A generic constraint-based local search library for the management of an electromagnetic surveillance space mission

Cédric Pralet and **Guillaume Infantes** and **Gérard Verfaillie**
ONERA - The French Aerospace Lab, F-31055, Toulouse, France
{Cedric.Pralet,Guillaume.Infantes,Gerard.Verfaillie}@onera.fr

Abstract

This paper presents what has been done at the French Aerospace Lab (ONERA) to deal with a scenario of space mission defined by the French Space Agency (CNES). This space mission is dedicated to the surveillance from space of ground electromagnetic sources. It involves two satellites: one for source detection and another one for data acquisition and download. It presents two sources of uncertainty: the presence or not of electromagnetic sources and, in case of presence, the volume of data generated by acquisition. Due to these uncertainties and to limited communication windows with ground control stations, online planning and scheduling (P&S) is necessary on board the second satellite to make consistent and optimal decisions in terms of data acquisition and download. In this paper we show how a generic constraint-based local search library can be used to build the onboard planning and scheduling component. This library, called *InCELL*, has been developed at ONERA. It allows temporal constraints, resource constraints, arithmetic and logical constraints, and optimization criterion to be quickly and incrementally evaluated at each step of a local search algorithm. Already experimented to deal with simpler scenarios, this is the first time it is experimented on a complex scenario involving agile satellites. We show also how the generic simulation tool *Ptolemy* can be used to simulate the space system and evaluate its P&S component.

Introduction

In the context of the CNES-ONERA *Agata* project about spacecraft autonomy (Charmeau and Bensana 2005), after working on a first mission scenario involving only one non agile Earth optical detection and observation satellite (Damiani, Verfaillie, and Charmeau 2004; Pralet and Verfaillie 2008), ONERA dealt with a more complex mission scenario defined by CNES and called *Agata-One*. The main objective was to assess whether or not the tools that were defined to deal with the first scenario can be easily adapted to deal with a more complex one.

The *Agata-One* scenario involves two agile Earth satellites placed on low altitude, circular orbits, on the same orbital plane. Agile satellites are able to perform very quick attitude movements along the three axes around their gravity center (roll, pitch, and yaw) generally thanks to gyroscopic actuators which are more efficient than usual reaction wheels. Thanks to regular roll attitude movements, the first satellite (S_1) scans a wide strip around its ground track. Thanks to its instruments, it is able to detect the presence of electromagnetic sources at the Earth surface and to localize them. In case of detection, it sends instantaneously information to the second satellite (S_2, which follows it at a small distance) via a permanent inter-satellite low-rate communication link. Satellite S_2 maintains a set of ground areas on which electromagnetic sources have been detected. Each time it overflies one of these areas, it can acquire data from it. To do that, it must perform a roll and pitch attitude movement to direct its acquisition instrument (a reception antenna) towards this area (the reception antenna is body-mounted on the satellite). When too many close areas must be handled, it must decide on those it will effectively handle and on the acquisition order. Once data from an area has been acquired, it is memorized in a mass memory and downloaded to ground reception and processing stations via a non permanent satellite-ground high-rate communication link. Downloading data to a ground station is only possible within one of the station visibility windows. Moreover, it is only possible when the satellite attitude is compatible with data download (as the reception antenna, the emission antenna is body-mounted on the satellite and, during the whole download period, the station must remain inside the satellite emission antenna cone). As for data acquisition, when too much data must be downloaded, satellite S_2 must decide on those it will download and on the download order. Fig. 1 summarizes the exchanges between satellites and ground reception stations.

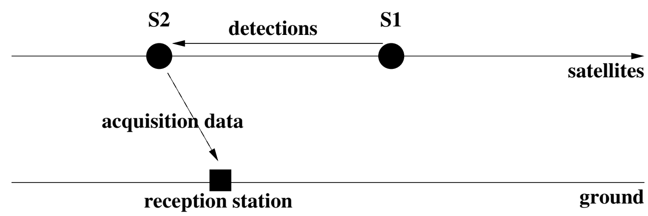

Figure 1: Exchanges between satellites and ground reception stations.

