualitative strategy. For example, if a strategy is "stop to refuel", a quantitative analysis may indicate where to stop and how much fuel to take on.

The IOPS Advisor, currently under development, is intended to represent schedules, failures, and repairs so that both sets of techniques can cooperate using the same representational constructs.

### Research Objectives

The primary scientific focus of this work is on representation. Specifically, we are determining how to represent schedules, schedule failures, and repair strategies so as to enable the IOPS advisor to:

- Identify and characterize schedule problems so as to determine the applicability of prior solutions or specific quantitative techniques.
- Acquire new descriptive features as they become necessary to discriminate among otherwise indistinguishable situations.
- Compare the applicability of multiple, competing solutions to the same problem.

### Knowledge Representation Issues:

The main knowledge representation issue, and the primary focus of our current activity, is to categorize and represent the heuristic knowledge used by controllers and OR analysts, specifically:

- How problems are detected and described.
- What problem-solving strategies exist.
- What aspects of a problem indicate the applicability of one strategy over another.

In order to gather a realistic set of failures and repairs, we have been observing controllers as they detect, diagnose, and repair schedule problems. Our initial study has suggested to us that controllers build and use sophisticated, high-level repairs from a small number of primitive operators. The primitives form the basic representation vocabulary used to describe actions, and it is anticipated that the list will be stable over time. The higher-level strategies, on the other hand, are more dynamic, and one of our tasks is to model the acquisition of new high-level strategies.

Typical primitive operators represent concrete actions like:

- Cancel a segment
- Delay a segment
- Divert a flight to a different airport
- Substitute one aircraft for another
- Substitute one crew for another
- Ferry an empty aircraft from one airport to another

Higher-level strategies, on the other hand, may involve both primary actions and secondary actions designed to mitigate the side-effects of the primary actions. Or, they might involve a series of steps taken to defer the impact of a problem, in the expectation that an opportunistic solution may present itself in the intervening time. Other high-level strategies include geographically localizing the impact of a problem or, conversely, diluting the impact of a problem by spreading a minor delay across several geographic points.

As we gather more high-level strategies from our observation of controllers and from our encoding of quantitative techniques, our plan is to encapsulate the strategies in knowledge structures that also include descriptions of appropriate situations for the strategies. The IOPS advisor will extract from the user a description of the current situation, propose repair strategies based upon the match between the current situation and the stored descriptions, and quantitatively evaluate the utility of situations generated by competing strategies. As it performs this selection and comparison, it can acquire, from the user, information about features of the world that determine the applicability of one strategy over another. These newly-acquired features can then become part of the selection criteria encoded with the strategies in memory.

### Knowledge acquisition

While the list of primitives is expected to remain relatively static, an important aspect of the IOPS Advisor

is that it will be able to acquire new descriptive features as it is used. If the system erroneously suggests a prior case as being a good match to the current situation, the user can correct this by supplying a descriptive feature that would differentiate the current situation from the case stored in memory. The error might have occurred either because the discriminating feature was not mentioned in the description of the current situation, or because it was not mentioned in the stored case. In the latter scenario, it can be added.

In general, a longer-range goal for the IOPS advisor is that, in having a human user interact with a planning tool, we have an opportunity to record information about plan accessing strategies, modification techniques and typical failures that can, in turn, become the heuristics used by a more autonomous system. A system that observed human schedulers in action and recorded their responses to specific planning problems, and which indexed those responses in memory using the functional criteria discussed above, would become a powerful expert assistant — an assistant with a good memory for what worked and what didn't in the past.

### Case-Based Planning Issues

While case-based planning addresses many of the qualitative problems in the irregular scheduling domain, much work must be done before a practical system could be put in the hands of a human scheduler. Fortunately, the core idea in case-based planning, that of incremental modification, is one aspect of the technology that could be usefully applied in the near term as a way to deal with the type of changes that have to be made to schedules during execution.

One of the recurring problems of automated planning is the issue of the repairs that have to be made during execution as a result of unforeseen circumstances. There are always unexpected problems that arise. Weather, crew sickness, and equipment failures cannot be predicted. Bottlenecks show up where none was suspected. Each of these classes of problems can be recognized using a specific set of symptoms, and each requires a specific type of repair.

Run-time repair and optimization, while useful, has to be traded-off against the overall stability of an existing plan. If a single aircraft is unexpectedly grounded, one form of optimization might be to rebuild the entire system schedule, minus that aircraft. But even if such a repair were computationally feasible, implementing it would be preposterous. A planner that deals with unexpected changes in the state of the world by completely replanning will be constantly creating new plans that will do little more than confuse the people that are using them. What is needed instead is incremental, local plan repair, coupled with local optimization. One wants to perturb the schedule as little as possible in the achievement of an acceptable response to an unexpected occurrence.

Much of the emphasis of CBR research to date has been on issues of plan indexing, retrieval and modification. While these issues are clearly present in this domain, our emphasis is primarily on plan evaluation through objective analytical (OR) tools which are also under development. Specifically, we are focusing on how to direct the search for relevant cases based on the OR model's assessment of the feasibility or "utility" of previously proposed solutions. Because the two sets of techniques tend to characterize the problems differently, integrating them is a challenge.

