

Postponing Decision-Making to Deal with Resource Uncertainty on Earth-Observation Satellites

Adrien Maillard and Gérard Verfaillie and Cédric Pralet

ONERA - The French Aerospace Lab, Toulouse, France

{firstname.lastname}@onera.fr

Jean Jaubert and Isabelle Sebbag

CNES, Toulouse, France

{firstname.lastname}@cnes.fr

Frédéric Fontanari

Airbus Defence and Space, Toulouse, France

frederic.fontanari@astrium.eads.net

Abstract

Agile optical Earth observation satellites are subject to several uncertainties: about the actual volumes of data to be recorded on board and then downloaded to ground stations and about the actual profiles of satellite attitude, onboard energy, and instrument temperature. To face these uncertainties, the current practice consists in using safety margins when scheduling acquisitions and downloads offline on the ground. Such a practice produces suboptimal schedules: acquisitions or downloads that are not selected when scheduling might be performed when observing on board the actual volumes of data or the actual levels of energy and temperature.

In [Maillard *et al.*, 2015], a flexible download scheduling approach has been proposed to deal with uncertainty about the volume of data that is generated by acquisitions and must be downloaded: download schedules are built offline on the ground; these schedules are then adapted online on board while guaranteeing that ground commitments to high-priority data downloads be satisfied.

In this paper, we propose to use a similar approach to face uncertainty about the available level of energy when scheduling acquisitions: acquisition schedules are built offline on the ground; these schedules are then adapted online on board by performing low priority acquisitions if and only if the current level of energy allows it while guaranteeing that performing all the future high priority acquisitions remains possible.

1 Introduction

Agile optical Earth observation satellites are space sensors which acquire data, compress and record it on board, and then download it to ground stations. Although predictive models

of the behavior of such satellites exist, these models cannot precisely predict the whole satellite behavior. A lot of uncertainties remain about:

- the future acquisition user requests;
- the presence or not of clouds over the ground area to be observed, which impacts the acquisition success or failure;
- the actual volume of data that must be recorded on board and then downloaded to the ground after the onboard execution of data compression algorithms;
- the actual satellite attitude trajectory, especially when the satellite must perform long attitude movements to move from an acquisition to the next one;
- the actual energy production by the solar panels, which depends on the satellite attitude, the actual energy consumption by the satellite and its instruments, and thus the actual energy profile;
- the actual instrument temperature profile (optical instrument and data emission antenna).

To face these uncertainties, safety margins are used when scheduling acquisitions and downloads offline on the ground. For example, minimum data compression rates, minimum energy productions, maximum energy consumptions, and maximum temperature evolutions are considered. The result is that schedules are artificially over-constrained and thus sub-optimal. Acquisitions or downloads that are not selected when scheduling might be performed when observing on board the actual volumes of data or the actual levels of energy and temperature. This incorrect assessment of resource production and consumption may lead to oversizing satellite characteristics at the design stage.

To overcome this difficulty, a first approach would consist in delaying acquisition and download decisions online on board when actual volumes of data and actual levels of energy and temperature are known. This is the approach that has been developed and experimented in the context of the autonomous Earth observation satellite EO-1 [Chien *et al.*,

2004]. Its main drawbacks are that acquisition and download scheduling is a computationally complex task, that computing resources available on board a spacecraft are very limited, and that time for online decision-making is limited too.

This is why an intermediate approach has been proposed in [Maillard *et al.*, 2015] to perform data download scheduling under uncertainty about the volumes of data to be downloaded. Flexible download schedules are produced on the ground when volumes are uncertain, together with commitments to high-priority data downloads. These schedules are then adapted on board when exact volumes are known. The onboard adaptation process guarantees that ground commitments to high-priority downloads be satisfied, but guarantees nothing about low-priority downloads. It tries only to perform the latter as well as possible. With such an approach, the onboard adaptation process remains computationally light and the main computing effort is supported by the ground.

In this paper, we present a similar approach that has been developed to perform acquisition scheduling under uncertainty about energy. Flexible acquisition schedules are produced on the ground when energy levels are uncertain, together with commitments to high priority acquisitions. The onboard adaptation process guarantees that ground commitments to high-priority acquisitions be satisfied: they will be performed exactly as foreseen in the ground schedule. However, it guarantees nothing about low-priority acquisitions: each of them will be performed if and only if the current level of energy allows it and allows all the future high-priority acquisitions to be performed as well.