It must be stressed that the attitude of satellite S_2 allow-

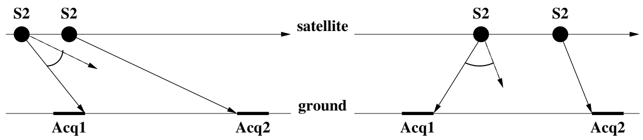

Figure 2: How the attitude movement to be performed by satellite S_2 to transit from a data acquisition to another one depends on the time at which the transition is triggered. In the second case (right), the angular movement to be performed is greater than in the first case (left).

ing it to direct its reception antenna towards a given area depends on the position of the satellite on its orbit and thus on time. In such conditions, the attitude movement necessary to transit from a data acquisition from a given area to a data acquisition from another area, and thus the time taken by this transition, depends not only on both areas, but also on the time at which the transition is triggered (time-dependent transition duration). See Fig. 2 for an illustration.

This mission scenario involves two main sources of uncertainty: the presence or not of electromagnetic sources and, in case of presence, the volume of data generated by acquisition (this volume is highly variable and can typically range from 1 to 1000). Due to these uncertainties and to the non permanent visibility of satellites by ground control stations, online decision-making on data acquisition and download is necessary on board satellite S_2. To make such decisions, it would be possible to use manually defined decision rules. However, decisions would be better informed if they could use the result of P&S: planning and scheduling regularly performed over a given horizon ahead, using the most up to date information about detections and data volumes; decisions made according to the first steps of the plans produced.

Due to the limited computing time available for P&S and due to the limited computing resources available on board (CPU and RAM), heuristic search (greedy and/or local search) seems to be the right option to build an anytime combinatorial search procedure, able to produce quickly good quality plans and to improve on them as long as time is available before making decisions. It is widely used in space missions that require online onboard P&S (Chien et al. 2000; 2005b; 2005a). We already used it in the context of Earth observation and surveillance missions (Lemaître et al. 2002; Beaumet, Verfaillie, and Charmeau 2011; Pralet and Verfaillie 2008; Pralet et al. 2011; Verfaillie et al. 2011). However, each time, we built specific heuristic search procedures, dedicated to the specific mission at hand and not directly reusable to handle other missions. To deal with the *Agata-One* scenario, we decided to change our approach and to use generic tools developed at ONERA in the context of the *Agata* project and, more specifically, the *Invariant-based Constraint EvaLuation Library* (*InCELL* (Pralet and Verfaillie 2013)).

InCELL draws its inspiration from the ideas of *Constraint-based local search* (CLS (Hentenryck and Michel 2005)). In CLS, the user defines a model of its problem in terms of decision variables, constraints, and optimization criterion. She/he defines also its local search procedure over the set of complete variable assignments (where every variable is assigned). Because the speed of each local move is one of the keys to local search success, the software uses so-called *invariants* which allow expressions and constraints to be quickly and incrementally evaluated after each move. In *InCELL*, multiple-input multiple-output invariants allow expressions, arithmetic and logical constraints, temporal and resource constraints to be expressed and efficiently handled. *InCELL* calls for *Simple Temporal Network* (STN (Dechter, Meiry, and Pearl 1991)) techniques which allow temporally flexible plans to be produced, and for *Time-dependent* STN (TSTN (Pralet and Verfailllie 2012)) techniques which allow time-dependent transition durations to be taken into account.

To deal with the *Agata-One* scenario, an *InCELL* model of the associated P&S problem (decisions about data acquisition and download by satellite S_2) was built, a simple non chronological greedy search procedure was designed, and the events that trigger a new call to P&S over a given horizon ahead were defined.

To simulate the space system and to evaluate its P&S component, an event-based model of the system, based on the notions of state, event preconditions and effects, and event activations, was built and implemented using the generic simulation tool *Ptolemy* (Eker et al. 2003). Whereas P&S allows only the utility of decisions over the planning horizon to be evaluated, this simulation allows the global utility of successive decisions over the simulation horizon to be evaluated.

Sect. 1 describes problem data and Sect. 2 presents the structure of possible decisions. In Sect. 3, a constraint-based model of the P&S problem is introduced. The main ingredients of the *InCELL* library, as well as its main reasoning mechanisms, are presented in Sect. 4. Sect. 5 describes the search procedure and Sect. 6 defines when P&S is called. Sect. 7 shows how the space system and its P&S component can be simulated and evaluated, using the *Ptolemy* tool.