### Operations Research Issues

Operations research analysts tend to think in terms of opportunities for optimization. One of our preliminary findings is that schedule planners do not readily identify these opportunities. Accordingly, an important aspect of the integrative research is to identify classes of situations in which particular optimization techniques are appropriate, and to select descriptive features that allow the system or planners to differentiate among these classes. We intend to codify this knowledge in the form of cases which couple the relevant optimization techniques with characteristic features of the appropriate class of situation.

### Case Study

The following hypothetical case study is based on observations of airline planners. The case illustrates the interplay between qualitative and quantitative reasoning described in this paper. Airports are designated by the following three letter codes: SFO = San Francisco, EUG = Eugene, and MED = Medford.

A runway construction project at EUG has imposed a weight restriction on departing flights. A departing flight EUG-SFO is over the weight limitation by approximately 20 passengers. The flight is scheduled to depart on time, however, inbound flow control is in effect at SFO (due to fog) and is imposing a 53 minute pre-takeoff delay on the EUG-SFO flight.

The planner generates some alternative solutions:

1. Move the excess passengers to a later EUG-SFO flight.
2. Have a flight enroute to SFO passing nearby EUG stop to pick up the excess passengers.
3. Remove enough fuel to carry the excess passengers, and stop at an intermediate point to refuel.

At this stage, the alternatives are qualitative: they simply match a problem with a strategy. Although in many cases this step of the solution process is trivial (e.g., weather-related IOP forces cancellations), we believe that in general this step is non-trivial and it is one aspect of the planner's job which distinguishes an experienced planner from an inexperienced one.

The next step of the planning process involves evaluating the relative merits of each proposed strategy with

respect to the planner's goals. In this case the planner chose not to solve the problem using strategy (1) because pushing the problem to a later flight would most likely cause weight restriction problems downline and would disservice the excess passengers. Strategy (2) was not chosen since it would involve delaying a large number of passengers on a different flight to accommodate a relatively small number of connecting passengers on the EUG-SFO flight. On further analysis of strategy (3), the controller determined that, since SFO air traffic control had already imposed a 53-minute delay on the inbound flight for reasons of airspace crowding, the flight could in fact refuel at MED and carry all passengers to SFO as planned without incurring additional delays. The cost of landing and departing at MED was considered negligible in comparison to the alternative costs of delaying passengers and causing misconnections of aircraft and people (although this calculation was not performed explicitly).

Notice that the planner's analysis in choosing among alternatives remains highly qualitative. The planner uses various sources of information to determine the viability of each approach, however, he rarely explicitly calculates the cost impact of various strategies. We believe that at this stage the planner could be greatly aided by OR models which:

- provide an objective analysis of the relative merits of each strategy based on utility measures.
- determine optimal implementations of high-level strategies, for example, given strategy (2), choosing an appropriate flight, or, given strategy (3), choosing an appropriate airport.

### Anticipated Results

Our key preliminary result is a growing catalogue of stereotypical problems and appropriate repair strategies, which form the backbone of a domain theory of schedule failure repair. We anticipate that a longer-term result of our research will be a working prototype of the IOPS Advisor System. This prototype will embody the failure descriptions and recovery strategies, as well as a set of features characterizing appropriate situations in which to apply specific quantitative optimization tools. The knowledge-based system will suggest strategies, given a description of the problem, while the OR components will be responsible for evaluating the costs and benefits of the proposed strategies and for determining specific implementations of the strategies.

### Evaluation

The bases against which we can evaluate the IOPS advisor project are:

- Does the system enable a controller to produce good schedule repairs? In particular, can a controller use the system's prepackaged strategies and OR evalu-

ation methods to improve upon solutions produced using the controller's own judgment?

- How good are the high-level strategies that the experienced planners employ? How often do controllers choose the best strategy? While the strategies obviously work, are they applied inappropriately? Does post-facto analysis repeatedly indicate that some other strategy might have been preferable?
- Are individuals able to make use of the canned strategies? Can one individual recognize and re-use canned strategies? Is there any transfer across individuals, such that one individual can use strategies developed by another? If so, how should the strategies be presented to the user?
- Can novices use the strategies and optimizations from the IOPS advisor to generate expert-like repairs? In general, how do solutions built by novices differ from solutions built by experts? Does the availability of a library of expert solutions improve a novice's performance?
- Does the integrative AI/OR approach provide a better method than either technique applied alone? Is it even possible to model the IOPS problem using either technique alone? What form would these models take (e.g. large scale linear programming, expert-system)? How would each of these approaches compare to the integrative approach?

### Summary

The airline irregular operation problem is representative of a general class of scheduling problems. An ideal solution would embody both the best quantitative techniques and the genuine expertise of skilled, experienced controllers. Traditionally, the two classes of solution have been described in such divergent terms as to make integration, or even comparison, difficult. By building a uniform representation of schedules, failures and repairs, our intention is to provide a framework for experimenting with qualitative and quantitative solutions and, ultimately, for integrating the two.