In Section 2, we present some facts about the physical space and ground system we consider that are useful to understand the acquisition and download scheduling problem and the approaches that can be proposed to handle it. In Section 3, we describe the current approach where all the decisions are made offline on the ground. Then, in Section 4, we present the flexible approach we propose where decisions are shared between ground (offline) and onboard (online). Section 5 presents the experimental results that have been obtained when comparing both approaches and Section 7 concludes with some proposals aiming at improving the flexible approach and at generalizing it beyond this specific space application.

2 Some useful facts about the space and ground system

The context of our work is a post-Pleiades project whose objective is to improve the currently operational French Pleiades Earth observation system in terms of imaging capacity, quality, and cost. We retain the main assumptions of this project.

User requests At any moment, users may emit observation requests. An (elementary) observation request is made of an emitting user, a priority level, a target ground area, and some observation constraints: observation period, maximum observation angle, minimum image quality... Only two priority levels are considered: high and low levels. High-priority requests are generally emitted by defense users, but may be

emitted by civil users in case of catastrophic events, such as natural disasters.

Data acquisition Data is acquired by the unique optical observation instrument which is body-mounted on the satellite and more precisely by the detector line it includes. A given ground area is thus acquired by scanning it. Scanning a given ground area takes some time and requires the satellite attitude to be permanently controlled to take into account the movement of the satellite along its orbit, the rotation of Earth on itself, and the required scanning. The satellite attitude is controlled by gyroscopic actuators which allow the satellite to move very quickly around its gravity center, along the roll, pitch, and yaw axes, while moving along its orbit. Gyroscopic actuators are used during acquisition, but also between acquisitions to move from the acquisition of an area to the acquisition of the next one. As a consequence, if we consider an acquisition schedule, that is a sequence of acquisitions, we must first check whether or not it is kinematically feasible, that is whether or not the attitude trajectory it requires is feasible, taking into account actuator limitations (maximum angle, maximum speed and acceleration ...). We must also check whether or not the detector line may be damaged by the sunlight (dazzling of the observation instrument due to an insufficient angle between the satellite observation and Sun directions). Producing a satellite attitude trajectory and checking it is computationally time-consuming.

Data recording Data generated by acquisition is compressed and then recorded on board in a mass memory whose size is limited. The actual compression rate depends on acquired data and is unpredictable.

Data downloading Recorded data is downloaded to ground data reception stations within station visibility windows. Data downloading uses an orientable emission antenna. To download data towards a given station, the antenna must be permanently trained on it. This requires the antenna to be permanently controlled to take into account the movement of the satellite along its orbit and around its gravity center, the rotation of Earth on itself, and the required training. The antenna must be also controlled to move from a station to a next one. All the constraints that must be verified by a download schedule, that is a sequence of downloads, are precisely described in [Maillard *et al.*, 2015].

Instrument switching To save energy, to keep their temperature below a given threshold, and to maximize their lifetime, the acquisition and download instruments (optical observation instrument and data emission antenna) are not maintained permanently active. Each of them has three modes: ON, standby (SDBY), and OFF, in a decreasing order of energy consumption. Each instrument is maintained ON only when necessary: effective acquisition or download. Between acquisitions or downloads, it is maintained in the SDBY mode if and only if the duration between the two successive acquisitions or downloads is less than or equal to a given threshold. Otherwise it is maintained OFF.




Figure 1: Typical energy profile.

Instrument temperature If we consider an acquisition and download schedule, including the associated instrument switching plan, we can build an approximation of the associated temperature profile of each instrument and check whether or not its temperature is always lower than or equal to the maximum acceptable level. Important margins are used in that domain because it is very difficult to model thermal interactions between the platform and the instruments.

Energy production and consumption Energy is consumed by the satellite platform, and by the acquisition and download instruments according to their current mode (ON, SDBY, or OFF). It is produced by solar panels which are body-mounted on the satellite. Because the satellite orbit is quasi-polar, low-altitude, and helio-synchronous, it alternates day and night periods (eclipse of the Sun by the Earth during night periods). Energy production is only possible during day periods and, within these periods, it depends on the satellite attitude trajectory. Produced energy is stored in batteries whose capacity is limited. See in Fig. 1 a typical energy profile. E_{max} is the maximum feasible energy level, equal to the battery capacity: the energy level cannot be greater than E_{max} . E_{min} is the minimum acceptable energy level: the energy level must not be lower than E_{min} . If we consider an acquisition and download schedule, including the associated instrument switching plan, we can build the associated energy profile and check whether or not the energy level is always greater than or equal to E_{min} .