1 Problem data

Permanent (static) problem data is the following:

- a finite set of ground areas that must be kept under surveillance;
- a finite sequence of priority levels;
- for each ground area, its priority level, its weight (to give more or less weight to areas of the same priority level), an acquisition duration, an expected, a minimum, and a maximum volume of data resulting from acquisition;
- a finite set of ground reception stations;
- a data download rate from satellite S_2 to any ground reception station.

Moreover, it is assumed that a function associates with each ground area a and each time t the attitude of satellite S_2 necessary to acquire data from a at time t, when acquisition is possible, and that another function associates with each pair of attitudes of satellite S_2 the minimum time necessary to reach the second one, starting from the first one.

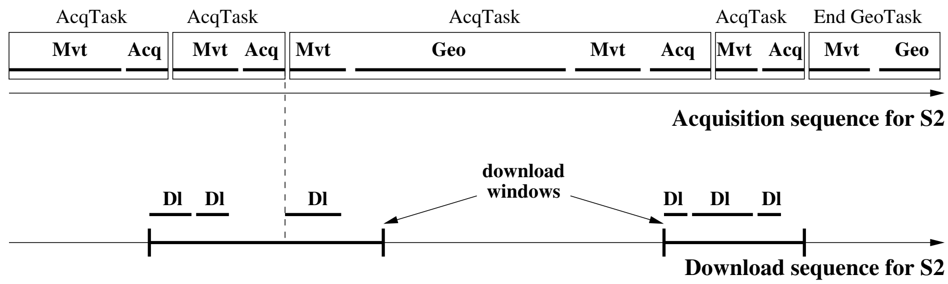

Figure 3: The two concurrent sequences of action on board satellite S_2. AcqTask = acquisition task; Geo Task = geocentric task; Mvt = attitude movement; Acq = data acquisition; Geo = geocentric pointing; Dl = data download.

Each time P&S is called, its complementary (dynamic) data is the following:

- a planning horizon ahead;
- an attitude of satellite S_2 at the beginning of the planning horizon;
- a set of ground areas from which electromagnetic sources have been detected by satellite S_1, but no acquisition by satellite S_2 has been performed yet;
- for each of these ground areas, its detection time and a finite sequence of acquisition windows by satellite S_2 over the planning horizon;
- a finite set of acquisitions that have been already performed, but whose data has not been downloaded yet (still present in memory);
- for each of these acquisitions, its detection and acquisition times and its actual volume in memory;
- a finite sequence of download windows by satellite S_2 over the planning horizon.

Acquisition windows are reduced in case of intersection with a download window, when acquisition is incompatible with download, in order to give priority to data download.

2 Possible decisions

On board satellite S_2, it is necessary to decide on two concurrent sequences of action:

- the sequence of data acquisitions;
- the sequence of data downloads.

The first sequence is made of acquisition tasks, each one following immediately the previous one. An acquisition task is, according to an HTN-like decomposition of tasks into sub-tasks (*Hierarchical Task Networks* (Nau et al. 2003)), itself made of:

- either an attitude movement immediately followed by a data acquisition;
- or an attitude movement immediately followed by a geocentric pointing, immediately followed by another attitude movement, immediately followed by a data acquisition (satellite geocentric pointing maintained towards Earth center is a waiting action, favourable to communication with Earth, data downloads, and energy recharging).

This sequence, possibly completed by an attitude movement followed by a geocentric pointing at the end of the planning horizon, entirely defines the attitude trajectory of satellite S_2.

The second sequence is made of data downloads, each one being performed within a download window. In this sequence, a download may not immediately follow the previous one. This is the case when it is necessary to wait for the end of an acquisition before downloading resulting data within a download window.

Fig. 3 illustrates the two concurrent sequences of action. Both sequences are not independent from each other because data download requires preceding acquisition.

3 A constraint-based model

P&S problem is a kind of over-constrained scheduling problem (over-constrained because it may be impossible to schedule all the candidate tasks (Kramer and Smith 2003)) which can be modeled using only constraints over intervals. An interval is defined by its presence, its starting date, its ending date, and its duration. Its presence is a boolean, equal to 1 if and only if the interval is effectively present in the schedule.

The model associates:

- with each ground area from which electromagnetic sources have been detected by satellite S_1, but no acquisition by satellite S_2 has been performed yet, an acquisition interval, a geocentric pointing interval, and a download interval;
- with each acquisition that has been already performed by satellite S_2, but whose data has not been downloaded yet, a download interval.

