Commentary on

"On-board decision-making on data downloads" by Verfaillie, Infantes, Lemaître, Théret and Natolot.

Robert Hawkins
Space Telescope Science Institute
Baltimore, MD, USA
rhawkins@stsci.edu

Introduction

The paper (Verfailleie, Infantes, Lemaître, Théret and Natolot. 2011) addresses the Spacecraft Data Download Problem. The problem is relevant in space-based missions where the spacecraft records data and needs to download it to ground based mission centers. In particular, this problem is relevant when the volume of these downloads is unknown and when the windows available are limited (insufficient to transport the entire data volume), for example EO missions, rovers and surveillance missions.

Three primary features characterize this problem:

1. Opportunity windows for data download from the spacecraft are limited.
2. The rate at which data can be downloaded is limited.
3. The volume of data per activity is unknown a priori.

The paper presents results contrasting the performance of ground-based plan-ahead approaches for the Spacecraft Data Download Problem with several approaches for on board decision-making.

Problem Formulation and Analysis

The paper begins with a formalization of the Spacecraft Data Download Problem. The primary components are activities, download windows (with associated ground stations), data rate information and various weighting, priority and utility functions that can be used to evaluate the problem results.

The problem is complicated by temporal constraints and the fact that data download volumes for activities are unknown. The temporal constraints are on the availability of the data – latest download time, utility of data decreasing with time, etc. Deferring downloads, and/or downloading to a remote station could render the data useless prior to arrival at the mission center. Each activity has a priority level, which significantly effects how the heuristics described perform.

Following formalization, the problem is then compared to existing problems from the literature. Generally, the problem can be seen as a variant of the (multi) knapsack problem, but is complicated by the temporal constraints presented and the unknown data volumes. When both are considered, the problem becomes close to the stochastic (multi) knapsack problem with objects with uncertain weights.

Finally, the problem is described as an MDP. The decision points are the starts of windows, ends of activities and ends of downloads. Time t and the data volume in memory for each activity define state. Actions are the activity and download window choices. Uncertainty is introduced due to the unknown data volumes, as previously described.

Experimental Approaches

The experimental section of the paper describes eight approaches to planning and onboard decision-making.

As a baseline, two ground-based approaches to planning the problem are described. Both of these approaches use as basis a traditional knapsack greedy algorithm. Since the data volumes of activities are unknown at decision time, assumptions must be made about the data volume. The two approaches, PG1 and PG2, differ in that the lowest priority activities are given, as part of the assignment process, an estimated volume in PG2, and maximum volume in PG1. The greedy algorithm, at each step, chooses an activity a from those that are downloadable using an available window. Amongst these activities, the activity with the highest priority that maximizes the utility density (ratio of utility at the earliest time with the duration of the download of a) is chosen.

Next, two on-board decision rule approaches are described. The first, DR1, uses a heuristic similar to that of the greedy algorithm above to choose an activity at each decision step, differentiated by the fact that it may only choose from activities immediately available (already ended) and from a currently available window. DR2, the second decision rule, allows for look-ahead in that it enables choosing a high priority activity that has not yet ended. The density ratio is slight changed in that it is the ratio of the utility to amount of time from the current time through the end of the download. Additionally, the rule will recursively fill the gap between the current time and the start of the previous activity chosen with additional downloads. This significantly helps low-priority data get downloaded.

The decision rules are followed by two on board implementations of the ground-based greedy algorithm that rebuild the plan online using knowledge about real volumes as they become available. The first of these, PB1, re-plans after every activity, while the second, PB2, re-plans at every decision step (as described in the MDP formalization).

The final two approaches use sampling-based approaches to planning on board. The first, SB1, modifies the previous on board algorithms by sampling possible volumes for uncompleted activities using known probability distributions and creating multiple plans over a limited horizon into the future, choosing as the next action the activity that appears first in the most plans produced. The final approach, SB2, extends the concept of SB1 further by evaluating each action at each step, sampling the future volumes and generating plans given that action, and producing a mean utility for each action based on its utility and that of the plans it could generate. The action with the maximum mean is then chosen.

Experimental Results

The eight approaches to the Spacecraft Data Download Problem were all implemented and tested on an experimental data set. In the experimental data set, there are 3 priority levels, ranging from highest (1) to lowest (3). The data set includes 1484 activities assigned to the different priority groups, with each activity given weight 1, and data volumes given a Gaussian distribution. On board planning is performed, for all approaches, with a 30-minute look-ahead horizon.

The eight approaches are compared and contrasted with relation to total downloads, utilization and utility (per priority level). In addition, the on board methods are compared with regard to computational processing time.

Ignoring computing time, the ground based planning methods are easily superseded by all of the on board methods, and amongst the on board methods, DR1 stands out as the inferior method. However the rest of the approaches do not differentiate themselves so clearly. DR2, PB1 and PB2 stand out as (seemingly) the most practical, when compute time is factored in, as SB1 and SB2 are orders of magnitude more expensive.

Comments

- What is the value of H_l (half life) used in the problem set?
- Have you studied the effects of having a non-uniform weight in a problem set?
- The problem description allows for more than three priority levels. Have you considered what the impact of a data set with for or more priority levels would be on the approaches you've described?
- What is the value of the utility function for an activity that is not downloaded?
- Adding the utility density function(s?) to the formalization would be helpful. Do these functions account for Dt_s?
- Showing relative performance to the on board approaches is clearly important. It would be nice to also have that performance grounded in real performance terms (both hardware and mission-relative), to get some context of utility. For example, if the spacecraft's processing were as fast as the described 3GHz processor, would the ~1.33s per decision cost of SB2 be practical in a given mission context?
- Application – as always is the question, how willing are space agencies to implement on board strategies such as this, where the choices that algorithms make determine what data the end user will see?

Conclusion

I found this paper to be illustrative of approaches that could be made to extending some planning and plan execution decision making processes into the on board realm. I'd like to see the results extended, and, in particular, given context with an actual space based mission to further explore their utility and applicability.

References

Verfailleie, Infantes, Lemaître, Théret and Natolot. 2011. On-board decision-making on data downloads. In *IWPSS 2011. Proceedings of the 7th International Workshop on Planning and Scheduling for Space, Darmstadt, Germany, June 8-10, 2011.*