Ground satellite control stations Ground control stations differ from data reception stations. They are used to send to the satellite acquisition and download schedules to be executed and to receive execution reports.

3 The current acquisition and download scheduling approach

The current approach consists in building offline on the ground an acquisition and download schedule, with precise activity start and end times, which verifies all the constraints related to satellite attitude, to acquisition and download, and to onboard energy, memory, and temperature. This schedule is then downloaded to the satellite using a ground control station visibility window and executed by the satellite without any change.

accepted	~ 14%
rejected - kinematic conflicts	~ 51%
rejected - energy high-level constraints	~ 25%
rejected - temperature high-level constraints	~ 9%
rejected - other constraints	~ 1%

Table 1: Causes of acquisition rejection in the planning process. Evaluation on 6 one-day scenarios from Pléiades satellites. The number of candidate acquisitions varies from 1416 to 2560.

The acquisition and download schedule building process is illustrated in Fig. 2. At the highest level, an acquisition schedule and then a download schedule are proposed, taking into account user request priorities and high-level aggregated constraints such as no more acquisition time than a given threshold over each satellite revolution or each set of consecutive satellite revolutions. The proposed schedules are then checked according to all the constraints. In case of any constraint violation, they are modified by removing for example some acquisitions or downloads. The process continues until a schedule that satisfies all the constraints is found. This process terminates because an empty schedule (no acquisition and no download) is always feasible if we assume a nominal satellite state and behavior.

To be sure that the proposed schedule is really executable by the satellite in spite of the numerous uncertainties described in Sect. 1, safety margins are used when checking constraints. For example, the expected energy production and consumption rates are replaced by minimum production and maximum consumption rates, resulting in harder constraints. These margins are used in high-level constraints. In Tab. 1, we can see that the proportion of accepted acquisitions over an initial pool of acquisitions is low (about 14%). Among the causes of rejection, about 34% are due to energy and temperature high-level constraints. This motivates a more flexible decision-making approach to deal with these constraints.

4 A flexible approach for acquisition scheduling

4.1 Some preliminary design choices

A first approach would be to compute a schedule with little uncertainty when a snapshot of the satellite state is downloaded onto the ground during control station visibility windows. Unfortunately, the scheduling process takes too much time and the final schedule is often verified by a human operator before being sent. The visibility windows would have to be bigger for that approach to work.

In [Maillard *et al.*, 2015], a flexible download scheduling approach has been proposed to deal with uncertainty about the volume of data that is generated by acquisitions and must be downloaded. Roughly speaking, this approach replaces the ground checking of the download and memory constraints by a lighter checking which considers only high priority downloads. For low-priority downloads, decision is delayed online on board and allows the ground download schedule to be modified while guaranteeing ground commitments to high-

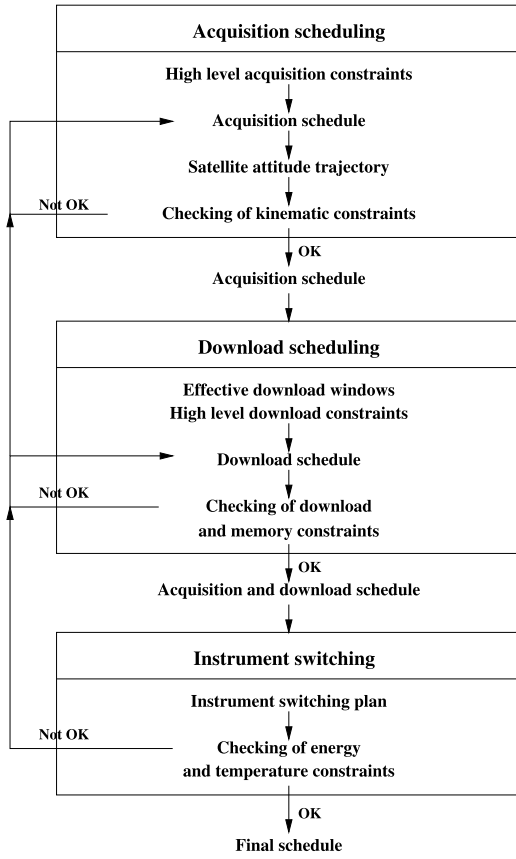


Figure 2: High-level view of the acquisition and download schedule building process.

priority downloads.