These intervals may be present or absent. Constraints to be satisfied are the following:

- each acquisition interval must be, when present, included in one of the acquisition windows of the associated ground area; its duration is the acquisition duration of the associated ground area, defined in the problem data;
- each download interval must be, when present, included in one of the download windows; if the acquisition has been already performed at the P&S time, download duration is equal to the actual volume in memory divided by

the download rate; if it has not been performed yet, it is equal to the maximum volume resulting from acquisition divided by the download rate (pessimistic assumption allowing the produced schedule to be surely executed);

- for each ground area from which electromagnetic sources have been detected, absence of the acquisition interval implies absence of the geocentric pointing and download intervals; presence of the geocentric pointing interval implies that it must precede the acquisition interval and follow the previous acquisition interval; presence of the download interval implies that it must follow the acquisition interval;
- there must be no overlapping between present acquisition and geocentric pointing intervals and enough time between successive intervals to allow attitude movements; moreover movements to or from geocentric pointings must be performed in minimum time in order to give geocentric pointing as much time as possible;
- there must be no overlapping between present download intervals.

The criterion to be optimized is a vector of global utilities, one per priority level. The global utility associated with a priority level p is equal to the sum of the local utilities associated with each of the ground areas of priority p. The local utility associated with a ground area is equal to its weight multiplied by two functions which both take a value between 0 and 1: a decreasing function of the time between detection and acquisition and another decreasing function of the time between acquisition and download. These functions tend to encourage quick acquisition and quick delivery of information on the ground. Two vectors of global utilities, resulting from two schedules, are lexicographically compared from the highest priority level to the lowest one.

4 The InCELL library

InCELL (*Invariant-based Constraint EvaLuation Library*) is a software library, dedicated to the quick incremental evaluation of expressions and constraints.

InCELL draws its inspiration from the ideas of *Constraint-based local search* (CLS (Hentenryck and Michel 2005)). In CLS, the user defines a model of its problem in terms of decision variables, constraints, and optimization criterion. She/he defines also its local search procedure over the set of complete variable assignments (every variable assigned). Because the speed of each local move is one of the keys to local search success, the software uses so-called *invariants* which allow expressions and constraints to be quickly and incrementally evaluated after each move. An invariant is a one-way constraint of the form $x \leftarrow exp$, where x is a variable and exp a function of other variables, such as for example $x \leftarrow \sum_{i=1}^{N} y_i$. On this example, when the value of y_j for some j is modified, it is not necessary to recompute $\sum_{i=1}^{N} y_i$ from scratch. It suffices to add to the previous value of x the new value of y_j, minus its old value. The only condition is the absence of cycles in the definition of invariants (no variable directly or indirectly function of itself).

InCELL extends the definition of invariants by allowing multiple-input multiple-output invariants. Invariants allow expressions, but also constraints, to be represented. Constraints, such as for example $\sum_{i=1}^{N} y_i \leq K$, are specific invariants whose evaluation stops when they are violated. In *InCELL*, a constraint optimization problem (variables, constraints, and criterion) takes the form of a DAG (*Directed Acyclic Graph*) of invariants. Each time the value of some atomic variables (variables that are not functions of other variables and are roots of the DAG) is modified, the DAG of invariants is lazily reevaluated according to a DAG topological order: any invariant is reevaluated only when necessary and at most once.

On top of these basic concepts and mechanisms, *InCELL* offers some constructs dedicated to scheduling: time point variables, interval variables (defined by two time point variables and a distance constraint between them), unary and binary distance constraints (of the form $x \leq K$ or $x - y \leq K$). All temporal constraints are managed using a special STN invariant (*Simple Temporal network* (Dechter, Meiry, and Pearl 1991)) which has as inputs a set of unary and binary distance constraints and as outputs the earliest dates of all the time point variables involved in the constraints. Classical STN techniques are used to handle the STN: constraint propagation, maintenance of propagation chains, decomposition of the distance graph into strongly connected components. Moreover, STN concepts and techniques are extended in *InCELL* to deal with so-called *time-dependent scheduling* (Gawiejnowicz 2008), that is with time-dependent distance constraints (Pralet and Verfailllie 2012) where the minimum distance is not a constant, but a function of the involved time points (of the form $x - y \leq F(x, y)$ with some assumptions about Function F). All the constraints over intervals, defined in the previous section, can be managed by *InCELL*, including the minimum transition times between successive acquisition and geocentric pointing intervals, thanks to time-dependent distance constraints.

