Commentary on "Using Generic Preferences to Incrementally Improve Plan Quality" by Gregg Rabideau, Barbara Engelhardt, Steve Chein

Alice F. Bowman

Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel, MD 20723
Alice.Bowman@jhuapl.edu

Summary

In an effort to address some deficiencies of automated planning, the authors of "Using Generic Preferences to Incrementally Improve Plan Quality", develop and test a method that allows the human planner to add preferences and optional goals. As defined by the authors, preferences are quality metrics for variables in completed plans and provide the mechanism for specifying which plan variables are important to plan quality. The authors claim this can dramatically increase plan quality with very little modeling effort. The methodology is developed as an extension to the **A**utomated **S**cheduling and **P**lanning **EN**vironment (ASPEN) platform. This platform develops plans based on discrete hard constraints and mandatory goals and resolves conflicts with an iterative repair algorithm. The resulting plans are feasible but may not be acceptable because they do not include what human planners call soft constraints and optional goals. With the inclusion of these additional considerations, the authors establish a basis for automatically generating high quality plans.

The authors extend the ASPEN platform to increase plan quality by the use of preference variables and iterative optimization. The human planner chooses the time period of the plan to be generated by ASPEN. Hard constraints and mandatory goals are defined as well as optional goals and preferred values of plan variables. Mandatory goals are considered a conflict until the activity that satisfies the goal has become part of the plan. Optional goals that are not satisfied degrade the overall plan score. A plan score is calculated based on the preferred value of specified variables (preference variables). The iterative repair process, a standard part of ASPEN, resolves hard constraints and unmet mandatory goals, and iterative optimization, an extension to ASPEN, strives to increase preference variable scores. Those that are low scoring are identified and addressed until a maximum score is attained. Moving activities in time, reordering activities, etc., can increase scores. Only the highest scoring plan is kept.

The authors describe five different types of plan variable classes that contribute to plan quality. For each of these basic plan variable classes, preferred values are specified (plan variables with preferences or preference variables) and real-valued scores are returned. The plan score is computed as the weighted average of scores for plan variables with preferences. In this way, preferences allow the user to evaluate feasible plans and make quantitative distinctions between different plans. In order to increase plan quality, the plan score must be increased. This is done by iterative optimization, which represents a series of improvements made by the improvement expert algorithms. An improvement expert algorithm is developed for each of the five different types of plan variable classes and is invoked depending on the type of preference variable that is low scoring. The improvement experts iteratively select and make improvements, and thereby search for more optimal plans. This is repeated for all of the low scoring preference variables until the best overall plan score is achieved or the specified time has been reached. Experimental results of this technique are presented for the New Millennium ST-4, New Millennium EO-1, Data Chaser, and Rocky-7 Mars Rover.

Commentary

This paper presents an interesting alternative approach of increasing plan quality by developing an extension to ASPEN, a plan generating platform. According to the authors, plan quality is increased in less time and with less computer resources than without the extensions. It does so by allowing the planner to specify soft constraints (or preferences) and optional goals.

Throughout the development of this extension, the authors make a point of working with mission planners so that their extensions to ASPEN are "user friendly". This is a very important point and is so often overlooked by software developers. In my experience the easier the program/algorithm is to use, the more it gets used and the more its tools are exploited. While other programs/algorithms may offer better features, it is the program/algorithm that is easiest to use and understand that becomes a mainstay of a mission.

Because the trend is toward smaller mission operations teams that are responsible for more tasks, it follows that the mission planner is expected to have more spacecraft mission expertise. In the past this knowledge may have been spread out among many planners. With the use of the described extensions to ASPEN, much of the required expertise to generate quality plans resides in the algorithms themselves, and thereby alleviates some of the burden of knowledge required by the now smaller planning team. This is beneficial from a quality point of view and a manpower point of view.

Some spacecraft mission concepts call for the operations team to swell and shrink based on whether the mission is in the cruise phase or data taking phase. Using an ASPEN type plan generation platform with the described extensions would be beneficial in that the expertise would be "on line" and those coming back onto the mission after an absence would be able to rely upon it. I expect that once a platform such as this is in use, mission planners will want to extend the capabilities of the algorithms to include more and more automated scheduling based on feedback from the spacecraft itself. According to the authors, these types of extensions could be accomplished.

If the intention of the authors is to have this paper read by persons whose backgrounds do not include much programming or AI experience, the addition of a diagram or flowchart outlining the process/data flow would be beneficial. This would make the relationships of the different variables and processes easier to understand. I find myself getting quickly bogged down in the descriptions of all the different variables and trying to remember how they relate as I continue reading through the paper.

As a mission planner, I would like to have a few lines summarizing what decision processes/compromises and quality checks went into the resulting plan. Currently, the only information that a planner receives on a resulting plan is a plan score. The authors state that in some cases resulting plans are a compromise between competing objectives and/or resource usage. No indication of this is given to the human planner, and this information may be very useful when combined with information outside of the scope of the plan generating algorithms. Quality issues are of major concern to any mission planner and in order to feel confident about using an automatically generated plan, some assurance has to be given that the plan was checked against some standard quality checklist.

From the point of view of a project leader, the authors might consider addressing what type of mission would benefit from this type of plan generation concept and also may want to include a section outlining the steps necessary to use this technique. Included in the discussion would be topics such as 1) how the complexity of the mission affects the development work required to use this technique, 2) what type of personnel background is required to prepare this technique for use, and what type of background is needed to use this technique on the mission, and 3) what savings are gained by using this concept (i.e., operations teams size, time required to generate plans automatically versus by standard techniques, etc.). In order to help project costs and time requirements, project leaders would also benefit from a general discussion of the work needed to ready this extended ASPEN platform for use on a mission. Since some missions may already be planning to use the ASPEN platform, the authors may want to point out that much of the modeling work required would already be done. On the other hand, missions not previously considering the ASPEN platform may have more costs associated with preparing the required modeling. The discussion may include computer platform requirements, time estimates to learn to use the extended ASPEN platform, computer models for all the subsystems, etc. This information can then be used to project costs and allow the project leader to weigh them against projected life cycle costs.

On a final note, I would find it very interesting if the authors could extend their presentation of experimental results to include not only results obtained by using the presented extensions to ASPEN but also the results obtained when normal planning tools for the specific NASA missions are used. This would give great insight to the merits of using their technique by allowing the reader to do a direct compare.

I find that this is a very interesting paper presenting many ideas that are very useful in mission operations. Although not a programmer, I can appreciate the difficult task that the authors have undertaken and understand how complicated it is to code for all the different thought processes and decisions that go into developing mission plans. I can also appreciate how hard it is to write a paper that can explain their work in terms that a person unfamiliar with programming and AI can easily understand. They are to be commended for the great work that they have done.