Our objective is to extend such an approach by proposing a flexible acquisition scheduling approach to deal with uncertainty about energy production and consumption. Whereas we focus on energy in this paper, the same approach may be applied to temperature. Roughly speaking, this approach would replace the ground checking of the energy constraints by a lighter checking which considers only high-priority acquisitions. For low priority acquisitions, decision would be delayed online on board.

However, acquisition scheduling is far more complex than download scheduling, mainly because of the satellite attitude trajectories to be built and checked. Space engineers consider that it would not be feasible to compute or modify such trajectories online on board, mainly due to the very limited resources on board the satellite (from 1000 to 10000 times less than on a classical ground computer).

In such conditions, we choose not to modify on board the acquisition schedule that has been built on the ground. The only freedom we grant the onboard software is the ability, for any low priority observation in the ground acquisition schedule, to perform it or not. If the onboard software chooses to perform it, the observation instrument must be set ON. If it chooses not to perform it, the observation instrument is set in the SDBY or OFF mode according to the duration between

the current acquisition and the next one in the ground schedule: SDBY mode when this duration is less than or equal to a given threshold, and OFF mode otherwise. Both modes allow energy to be saved.

In the following two sections (Sect 4.2 and 4.3) we explain how flexible acquisition schedules can be built on the ground and then executed on board.

4.2 Building a flexible acquisition schedule on the ground

Algorithm 1: Backward computing of required energy levels

```

1 Flexible-energy-checking( $Seq, EnHs, RE_nHe$ )
2 begin
3    $N \leftarrow |Seq|;$ 
4    $Ren_N \leftarrow RE_nHe;$ 
5    $j \leftarrow N;$ 
6   for  $i = N - 1$  to 1 do
7      $\langle Fea_i, Ren_i \rangle \leftarrow$ 
8       Required-energy-computing( $i, j$ );
9     if  $Fea_i = failure$  then return  $\langle failure, i \rangle;$ 
10    if  $Pr_i = high$  then
11       $j \leftarrow i;$ 
12    else
13       $Mo_i \leftarrow$  Not-performed-mode( $i$ );
14  if  $Ren_1 > EnHs$  then
15    return  $\langle failure, 0 \rangle;$ 
16  else
17    return  $\langle success, 0 \rangle;$ 
  
```

A flexible acquisition schedule is a classical acquisition schedule, that is a sequence Seq of acquisitions, where a required level of energy is associated with each acquisition. The required level of energy associated with an acquisition $i \in Seq$ is the minimum level of energy that is necessary to perform i and all the mandatory activities that follow i , that is all the high-priority acquisitions that follow i in Seq and all the future required activities such as downloads. This approach and the algorithms could be extended to more than two priority levels. For a given acquisition a , we would compute a required level of energy such that all future observations of higher priority can be acquired. For the sake of clarity, we consider only two priority levels here.

Because the download schedule is also flexible, we do not exactly know what will be actually downloaded (other than the certainty that high-priority downloads will be executed). To overcome this difficulty, we make a pessimistic assumption: all the download windows or at least those used to download high-priority acquisitions will be completely used; hence, the antenna will be maintained ON over these windows and SDBY or OFF between these windows, according to the rules presented in Par. **Instrument switching** of Sect.2.

Required levels of energy can be efficiently computed by a backward procedure which browses the sequence Seq of acquisitions from the last to the first. The pseudo-code of

this procedure is given in Algorithm 1 (Function **Flexible-energy-checking**). Its input is made of three elements:

- a sequence Seq of acquisitions over a temporal horizon $[Hs, He]$ at which we add two fictive acquisitions at the beginning and the end of the horizon (Hs and He), whose duration is null and priority is high; acquisitions in Seq , including the two fictive ones, are numbered from 1 to $|Seq|$; an observation instrument switching plan can be deduced from the sequence Seq by following the rules presented in Par. **Instrument switching** of Sect.2;
- a level $EnHs$ of energy at the beginning Hs of the horizon;
- a required level $REnHe$ of energy at the end He of the horizon.

With each acquisition i in Seq are associated four parameters:

- its priority Pr_i , either high or low;
- its starting time St_i ;
- its mode Mo_i , initialized at ON, but possibly modified by the algorithm;
- its required level of energy Ren_i , computed by the algorithm.