InCELL allows also resource constraints to be defined and profiles of resources (with piecewise constant or linear evolutions) to be quickly and incrementally maintained, taking into account the earliest dates produced by the STN. This would allow memory (piecewise constant evolution) and energy (piecewise linear evolution) constraints to be managed. However, these constraints are ignored in our problem: energy because it is not limiting and memory because of the uncertainty about the volume of data generated by acquisition. When planning acquisitions, we prefer not to limit acquisitions because of possible large volumes of data. However, when executing the acquisition plan, before triggering an acquisition, in case of possible memory overflow, we remove from memory lower priority data and, when it is not sufficient, we cancel the acquisition.

One of the key features of *InCELL* is its ability to work on dynamic models (a new model each time P&S is called) using a unique static model which is recycled to build dynamic models. This allows any dynamic memory allocation to be avoided on board: a key requirement when building embedded reactive control software.

See (Pralet and Verfaillie 2013) for more details about the

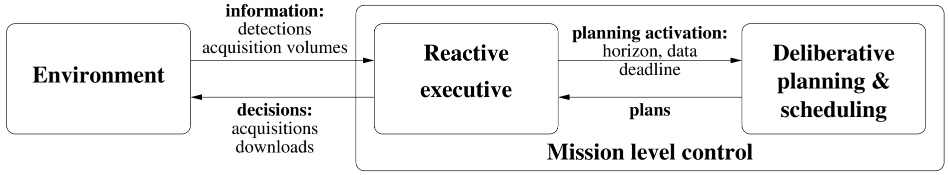

Figure 4: Exchanges between the environment, the reactive executive, and the deliberative P&S component.

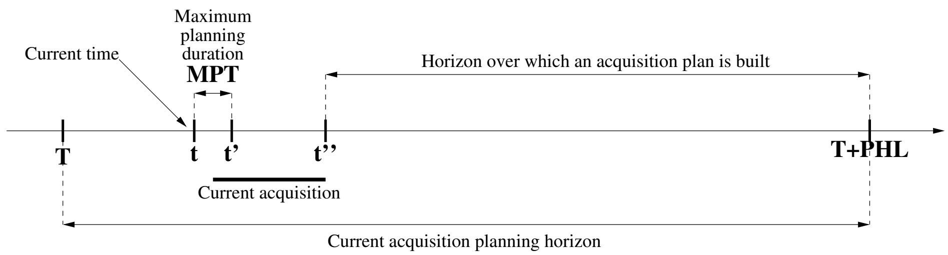

Figure 5: Horizon over which an acquisition plan is built at a given time t on board satellite S_2.

InCELL library.

5 A non chronological greedy search

To build online acquisition and download plans on board satellite S_2, we defined a very simple greedy search procedure, although more sophisticated local search procedures could be considered (Aarts and Lenstra 1997).

At each step of this procedure, an acquisition (resp. download) of highest priority level and of highest utility at this priority level is selected and added to the acquisition (resp. download) plan, when addition is possible. It is added in the best acquisition (resp. download) window and at the best position in the acquisition (resp. download) sequence in terms of utility. Once the acquisition sequence is defined, geocentric pointings are added between acquisitions, when possible.

6 Calls to planning and scheduling

In case of online P&S, it is not only necessary to define the P&S model and the reasoning and search mechanisms. It is necessary to define when the executive calls to P&S, in order to get a plan over some planning horizon ahead and to follow it until a new call to P&S. See Fig. 4 for a global view of the exchanges between the environment, the reactive executive, and the deliberative P&S component.

In our problem, as far as acquisitions are concerned, we define the length PHL of the planning horizon (horizon over which P&S is called; typically some hours) and the maximum planning time MPT (maximum time taken by P&S; typically some seconds). The planning horizon of length PHL is regularly shifted (typically every half an hour). P&S is called again when a new acquisition opportunity appears over the planning horizon. This happens either when electromagnetic sources are detected by satellite S_1 on some ground area, or when the planning horizon is shifted and a new acquisition window for some ground area appears over the new planning horizon. In such a case, we consider the current time t, the time $t' = t + MPT$ at which a plan will be surely available, the time t'' from which decisions can be made, taking into account acquisition or attitude movement possibly in progress at t' (acquisitions and attitude movements are not interruptible, but geocentric pointings are), and we call to P&S over the planning horizon from t''. See Fig. 5 for an illustration.

