

Commentary on “SWIM: An AI-based System for Organisation Management”

Robert D. Rasmussen

Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, California 91109
robert.d.rasmussen@jpl.nasa.gov

Abstract

This is a short commentary on a paper appearing in this workshop by Pauline M. Berry of SRI International and Brian Drabble of the University of Oregon. Their paper reports the development of a workflow management system using AI techniques for reactive control, scheduling and continuous execution. While developed for application in information surveillance and reconnaissance, the authors claim applicability to a variety of domains, including space.

Background

The commentary in this short report comes from the perspective of someone steeped in interplanetary space missions. This undoubtedly distorts the perspective that can be brought to general applications of planning and scheduling technologies for space, because in the experience of this commentator the emphasis so far for AI technologies in this narrower realm has been on increased autonomy and other improvements directed toward the flying vehicle.

The greatest challenge in commenting on the subject paper has been to stretch this point of view to accommodate the authors' vision—something that might be more readily achieved by others. The following commentary is offered in that spirit.

Setting a Context

Discussion of applications for planning and scheduling in the space community has generally involved the following key objectives:

- Automation of routine functions for reduced operations cost
- Reduced development cost and schedule through higher level programming and model based reasoning
- Reduced complexity and greater efficiency in the deployment of limited resources
- Remote system autonomy in uncertain environments and continued operation in the presence of faults
- Improved mission return through in situ analysis and response

In many cases, these are old problems that are well understood. They generally involve the minutiae of a complex, tightly coupled system with many parts (like a spacecraft, rover, or ground station), and organizing all their many constraints and interactions within a context of competing demands on the system. In a sense, one can view the management process for such a system as omniscient, in that an attempt, at least, is made to understand everything and be ready for anything. The principle issues in applying automated planning and scheduling techniques in this domain are in marshalling all the details needed to describe a system, in making it easy for a team of diverse talents to do so and to operate the resulting system, in getting these complex technologies to meet embedded, real-time performance demands, and in convincing a very conservative culture to place their trust in a technology whose products often defy simple explanation. Answers may not be easy, but at least the criteria they need to meet are evident.

When we move into fresher territory and there are no clear solutions to the problems at hand, short of applying AI techniques, a host of difficult issues arise. Most unsettling is the fading option of total understanding and control. We accept this reluctantly, but not without a reserved notion that, even in these cases, we retain a perimeter around the problem that we can police. This is possible because we have still tended to focus on fairly narrow pursuits, such as the operation of a single vehicle in an uncertain environment, or at worst a few systems, most in relatively certain environments. Moreover, we've tended to assume that uncertainty was largely resolvable, once encountered. Discussion of planning and scheduling issues in space applications has therefore tended to center around a comparatively disciplined set of approaches.

A Place for SWIM?

SWIM, on the other hand, having addressed the world of information surveillance and reconnaissance for DARPA, has grown out of quite different situation in which the set of assets to be managed is potentially vast, each asset or objective may itself be a complex system with idiosyncratic behavior (e.g., it can include people), goals

are often only subjectively defined, constraints may be soft, and so on. In some ways, this is a much harder problem. Consequently, it seems forced to relinquish a level of rigor that we have assumed for space applications.

Given this disparity, it was at first difficult to grasp how one might bridge the apparent gap between the types of problems SWIM purports to address and those more typically bandied about when space is the subject. However, following some dialog on this issue, the authors have offered in their paper a number of examples tailored to space applications that clarify the potential role for this sort of approach. These examples encompass everything from planning procurements to conducting tests—not the usual sort of thing that comes up in debates over AI applications in space.

Put in this light, the workflow management approach described by the authors, while not of the prevailing kind we generally expect, would nevertheless provide some welcome discipline in processes that often appear to sustain themselves on much more tangled fodder. That is, rather than offering less rigor when considered for unintended applications, SWIM appears to offer more rigor than usual when applied as intended. The introduction of a system like SWIM to a forum of space practitioners, therefore, opens a new avenue of discussion largely unexplored—especially in the unmanned sphere where autonomy is foremost in the common dialog.

With this insight it is worthwhile to explore which aspects of space applications might have the right tenor for this approach. It is hard to imagine, frankly, whether some of the processes we presently suffer could comfortably tolerate the scrutiny such a system would impose. One can certainly see opportunities in project development and test, or in the complex logistics of something like a space station.

More tractable, however, may be some real technical challenges just on the horizon that will arise when present ambitions for ubiquitous presence throughout the solar system begin to be met. These systems begin to exhibit characteristics similar to those described in the paper (though smaller in scope). One can imagine, for instance, several aerobots drifting through the atmosphere of Venus, where it is the collective behavior of several vehicles that must be managed, while for each vehicle individually, behavior is potentially erratic, objectives are fuzzy, at best, and assignments may change from one vehicle to another as atmospheric whims dictate.

Closer to home, fleets of earth observing and communication platforms grow harder to manage every year as numbers increase in bounds. As it becomes necessary to address these units as one large coordinated asset, workflow management systems may become invaluable.

It really doesn't take too much effort to find examples where this sort of thing could become important for space.

A Brief Assessment

SWIM may very well be a contender in this arena. Making such an assessment from the paper, however, is difficult since most of the interesting questions and issues that one might wish to explore cannot be summarized in a short paper. The authors properly motivated the work, and provided a competent summary of the general structure, but SWIM is a large system with many components, each of which could be given only terse treatment.

Details on results so far were likewise short. The work has been performed mainly for the information surveillance and reconnaissance community, so more detail about this application domain would probably have been hard to extrapolate to space. In any event, a practical assessment by the customer of this work would be more illuminating as an indicator of proper bearing. In this vein, the authors did offer a brief discussion of some of the concerns they have in introducing such as system and clearly indicate that much work remains.

What can be said is that where insight into the philosophy behind the approach could be gained, the ideas presented seem to be on track. Recognizing the importance of a continuous planning and execution framework, for instance, was gladly received.

It should also be mentioned that incorporated components, such as the Procedural Reasoning System, the Squeaky Wheel Optimizer, and others, will get a good workout in this architecture, and if nothing else will offer further validation for these items.

Conclusion

It is this commentator's opinion that there is room for systems such as SWIM in the discussion of planning and execution technologies for space. Furthermore, there are elements of SWIM which clearly have potential application in more commonly discussed space applications. There may even be straightforward modifications of SWIM to these applications that can exploit this potential. This would be worth exploring.

Regardless of the technology inside, though, the real key for any such architecture is successful adaptation to a real application fielded in the real world with real users. This should be on the road map for any serious technological endeavor. It will be interesting to see how SWIM evolves in this regard.