Comments on
Automated Planning for Spacecraft and Mission Design

Peter Norvig

NASA Ames Research Center
Moffett Field California 94035
Pnorvig@arc.nasa.gov

Abstract

The paper *Automated Planning for Spacecraft and Mission Design* by Smith and Stebbins presents the idea of using automated scheduling in the design phase rather than the operations phase of a spacecraft mission. This raises a number of interesting questions.

Introduction

The use of automated planning and scheduling tools for spacecraft has traditionally been restricted to the operations phase: the flight operations team determines a set of goals for the spacecraft and use a planner to generate a plan to achieve the goals. The Smith and Stebbins paper presents the idea of using a planner in the design phase to evaluate potential spacecraft designs. This paper comments on the proposal, in terms of the nature of the simulation environment in which plans are evaluated, the choices for design optimization techniques, and the very need for mission plans.

Simulation Environment

In their introduction, the authors say how simulations and plans are both used to gain information about a spacecraft design. Clearly, one is not much good without the other: only with a high-fidelity simulation of a realistic plan can we gain information. If we can achieve certain guarantees about the results of simulations, then simulation can become a more integral part of the design process, as proposed, for example, in NASA's Intelligent Synthesis Environment (ISE) program. Without the guarantees, simulation can only be a heuristic aid.

Note that the simulation may miss important interactions. For example, it may consider only the topological connections of spacecraft components, and thus fail to point out flaws resulting from thermal or electrical interference between two components that are placed too close together. The designer needs to know what sort of interactions the simulation will catch and what sorts it will not.

A simulation run that verifies nominal operation is useful, but is not sufficient for verifying a spacecraft design. We want our spacecraft to be robust under failure, and that requires simulations of low-probability failure events. One approach is to explicitly introduce hardware failures of various kinds into the simulation of a mission, and report the results over a suite of failure conditions. A more ambitious approach is to have a model of failure probabilities and run the simulation with the possibility of failure at every time step. The accuracy of this approach can be improved by using likelihood weighting to artificially over-sample the low-probability failures.

Optimization

One exciting possibility mentioned in the paper is to use the planner and simulation environment to guide a search through the space of spacecraft designs in an attempt to optimize the design. This raises the question of what communication is allowed between the planner, simulator, and optimizer modules.

It is certainly possible to do optimization without any communication between modules other than their advertised inputs and outputs. Pick a point in the space of design parameters, evaluate it by running the simulator on one or more missions, and measure the resulting performance. Repeat on other design points (guided by the performance measure of previous points) until a sufficiently high-scoring design is found.

However, it may be possible to do much better if the modules are allowed to communicate more deeply. For example, suppose the planner is forced to make a decision because of a resource limitation imposed by a design parameter (for example, insufficient power to perform a maneuver). The planner could record this limitation, and suggest to the optimizer that a future trial be attempted in which the power parameter is set just high enough to allow the maneuver. Alternatively, if the planner is capable of representing conditional plans, it could do so at this point, with a condition based on the designed power level. Equivalently, the simulation could be split into two paths with different power level values, and both paths could be followed. One would need to take care in introducing conditionals or splits, so as not to end up with an exponential run time.

Another question is what kind of optimization technique to use. There are a wide variety of options: hill-climbing, simulated annealing, genetic algorithms, tabu search, linear programming, etc. Which is best? Does the ability of the Aspen planner to introduce special-purpose combinatorial solvers come into play in the choice?

Who Needs Mission Plans?

If a spacecraft has been designed and verified using this approach of simulation and planning, we end up with a different view of what the spacecraft does. Rather than looking at the spacecraft as being capable of executing the set of individual commands that are determined by the flight operations team to achieve the mission plan, we can begin to look at it as capable of achieving mission goals, within the design space parameters. This opens the way for an operations phase that matches the design phase in the way automated planners are used to command the spacecraft.

This also raises the question of whether the planner should be located onboard or on the ground. If an onboard planner has the same inputs as a ground-based planner, then uplinking a goal rather than a sequence of commands can be seen as a form of data compression: the goal is a compression of the possibly large conditional plan that the goal generates. If an onboard planner has access to additional or more up-to-date state information that the ground-based planner does not have, then we have a situation where the plan that would be generated and executed onboard is different from what the ground operators anticipated. For this to be acceptable would require great confidence in the level of testing and analysis done in the design phase.

There is a compromise between onboard and ground-based planning, in which a certain level of control is allocated to onboard processes – a kind of heightened safe mode in which the spacecraft can recover from a variety of problems – and the remainder is determined on the ground. This approach too would depend on thorough testing and analysis, starting in the design phase as Smith and Stebbins envision.