ULISSE: A Planning Trojan Horse to Enter the International Space Station?

Derek Long
University of Strathclyde, Glasgow, UK

1 Introduction

In their paper, "Supporting Incremental Planning Processes within the ULISSE Framework", Cesta et al. describe the framework and infrastructure that supports the exploitation of science-gathering systems on board the International Space Station (ISS). This consists of a hierarchically structured series of units that take responsibility for planning the activities associated with experiments, with User Support and Operation Centres (USOCs) providing plans for individual experiments and the Columbus European Planning Team (EPT) managing the integration of these plans into a single schedule.

Inevitably, these complex tasks require some automated support. The authors observe that one of the aspects of the ULISSE project is the provision of a Planning and Validation Tool (PVT), which is intended to provide some of this support to USOCs (perhaps also to the EPT, although this is not stated). The authors focus on the Planning and Scheduling Service (PSS) within the PVT: a relatively small part of the ULISSE project as a whole. More specifically, they examine how this service can be used to aid in planning activities for the Fluid Science Laboratory, selected as a case study because of its complexity. The PSS uses as its underlying model a timeline representation developed by the authors in earlier work (Cesta and Fratini 2008).

2 Any Language as long as it's Greek

Two of the challenges in transitioning research into application, from the position of University researcher, are achieving both the degree of knowledge and expertise in the domain of the problem application (as opposed to the research area that is to be exploited in its solution) and achieving the level of trust on the part of the potential users to convince them to invest the effort in reaching plausible application of the research ideas. This paper clearly demonstrates the extent to which the first of these challenges absorbs effort on the part of the research team: there is a vast depth of background knowledge presented in summary in the paper and it is obvious that, in order to make progress with raising trust amongst the potential users, the research team had to make considerable investment of time and energy in acquiring this background. A very significant service that the team could perform for other researchers interested in making progress in this application area is to make this knowledge as accessible and succinct as possible. As with many technical areas, the domain is filled with alien acronyms and specialist terms and these act as a barrier to entry for those not given access to the necessary background materials.

The authors explain, at some length, the operations of the FSL and the constraints that govern the construction of plans for its use. It is not clear from this account which of these constraints makes planning for the FSL a difficult task (perhaps they are all tight constraints, but it is not obvious).

The PVT provides the PSS for plan and schedule synthesis and a Formal Verification Service (FVS), which is currently based on use of the Murphi model checker (Della Penna et al. 2003). Was the decision to use this path a consequence of the organisation of the project or was it based on an analysis of the best path to follow?

The authors use, of course, their prior art in the construction of a model of the FSL: the domain description language for their timeline formalism, developed for earlier projects (Cesta and Fratini 2008). The authors claim as the advantages of this that they can "use abstract and general methods to define algorithms for solving specific P&S problems" and "exploit a domain independent, general purpose planner and scheduler to efficiently solve well identified specific sub-problems." There are, of course, alternatives: the action-centred formalism of PDDL has supported the construction of many more domain independent, general purpose planners than any other planning domain language. One can argue about the relative utilities of alternative formalisms and generate a great deal of heat without much light, but ultimately this decision appears to come down to a matter of personal preference: the evidence in terms of supporting algorithm development and weight of research interest is that following standards always encourages direct comparison and scientific evaluation of claims about systems. The planning systems that are least widely adopted are those that can only use domains written in languages known to a few. The counterpoint to this is that it is often these planners that are developed for application and that are therefore the most successful in deployment and service.

There are always compromises in turning research into application and planners are no exception. The authors would no doubt claim that PDDL is too restrictive and their application demands a more expressive language. It is unfortunate that, as with so many applications, the separation between the application and the large body of research begins so early in the investigation of the application: the consequence is that the application is of limited interest to the research community, since it is inaccessible to them — the models are written in a formalism none of the rest of the community uses.

3 A Trojan Horse?

As indicated above, one of the major challenges in transferring research into application is in achieving a level of trust amongst users. Case studies, even with limited goals, can be critical opportunities to gain acceptance and to insert technologies that become part of the accepted scenery: a Trojan Horse by which to introduce new and more interesting solutions to deeper problems. It is not entirely clear how much of an opportunity this work represents. How far do the authors believe that this case study might allow them to progress in pushing acceptance of planning and scheduling technology and will it be an entry to the USOCs, the EPT or the ISS?

4 In Summary

The authors present a rich and complex potential application for planning. They show how they have approached the problem themselves, using their prior art based on their own modelling language and tools. I will be interested in hearing during the discussion of this paper:

- Where is the core of the planning problem for FSL? Amongst the various constraints, which ones really limit the range of feasible solutions?
- Is it inevitable that the domain models are essentially inaccessible to researchers? How much opportunity is there for this work to reach out to the broader planning community? How fundamental is the role of Murphi to this work?
- To what extent does this work form a basis for future developments? Is the work creating an opportunity for future exploitation of P&S technologies in the ISS programme?

References

Cesta, A., and Fratini, S. 2008. The Timeline Representation Framework as a Planning and Scheduling Software Development Environment. In *Proc. of 27th Workshop of UK Planning and Scheduling SIG*.

Della Penna, G.; Intrigila, B.; Melatti, I.; Minichino, M.; Ciancamerla, E.; Parisse, A.; Tronic, E.; and Zilli, M. 2003. Automatic Verification of a Turbogas control System wuth the Murphi Verifier. In *Proc. of 6th International Workshop on Hybrid Systems: Computation and Control (HSCC)*.