An execution of Algorithm 1 is presented in Fig 3. A subset of the sequence of acquisitions is represented from Acq. 4 to Acq. 9. High (resp. low) priority acquisitions are represented by bold (resp. thin) rectangles. For example Acq. 8 is of high-priority and Acq. 9 of low priority. The associated observation instrument switching plan is represented with segments of large (resp. intermediate and small) width for ON (resp. SDBY and OFF) modes. For example, the instrument is maintained in the SDBY mode between Acq. 7 and Acq. 8, but in the OFF mode between Acq. 8 and Acq. 9.

At step 1, we assume that the required level of energy at the beginning of Acq. 8 (of high priority) has been already computed and we compute the required level at the beginning of Acq. 7 (of low priority), assuming the execution of Acqs. 7 and 8 (line 7 of Algorithm 1). A failure is returned when a failure occurs while computing (unfeasible acquisition; line 8). Then, at step 2, we compute the required level at the beginning of Acq. 6 (of low priority), assuming the execution of Acqs. 6 and 8, but not of Acq. 7 which is of low priority and thus not mandatory (the mode of Acq. 7 is set to SDBY; line 12). The same way, at step 3, we compute the required level at the beginning of Acq. 5 (of high priority), assuming the execution of Acqs. 5 and 8, but not of Acqs. 6 and 7 (the mode of Acq. 6 is set to OFF). Because Acq. 5 is of high priority its required level of energy can be used as a reference level to compute required levels for the previous acquisitions (line 10). This is what is done at step 4 where we compute the required level at the beginning of Acq. 4 (of low priority), assuming the execution of Acqs. 4 and 5. When reaching the beginning of the planning horizon, a failure is returned when the required level of energy is greater than the initial level (unfeasible schedule; line 14). Otherwise, a success is returned.

Algorithm 2: Backward computing of required energy levels

```

1 Required-energy-computing( $i, j$ )
2 begin
3    $ren \leftarrow Ren_j$ ;
4    $t \leftarrow St_j$ ;
5   while  $t > St_i$  do
6      $ren \leftarrow \max(Emin, ren - \mathbf{Power}(t) \cdot \Delta t)$ ;
7     if  $ren > Emax$  then return  $\langle failure, ren \rangle$ ;
8      $t \leftarrow t - \Delta t$ ;
9   return  $\langle success, ren \rangle$ ;

```

Function **Not-performed-mode** returns the mode (SDBY or OFF) that has been initially planned just after Acq. i : we assume that, if i is not performed, the observation instrument will be set in that mode from the beginning of Acq. i in order to save energy.

Function **Required-energy-computing** computes the required level of energy at the beginning of Acq. i assuming that the required level of energy at the beginning of Acq. j (of high priority) has been already computed. Such a computation is performed backward step-by-step from the beginning of Acq. j to the beginning of Acq. i . It must be performed step-by-step because the production and consumption of energy are variable. The production depends on the fact that the satellite is in eclipse or not and, if is not in eclipse, on the satellite attitude. The consumption depends on the mode of the instruments (observation instrument and emission antenna) which may be ON, SDBY, or OFF. We assume a step of length Δt , small enough to consider that energy production and consumption are constant over any interval of length Δt .

The pseudo-code of this function is given in Algorithm 2. At each step, the required level of energy at time $t - \Delta t$ is computed from the required level of energy at time t already computed. Function **Power** returns the power produced between $t - \Delta t$ and t (assumed to be constant). This power takes into account the production of the solar panels and the consumption of the observation instrument, of the emission antenna, and of the satellite platform. Hence, it may be positive or negative.

The forward evolution of the energy level en on board the satellite from $t - \Delta t$ to t is given by Eq. 1, where we assume that $Emin \leq en(t - \Delta t) \leq Emax$:

$$en(t) = \min(Emax, en(t - \Delta t) + \mathbf{Power}(t) \cdot \Delta t) \quad (1)$$

This equation takes into account that the energy level cannot be greater than $Emax$. A failure occurs when the energy level is lower than $Emin$ (unacceptable level).