As far as downloads are concerned, P&S is called MPT before each download window (or group of windows that overlap or are very close to each other) over the whole window (or group of windows).

7 Simulation

We used the simulation tool *Ptolemy* to simulate the space system. *Ptolemy* (Eker et al. 2003) (see http://ptolemy.eecs.berkeley.edu/) is a generic tool dedicated to the simulation of dynamic systems, with an emphasis on hybrid simulation. Among many other possibilities, it is possible within *Ptolemy* to simulate a system whose dynamics involves both discrete events and continuous evolutions of resources. To express in *Ptolemy* the dynamics of the space system, we particularly relied on the *Ptera* framework (Feng, Lee, and Schruben 2010) which is based on the notions of state, events, event preconditions and effects, and conditional activations by events of other events (possibly with some delay and some probability). See Fig. 6 for an illustration of the several temporal horizons that are handled in the simulation (simulation, commitment, planning, and decision horizons).

This simulation was run on scenarios built by CNES. Fig. 7 shows a screenshot of the simulation tool at the end of a five day simulation horizon, where one can see:

- at the top left, the current acquisition requests over the whole world (small circles) and the visibility circle of the unique ground reception station (in blue);
- at the top right, an artist view of satellites S_2 (in front) and S_1 (behind) with the pointing direction of the former;

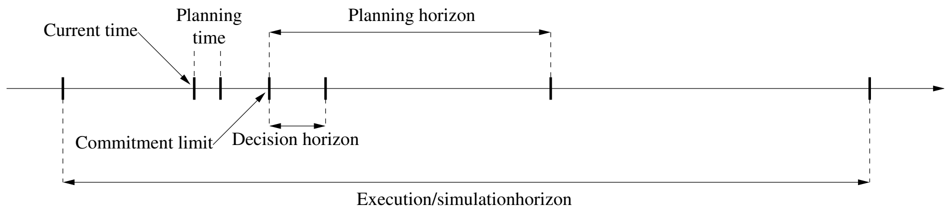

Figure 6: Illustration of the several temporal horizons that are handled in the simulation.

