Commentary on: A local approach to automated correction of violated precedence and resource constraints in manually altered schedules

Roman Barták and Tomáš Skalický

by

Thomas Starbird
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive, M/S 301-250D
Thomas.Starbird@jpl.nasa.gov

The paper by Roman Barták and Tomáš Skalický deals with a particular instance of the general problem of automatically correcting a plan that has conflicts in it introduced by manual alteration of a previously correct plan. This is an important problem, because planning situations often have many solutions that are legal (conflict-free), but that differ in optimality in ways that humans are best to judge. I have the following scenario in mind: have the automated system construct a legal plan; have a human alter it to make it better in some human-judged aspect, but with the possibility of introducing illegalities; and then let the system correct those illegalities. The paper supports this type of scenario. In particular, the paper sets as a goal that the correction step should make minimal changes to the human-generated version. Minimality of change would presumably raise the probability of preserving the desirable feature that the human's manual intervention placed into the plan.

Minimality of change also makes the final plan easier to understand by the human. One of the lessons learned in preparation for Mars Exploration Rover (MER) prime mission operations was that an automatically generated plan could be difficult to grasp and to present for approval to the scientists and engineers. The goal of minimizing changes to the plan is important, assuming that the human wants or needs to understand the changes or the resulting plan. In theory there could be cases where the human can give complete control to the automated system, accepting the plan essentially without looking at it. Presumably the automated system would indeed deliver plans that are correct with respect to the models and constraints known to the system. But humans often want or need to understand the plan, either as a double-check of its correctness in aspects known to the system, or more importantly to be able to check other aspects, or to convey the plan (or the changes to it) to people who are interested in or affected by what actions will result from executing the plan. Certainly in the space missions I've been associated with, the Mission Manager, at least, needs to know what the plan is all about.

So the general topic of the paper is a useful one.

The specific problem treated in the paper makes strong simplifying assumptions. Constraints on timing of activities are limited to precedence relations and do not include deadlines. Concerning resources, only unary resources are included (each resource can support at most one activity at a time, though one activity may require more than one resource). These assumptions enable elegant algorithms that are easy to state and prove (once you find them!). There may be applications in "real life" where these assumptions hold. The problems encountered in the Jet Propulsion Laboratory (JPL) missions that I am familiar with, however, have more kinds of constraints. Thus the algorithms in the paper are not directly applicable to them, at least not for a complete solution. In particular, many of the problems faced in JPL missions are ones where there is no solution attainable by altering the start times of activities. Often the challenge is to determine which activities to delete in order to get a legal plan.

It is not clear to me how possible it is to generalize the paper's algorithms to more general cases. There are two possibilities. It may be that generalization is direct, that is, that using the paper's algorithms as a start, elaborations could be found to handle more general cases. Another possibility is that the paper's algorithms could be used as is to correct the problems it addresses and that other methods could be applied before or after to address other aspects. But of course at first glance it would seem that any changes made by one or the other could be contrary.

As the authors mention, while the algorithms were constructed with the general objectives of keeping the total time span of the plan as small as possible and with changing the plan as little as possible, there is no claim that the algorithms are provably optimal in either of these senses. But that's okay I think. Experimenting with the algorithms in various problem situations would determine whether those general objectives are commonly met.

The paper is well written. The algorithms are presented succinctly and their proofs of correctness are clearly and completely presented. I enjoyed reading the paper.

Acknowledgement

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.