From this forward equation, it is possible to deduce an equation which gives the backward evolution of the required level of energy ren on board the satellite from t to $t - \Delta t$. If we assume that $Emin \leq ren(t) \leq Emax$:




Figure 3: Illustration of the backward procedure used to compute required levels of energy.

$$\begin{aligned}
 ren(t - \Delta t) &= \min(e \mid (\min(Emax, e + \mathbf{Power}(t) \cdot \Delta t) \\
 &\quad \geq ren(t)) \wedge (e \geq Emin)) \\
 &= \min(e \mid (Emax \geq ren(t)) \\
 &\quad \wedge (e + \mathbf{Power}(t) \cdot \Delta t \geq ren(t)) \\
 &\quad \wedge (e \geq Emin)) \\
 &= \min(e \mid (e \geq ren(t) - \mathbf{Power}(t) \cdot \Delta t) \\
 &\quad \wedge (e \geq Emin)) \quad (2) \\
 &= \max(Emin, ren(t) - \mathbf{Power}(t) \cdot \Delta t)
 \end{aligned}$$

Eq. 2 is used at line 6 of Algorithm 2. A failure is returned when the required energy level is greater than $Emax$ (unsatisfiable requirement; line 7).

4.3 Executing a flexible acquisition schedule on board

On board, all high-priority acquisitions are executed. Before each low-priority acquisition i , if the current level of energy is greater than or equal to the required level of energy Ren_i computed on the ground, acquisition i is executed. Otherwise the observation instrument is put into the mode (SDBY or OFF) that has been initially planned just after acquisition i in order to save energy.

5 Experiments

5.1 Approaches to be compared

To assess the positive impact of sharing decision-making between ground and onboard planning, we implemented and compared the following three approaches into a satellite simulator already equipped with various download scheduling methods and algorithms:

1. planning on the ground with coarse energy checking (**GroundCoarse**): acquisition planning is performed on the ground with an algorithm which verifies high-level constraints limiting the acquisition duration over

each orbit to indirectly enforce energy constraints. If the schedule is inconsistent with these constraints, the least valuable acquisitions are removed and the process is started again. On board, the plan is executed without any change. This method is close to the currently operational method.

2. planning on the ground with fine energy checking (**GroundFine**): acquisition planning is performed on the ground with an algorithm which simulates the execution of the schedule step-by-step with maximum power consumption and minimum power production for all acquisitions (high and low-priority). During simulation, when the onboard level of energy is strictly lower than $Emin$, acquisitions are removed from the plan and the process is started again. On board, the plan is executed without any change.
3. planning shared between ground and onboard (**Flexible**): acquisition planning is performed on the ground with an algorithm which computes backward, for all low-priority acquisitions, the level of energy that is required to perform all the future high-priority acquisitions (see Algorithm 1). During backward computing, if the required level of energy is greater than $Emax$, acquisitions are removed from the plan and the process is started again. On board, low-priority acquisitions are executed if and only if the current level of energy is greater than the required level computed on the ground.

5.2 Scenarios

We used one realistic scenario produced by Airbus Defence and Space. This scenario covers one day. It involves 5 users, 2 priority levels (high and low), 3 ground reception stations, 20 visibility windows, and 1364 acquisitions (247 high-priority acquisitions and 1117 low-priority acquisitions).

Uncertainties During day periods, power production depends on the satellite attitude. This attitude $att(t)$ at time t is subject to errors $\epsilon(t) \in \mathbb{R}^2$ on pitch and roll angles. Let $\overline{att(t)}$ be the worst attitude at time t regarding the angle with the optimal attitude $hpAtt(t)$ where solar panels are directed towards the Sun:

$$\overline{att(t)} = \arg \max_{\|a-att(t)\| \leq \epsilon(t)} (angle(a, hpAtt(t))) \quad (3)$$

The minimum power production $Pp(t)$ at time t is given by the following equation:

$$Pp(t) = \max(0, P_{sun} \cdot \cos(angle(\overline{att(t)}, hpAtt(t)))) \quad (4)$$

We also consider that power consumptions by the platform, the observation instrument and the emission antenna are uncertain. We evaluated the same scenario with three uncertainty hypotheses: maximum consumptions 10%, 20%, or 30% higher than the nominal consumptions.

5.3 Results and discussion

Priority	Uncertainty	GroundCoarse	GroundFine	Flexible
High	10%	247	247	247
	20%	247	247	247
	30%	247	247	247
Low	10%	976	1053	1086
	20%	938	1010	1061
	30%	898	927	1040

Table 2: Number of acquisitions that are really executed.

