Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Flexible Scheduling for an Agile Earth-Observing Satelllite

Adrien Maillard *
ONERA - The French Aerospace Lab
FR-31055 Toulouse, France
adrien.maillard@onera.fr

1 Introduction

Earth observation from space allows us to better understand natural phenomenas such as marine currents, to prevent or follow natural disasters, to follow climate evolution and many other things. To achieve that, there are a great number of artificial satellites orbiting Earth, equipped with high-resolution optical instruments and communicating with a network of ground stations. A satellite is said to be agile when it is able to move quickly around its gravity center along its three axes while moving along its orbit, thanks to gyroscopic actuators. It is equipped with a body-mounted optical instrument. To observe a ground area with the instrument, the satellite must be pointed to it. In practice, users submit observation requests to a mission center, which builds activity plans which are sent to the satellites. These plans contain several types of actions such as orbital maneuvers, acquisition realisations and acquisition downloads towards ground stations. Many techniques are used to synthesize such activity plans. Until now, plans are computed offline on the ground and converted into telecommands that the satellite executes strictly, without any flexibility. However, the satellite evolves in a dynamic environment. Unexpected events occur, such as meteorological changes or new urgent observation requests, that the system must handle. Moreover, resource consumption is not always well known. Until now, to ensure that plans will be executable on board with these uncertainties, they are built with worst-case hypothesis on resources consumption.

The objective of this work is to give more autonomy to the satellite without compromising the predictability that is needed for some activities. On the ground, we have high computing power and high uncertainty, while on board we have very low computing power and low uncertainty. The main idea is to share decision-making between ground and board to take advantage of the high computing power on the ground and of the low uncertainty on board. First we apply this idea to download scheduling which consists in scheduling file downloads during ground station visibility windows. Second, we apply this idea to observation planning.

2 Flexibility for Data-download scheduling

Because of the use of more and more sophisticated onboard compression algorithms, the amount of data that results from an acquisition and thus is recorded on board and must be downloaded to the ground is more and more unpredictable. It depends on the data that has been acquired. In the case of optical instruments, the presence of clouds over the observed area allows high compression rates and results in a low amount of data to be recorded and downloaded.

As satellites are not continuously accessible by a ground control station, generated volumes are known only on board. This leads to think that decisions about downloads should be made on board just before any reception station visibility window with the exact knowledge of the volumes of the already recorded data. However, planning onboard is time-consuming and requires high-computing resources. Moreover, the resulting plans are unpredictable [Pralet *et al.*, 2014]. This is problematic especially for high-priority acquisitions, because users who request an image may want to know when data will be downloaded. In such conditions, it is still possible to build data download plans on the ground with the assumption of maximum volumes (minimum default compression rate). The resulting plans are then always executable on board, but suboptimal due to an under-use of the station visibility windows.

In this work [Maillard *et al.*, 2015], we introduced more autonomy on board satellites while keeping some predictability. The data download problem is a complex scheduling problem with temporal and non-temporal constraints, close to *RCPSP/max* or *flexible job-shop* problems. However, no existing technique can be directly applied to this problem because of specific properties such as time-dependant temporal constraints or a sharing objective in the criterion. A complete plan is built on the ground with maximum volumes for high-priority acquisitions and expected volumes for low-priority acquisitions. For that, we use a Squeaky Wheel Optimization scheme built upon a non-chronological greedy algorithm. The numerous problem constraints are checked with the constraint library InCELL [Pralet and Verfaillie, 2013]. Because of the use of expected volumes for low-priority acquisitions, the plan is not directly executable onboard. Some volumes may be higher than expected and then the plan needs to be repaired. Onboard plan adaptation is performed before each group of sufficiently close visibility windows. A fast greedy repair

---

$*$. This work is supervised by Cédric Pralet and Gérard Verfaillie at ONERA, by Jean Jaubert at CNES and Frédéric Fontanari at Airbus Defence and Space.

procedure builds a consistent plan from this ground plan by removing low-priority acquisitions when they endanger future high-priority downloads. For that, we use lower bounds on the latest starting dates of high-priority downloads. These bounds are computed on the ground and allow to quickly check whether or not the real volume of low-priority download endangers next high-priority downloads. With this look-ahead like mechanism, we ensure that all high-priority downloads will be done. Sometimes, real volumes are lower than expected. In this case the onboard algorithm tries to insert downloads that are not in the ground plan or to move forward future downloads.