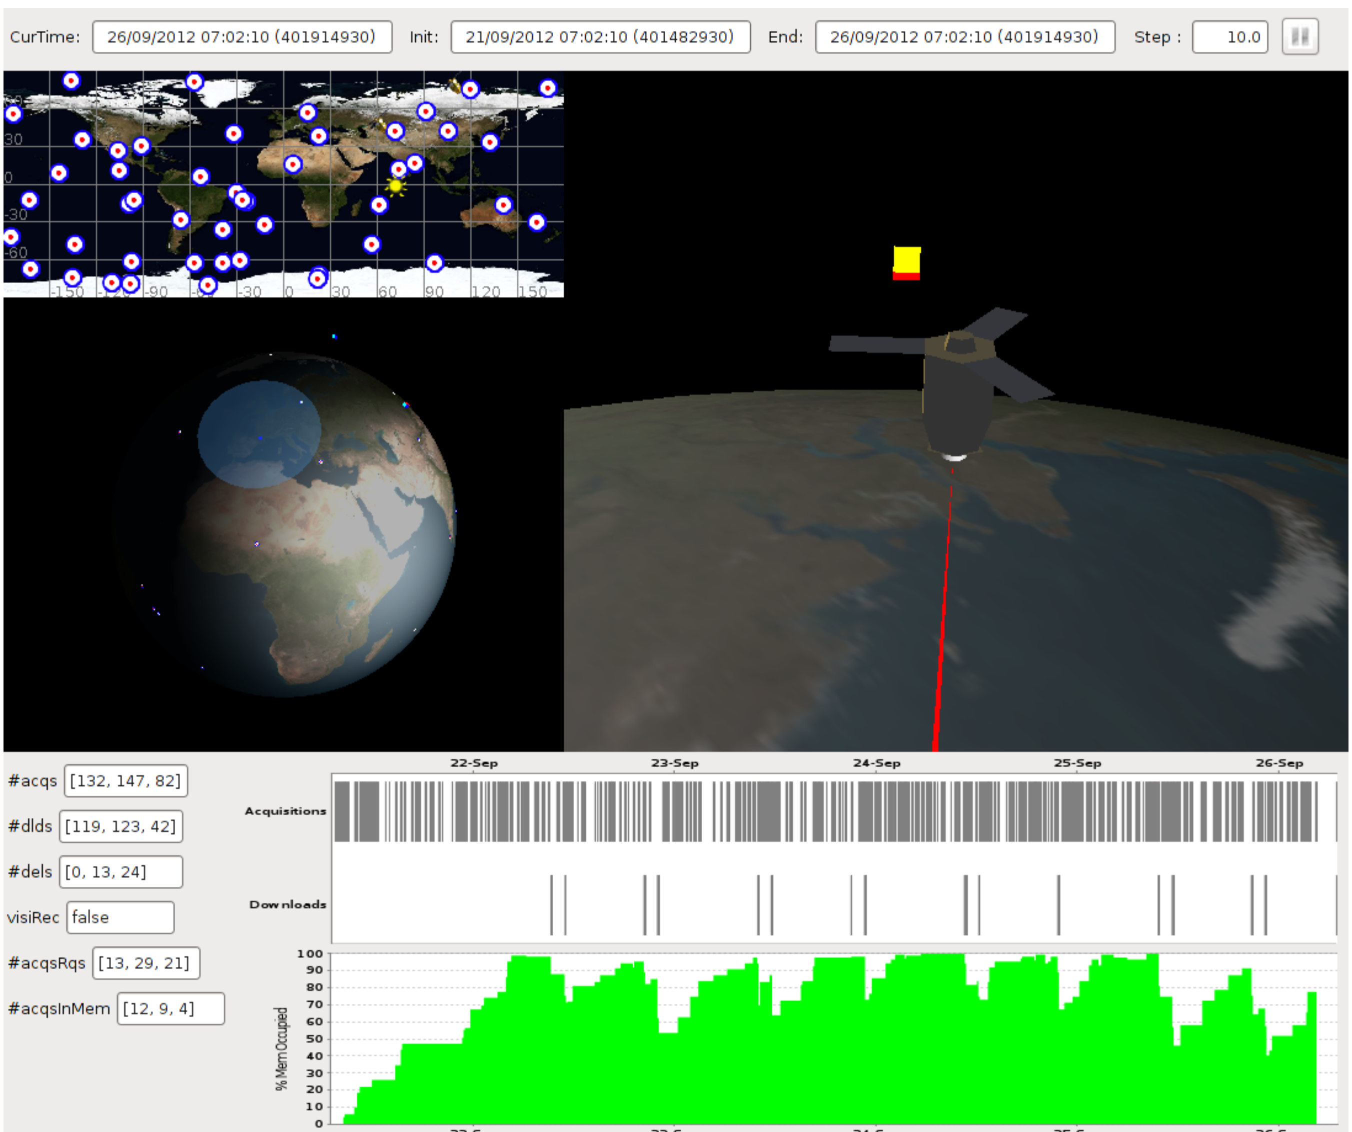

Figure 7: Screenshot of the *Agata-One* simulator at the end of a five day simulation horizon.

Prio	1	2	3
NAcq	132	147	82
NDl	119	123	42
NRm	0	13	24

Table 1: Global results over the five day simulation horizon: **Prio** = priority level, **NAcq** = number of performed acquisitions; **NDl** = number of downloaded acquisitions, **NRm** = number of performed acquisitions that have been removed from memory to free space.

- at the bottom right, the sequence of acquisitions, the sequence of downloads, and the evolution of memory on board.

Over this simulation horizon, acquisition planning is called 680 times, each time over a four hour horizon ahead. Download planning is called 16 times, each time over the next download window. Each time it is called, acquisition (resp. download) planning must manage some tens of acquisitions to be performed (resp. downloaded). Acquisition (resp. download) planning takes on average 432 ms (resp. 2258 ms) on an i5-520 Intel processor with 1.2 GHz and 4 GBRAM.

The global results per priority level are shown on Tab. 1 The mean utilization percentage of the downloads windows is of 85.39%.

A demonstration of the space system simulation is presented in the ICAPS 2013 Application Showcase.

Conclusion

The first result of this study is the demonstration that the generic *InCell* library allows the planning problem associated with a new complex space mission to be easily modeled and efficiently solved. Beyond the necessary improvements of the library in terms of modeling power and algorithm efficiency, the next steps should be the management of other missions, the implementation of the executive and of the P&S component on actual space processors, and the effective use on board an autonomous spacecraft, for example to manage data downloads in presence of uncertainty about volumes.