We measured only the number of executed acquisitions. In Tab. 2, we see that whatever the approach, all the high-priority acquisitions are executed. Concerning the low-priority acquisitions, we can see that the Flexible approach allows the highest number of low-priority acquisitions to be executed. Without any surprise the GroundFine approach which simulates finely the energy profile is more effective and allows more low-priority acquisitions to be executed than the GroundCoarse approach which enforces stronger high-level constraints.

Fig. 4 shows the onboard energy levels during a 24-hour scenario from different perspectives. The GroundFine method computes a first predicted energy profile (dotted line on the figure) based on worst energy consumptions and productions. A procedure then removes acquisitions until this profile does not go under E_{min} , that is 85% in this case. We see that the predicted and real profiles produced by the Flexible approach are quite close at the beginning and at the end of the day because acquisition and download activities mainly occur between 6am and 6pm in this scenario.

Coupling with the data download planning The acquisition scheduling problem is strongly connected with the data download problem (see Fig. 2). In the flexible approach, we made the strong assumption that all visibility windows were used for downloading. Whereas this is close to reality in over-subscribed scenarios such as defense scenarios where there

are few ground stations and then few visibility windows, it is far from true in civil-oriented scenarios where there are a large number of ground stations that cover all continents and are generally underused. This assumption makes the acquisition schedule very pessimistic. Work remains to be done to use results from the download planning phase during the computation of required energy levels. One possibility would be to consider that downloading high-priority acquisitions is the only mandatory requirement. In that case, only the visibility windows needed for that would be taken into account when simulating the schedule. Following the same Flexible approach, required energy levels would be computed on the ground for each group of visibility windows (a group constitutes an onboard download planning horizon). On board, before any of these groups, if the actual level of energy is sufficient, then the whole group could be used for downloading. Otherwise, only high-priority downloads would be executed to save energy.

6 Related Work

The problem of selecting and scheduling acquisitions for Earth observation satellites has been extensively studied (see the survey by [Globus *et al.*, 2004]). This problem can be viewed as an oversubscribed scheduling problem [Barbulescu *et al.*, 2006] where, in addition to scheduling, acquisitions have to be selected.

Selecting and scheduling acquisitions for agile Earth-observation satellites has been studied by [Grasset-Bourdel *et al.*, 2012] and [Beaumont *et al.*, 2011]. One of its particularities is the presence of time-dependent constraints due to the platform agility (activity or transition durations dependent on their start time [Gawiejnowicz, 2008; Pralet and Verfaille, 2013]). Authors have highlighted that this problem has features from AI (for the hierarchical decision process or the planning subproblem) and OR (for the pure scheduling subproblem).

However, these approaches are often pure ground or pure onboard approaches and, whereas energy is always mentioned, the uncertainty about it is not explicitly taken into account. For any kind of continuous resource (including time), handling uncertainty remains a challenge in AI applications [Bresina *et al.*, 2002].

From a more theoretical point of view, the approach proposed in this paper is close to the one developed by [Fox and Long, 2002] and [Gough *et al.*, 2004]. In this latter approach, which is an alternative to the computationally expensive *contingent planning*, deterministic plans with conservative estimates of a continuous resource are built. Opportunistic loop branches are then developed at some steps of the plan, based on the assumption that the continuous resource level may be higher at execution than predicted when planning. A loop branch will be executed if and only if the current level of resource is sufficient. This approach, like ours, allows wastage of resource to be avoided.

7 Conclusion

In this paper, we showed how uncertainty about energy levels can be managed by using a mixed architecture where