Ground planning over a large horizon with sufficient computing power and time allows to produce a plan with a good quality and to take into account several criteria such as sharing between users or information age which could not be possible on board. Board repairing allows to take advantage of real volumes and to update the plan in a limited time (compared to full board planning). The combination of these two planning phases into a flexible scheme makes this approach a very competitive one, even better in some cases than pure onboard planning in terms of criterion and much more predictable for high-priority downloads.

3 Flexibility for Observation scheduling

When we studied flexibility for the data-download problem, we assumed that an observation plan was computed beforehand. We now want to introduce flexibility in the observation plan as well. During an observation, the satellite scans its objective. Because the optical instrument is body-mounted, the satellite must perform attitude movements to perform this scanning. Computing these attitude movements is not feasible onboard, due to limited computing capabilities. Then, we have to ban pure onboard planning. The onboard software can only execute plans or modify them lightly in a way that does not change the satellite attitude.

Nowadays, observation plans are built on the ground. Attitude constraints must be satisfied. A download plan is then built, and the whole plan is simulated. In this procedure, energy, temperature and memory constraints are checked. If they are violated, some of the low-priority acquisitions are removed and the constraint-checking process is started again. There are uncertainties about energy consumption, energy production, temperature evolution and observation volumes. To ensure that the observation plan will be executable without violating any constraint, upper bounds on resource usage are used when checking constraints. Actual energy consumption is often lower than maximum and actual energy production is often higher than minimum, and so on. To increase the system capability, the idea is to remove such constraint checks on energy from the ground planning process for low-priority acquisitions. These checks are replaced by the production of minimum levels of energy that must be present on board at the beginning of each low-priority acquisition to ensure that, even if the low-priority acquisition is performed, future high-priority acquisitions can be done. These levels are produced with a backward simulation scheme that starts from the end of the horizon and computes realization bounds on the levels of energy.

Observation plans produced on the ground are then conditional plans involving conditions for triggering low-priority acquisitions. Once on board, before each low-priority acquisition, if the actual level of energy is higher than the required level computed on the ground, the acquisition is performed. If not, the telescope is not turned on, energy is then saved.

Compared with the current approach, this approach avoids wastage of resource and allows more acquisitions to be executed. Remaining work include computing alternative satellite attitude movements to save more energy in case of acquisition cancelling (attitude movements needed for the acquisition are still performed), extending this approach to other resources such as temperature, and generalizing to more generic problems.

4 Related works and conclusion

The problem of adapting a plan to uncertainties by delaying decision-making is not novel. In scheduling, it is the *flexible scheduling* paradigm [Billaut *et al.*, 2010]. These approaches are linked with *contingent planning* (and its real-world applications on telescope scheduling in *just-in-case planning*) where several plan branches potentially leading to different ends are computed offline. During the online phase, depending on actual conditions, branching in the plan is performed. A derived paradigm, *opportunistic planning* [Fox and Long, 2002], where a nominal plan is computed and augmented with loop branches, is maybe closer to our approach. In this paradigm, the main branch will always be executed because opportunistic branches are always coming back to it. In our case, the main plan is the high-priority plan made of high-priority acquisition and downloads.

Références

[Billaut *et al.*, 2010] Jean-Charles Billaut, Aziz Moukrim, and Eric Sanlaville. *Flexibility and robustness in scheduling*, volume 56. Wiley. com, 2010.

[Fox and Long, 2002] M. Fox and D. Long. Single-trajectory opportunistic planning under uncertainty. In *Proc. of the UK Planning and Scheduling SIG meeting*, Delft, The Netherlands, 2002.

[Maillard *et al.*, 2015] Adrien Maillard, Cédric Pralet, Jean Jaubert, Isabelle Sebbag, Frédéric Fontanari, and Julien L'Hermitte. Ground and Board Decision-Making on Data Downloads. In *Proc. of the 25th International Conference on Automated Planning and Scheduling (ICAPS-15)*, Jerusalem, Israël, 2015.

[Pralet and Verfaillie, 2013] Cédric Pralet and Gérard Verfaillie. Dynamic online planning and scheduling using a static invariant-based evaluation model. In *ICAPS*, 2013.

[Pralet *et al.*, 2014] C. Pralet, G. Verfaillie, A. Maillard, E. Hebrard, N. Jozefowiez, M.-J. Huguet, T. Desmousceaux, P. Blanc-Paques, and J. Jaubert. Satellite Data Download Management with Uncertainty about the Generated Volumes. In *Proc. of the 24th International Conference on Automated Planning and Scheduling (ICAPS-14)*, Porthmouth, NA, USA, 2014.