References

Aarts, E., and Lenstra, J., eds. 1997. *Local Search in Combinatorial Optimization*. John Wiley & Sons.

Beaumet, G.; Verfaillie, G.; and Charmeau, M. 2011. Feasibility of Autonomous Decision Making on board an Agile Earth-observing Satellite. *Computational Intelligence* 27(1):123–139.

Charmeau, M.-C., and Bensana, E. 2005. AGATA: A Lab Bench Project for Spacecraft Autonomy. In *Proc. of the 8th International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS-05)*.

Chien, S.; Knight, R.; Stechert, A.; R.Sherwood; and Rabideau, G. 2000. Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling. In *Proc. of the 5th International Conference on Artificial Intelligence Planning and Scheduling (AIPS-00)*, 300–307.

Chien, S.; Cichy, B.; Davies, A.; Tran, D.; Rabideau, G.; Castano, R.; Sherwood, R.; Mandl, D.; Frye, S.; Shulman, S.; Jones, J.; and Grosvenor, S. 2005a. An Autonomous Earth-Observing Sensorweb. *IEEE Intelligent Systems* 20(3):16–24.

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davies, A.; Mandl, D.; Frye, S.; Trout, B.; Shulman, S.; and Boyer, D. 2005b. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication* 2:196–216.

Damiani, S.; Verfaillie, G.; and Charmeau, M.-C. 2004. An Anytime Planning Approach for the Management of an Earth Watching Satellite. In *Proc. of the 4th International Workshop on Planning and Scheduling for Space (IWPSS-04)*.

Dechter, R.; Meiry, I.; and Pearl, J. 1991. Temporal Constraint Networks. *Artificial Intelligence* 49:61–95.

Eker, J.; Janneck, J.; Lee, E.; Liu, J.; Liu, X.; Ludvig, J.; Neuendorffer, S.; Sachs, S.; and Xiong, Y. 2003. Taming Heterogeneity: the Ptolemy Approach. *Proceedings of the IEEE* 91(1):127–144.

Feng, T.; Lee, E.; and Schruben, L. 2010. Ptera: An Event-oriented Model of Computation for Heterogeneous Systems. In *Proc. of the 10th International Conference on Embedded Software (EMSOFT-10)*, 219–228.

Gawiejnowicz, S. 2008. *Time-dependent Scheduling*. Springer.

Hentenryck, P. V., and Michel, L. 2005. *Constraint-based Local Search*. MIT Press.

Kramer, L., and Smith, S. 2003. Maximizing Flexibility: A Retraction Heuristic for Oversubscribed Scheduling Problems. In *Proc. of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, 1218–1223.

Lemaître, M.; Verfaillie, G.; Jouhaud, F.; Lachiver, J.-M.; and Bataille, N. 2002. Selecting and scheduling observations of agile satellites. *Aerospace Science and Technology* 6:367–381.

Nau, D.; Au, T.; Ilghami, O.; Kuter, U.; Murdock, W.; Wu, D.; and Yaman, F. 2003. SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research* 20:379–404.

Pralet, C., and Verfaillie, G. 2008. Decision upon Observations and Data Downloads by an Autonomous Earth Surveillance Satellite. In *Proc. of the 9th International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS-08)*.

Pralet, C., and Verfaillie, G. 2013. Dynamic Online Planning and Scheduling using a Static Invariant-based Evaluation Model. In *Proc. of the 23rd International Conference on Automated Planning and Scheduling (ICAPS-13)*.

Pralet, C., and Verfailllie, G. 2012. Time-Dependent Simple Temporal Networks. In *Proc. of the 18th International Conference on Principles and Practice of Constraint Programming (CP-12)*, 322–338.

Pralet, C.; Verfaillie, G.; Olive, X.; Rainjonneau, S.; and Sebbag, I. 2011. Planning for an Ocean Global Surveillance Mission. In *Proc. of the 7th International Workshop on Planning and Scheduling for Space (IWPSS-11)*.

Verfaillie, G.; Infantes, G.; Lemaître, M.; Théret, N.; and Natolot, T. 2011. On-board Decision-making on Data Downloads. In *Proc. of the 7th International Workshop on Planning and Scheduling for Space (IWPSS-11)*.