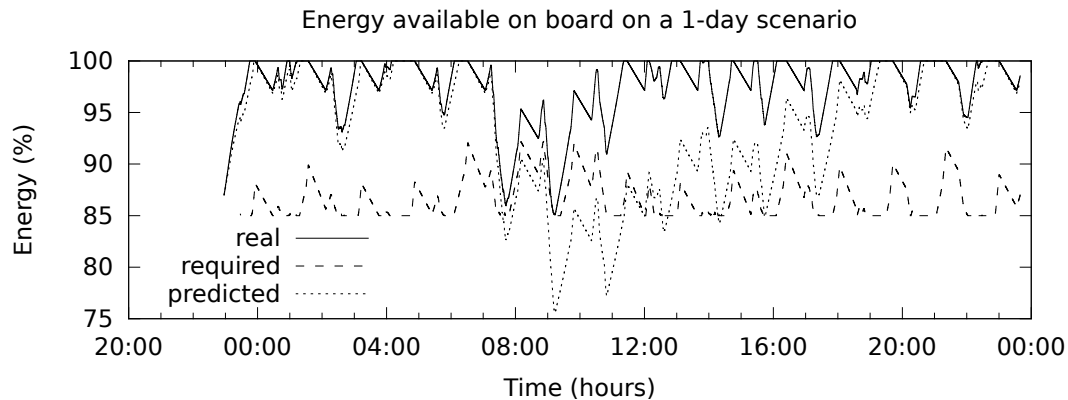


Figure 4: Satellite energy profiles (in percentage of the maximum energy level E_{max}) over a 24-hour scenario ($E_{min} = 0.85 \cdot E_{max}$). The dotted line is the first ground prediction with margins produced by the GroundFine approach. The dashed line represents the required levels of energy computed on the ground by the Flexible approach. The solid line represents the actual energy profile resulting from the onboard execution of the Flexible approach.

decision-making about energy consumption is postponed as much as possible. Commitments to high-priority acquisitions are guaranteed. For each low-priority acquisition, a backward procedure computes a lower bound on the level of energy required to complete this acquisition and all future high-priority acquisitions. On board, low-priority acquisitions are executed or not by comparing the lower bound to the current level of energy available. We compared this approach with the ones currently in use and showed that it avoids wastage of resource and allows more acquisitions to be executed. Future work includes extending this mechanism to others parameters such as temperature, coupling flexible acquisition planning with flexible download planning, and generalizing this approach to more generic problems. This is an early-stage design study: implementing and testing this approach for a real satellite mission would demonstrate its interest.

References

- [Barbulescu *et al.*, 2006] L. Barbulescu, A. Howe, L. Whitley, and M. Roberts. Understanding Algorithm Performance on an Oversubscribed Scheduling Application. *Journal of Artificial Intelligence Research*, 27:577–615, 2006.
- [Baumet *et al.*, 2011] G. Baumet, G. Verfaillie, and M.C. Charneau. Feasibility of Autonomous Decision Making on board an Agile Earth-observing Satellite. *Computational Intelligence*, 27(1):123–139, 2011.
- [Bresina *et al.*, 2002] J. Bresina, R. Dearden, N. Meuleau, S. Ramkrishnan, D. Smith, and R. Washington. Planning under Continuous Time and Resource Uncertainty: A Challenge for AI. In *Proc. of the 18th International Conference on Uncertainty in Artificial Intelligence (UAI-02)*, pages 77–84, Edmonton, Alberta, Canada, 2002.
- [Chien *et al.*, 2004] S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, R. Lee, D. Mandl, S. Frye, B. Trout, J. Hengemihle, J. D’Agostino, S. Shulman, S. Ungar, T. Brakke, D. Boyer, J. VanGaasbeck, R. Greeley, T. Doggett, V. Baker, J. Dohm, and F. Ip. The EO-1 Autonomous Science Agent. In *Proc. of the 3rd Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-04)*, pages 420–427, New York City, USA, 2004.
- [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.
- [Gawiejnowicz, 2008] S. Gawiejnowicz. *Time-dependent Scheduling*. Springer, 2008.
- [Globus *et al.*, 2004] A. Globus, J. Crawford, J. Lohn, and R. Morris. A Comparison of Techniques for Scheduling Earth Observing Satellites. In *Proc. of the 16th Conference on Innovative Applications of Artificial Intelligence (IAAI-04)*, San Jose, CA, USA, 2004.
- [Gough *et al.*, 2004] J. Gough, M. Fox, and D. Long. Plan execution under resource consumption uncertainty. In *Proc. of the ICAPS Workshop on Connecting Planning Theory with Practice*, Whistler, British Columbia, Canada, 2004.
- [Grasset-Bourdel *et al.*, 2012] R. Grasset-Bourdel, G. Verfaillie, and A. Flipo. Action and Motion Planning for Agile Earth-observing Satellites. *Acta Futura*, 5:121–131, 2012.
- [Maillard *et al.*, 2015] A. Maillard, C. Pralet, J. Jaubert, I. Sebbag, F. Fontanari, and J. 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] C. Pralet and G. Verfaillie. Time-dependent Simple Temporal Networks: Properties and Algorithms. *RAIRO Operations Research*, 47(2):173–198, 2013.