Advances in Planning and Scheduling of Remote Sensing Instruments for Fleets of Earth Orbiting Satellites

Jennifer Dungan, Jeremy Frank, Ari Jónsson*, Robert Morris, David E. Smith
Computational Sciences Division
NASA Ames Research Center, MS 269-2
{frank,jonsson,morris,de2smith}@ptolemy.arc.nasa.gov
Moffett Field, CA 94035

Abstract

This paper describes a system for planning and scheduling science observations for fleets of Earth observing satellites. Input requests for imaging time on an Earth observing satellite are specified in terms of the type of data desired, the location to be observed, and an objective priority of satisfying the request. The problem is to find a sequence of start times for observations and supporting activities such as instrument slewing and enforcement of instrument thermal duty cycles, that satisfy a set of temporal and resource constraints describing the physical operation of the spacecraft. We assume that there are more requests that can possibly be serviced over a given scheduling window, and that images may vary in their scientific utility, leading to an optmization problem. This paper presents an approach to solve this problem employing 1) a formal declarative model of the problem, 2) stochastic sampling methods to find plans, and 3) special purpose heuristics based on a generalized contention measure.

Keywords: Planning, scheduling, stochastic search, constraint satisfaction, Earth observing satellites.

1 Introduction

NASA's growing fleet of Earth Observing Satellites (EOSs) employ advanced sensing technology to assist scientists in the fields of meteorology, oceanography, biology, geology, and atmospheric science to better understand the complex interactions among Earth's lands, oceans, and atmosphere. Currently, science activities on different satellites, or on different instruments on the same satellite are scheduled independently of one another, requiring the manual coordination of observations by communicating teams of mission planners.

As the number of EOSs and demand for observation time on them increase, it will no longer be viable to manually plan coordinated science activities. A more realistic vision for science observation management is to allow customers of the data (the scientists themselves) to request data products from a centralized EOS science observation management system instead of directly from an individual satellite or mission. Automated techniques will assist in the process of determining the resources that are involved in collecting data, storing the data temporarily on board satellites, and transmitting the data back to Earth. This will enable more efficient management of the fleet of satellites as well as the communication resources that support them.

Previously reported work on EOS scheduling problems includes both theoretical investigations using abstract models, as well as operational schedulers for ongoing EOS missions. Very few approaches consider multiple satellites or the coordination of observations. The theoretical studies on managing a single satellite usually involve simplified models of the satellite. For example, Lemâitre et al. [2], Pembert on [4] and Wolfe and Sorensen [7] do not discuss on-board data storage or communications system management.

There are several operational systems for ongoing EOS missions. The ASTER scheduler described in [3] and the Landsat 7 scheduler [5] are two examples. These schedulers have quite detailed models of the satellites and the communications environment. However, they do suffer from some limitations, which are discussed in [1].

Many of the search algorithms described in the literature are incomplete algorithms. The primary reason for focusing on such algorithms is that, even for small numbers of satellites, the problems are large enough that solving them optimally is impractical. The usual approach is to greedily select the next highest priority request to try and schedule, and reject it if there is nowhere for it to go. The ASTER scheduler [3] works exactly this way, as does an approach described by Wolfe and Sorensen [7].

In Frank *et. al.* [1], an automated approach for science observation management for multiple satellites based on constraint-based interval planning was proposed. The current paper summarizes results obtained using this approach, and discusses extensions and refinements that have emerged from the results of experiments undertaken using the approach. Specifically, the remainder of this paper is structured as follows:

- The satellite instrument and resource model used by the system for science observation planning is described in section 2;

*Research Institute for Advanced Computer Science (RIACS)

- A two-phased search technique for generating high quality observation schedules based on the objective of maximizing the number of high-priority requests scheduled, which combines stochastic greedy scheduling with constraint-based planning is discussed in section 3; and
- The heuristic used in the greedy scheduling phase, which selects both the observations, and the times for them based on the contention for time and other satellite resources is summarized section 4.

2 Modeling The EOS Science Observation Scheduling Problem

An EOS observation scheduling problem consists of a set of satellites, each in a particular orbit around the earth, and each with heterogeneous capabilities involving a suite of instruments and resources for downlinking data. Some satellites will have pointable instruments, providing increased flexibility in the locations they can observe at any point in an orbit. The problem also contains a set of requests, each consisting of the location to be observed, the type of data desired, and a priority, corresponding to the scientific utility of the data. Solving the EOS observation scheduling problem consists of generating a sequence of observations to be acquired by each available imaging instrument on each of a set of satellites, along with supporting activities such as instrument slews, instrument shut-down to handle thermal duty cycles, and transmissions of data back to Earth to empty the SSR.

The Constraint Based Interval Planning (CBI) framework [6], as implemented in the EUROPA planning system, was employed to solve the EOS observation scheduling problem. A general description of the modeling paradigm appears in [1]. A CBI model for the EOS domain contains a declarative description of satellite sensing instruments and resources for storing and transmitting data, as well as its orbital track. The model also describes the constraints each observation plan sequence must satisfy. These include requirements on the instruments used to collect the data, including those associated with the duty cycle for the instrument. The model also characterizes duration and ordering constraints associated with the data collecting, recording, and downlinking tasks. In addition, SSR capacity, and constraints on communications equipment such as satellite antennae and ground stations must be satisfied. There may also be set-up steps associated with particular operations, like establishing a data link prior to downlink, or aiming an instrument prior to data acquisition. These steps generate further temporal and ordering constraints.

Figure 1 visually depicts how all of these aspects are combined in a simple model. This model shows the interaction of an instrument and an SSR. The instrument transitions between `Pointing, Idle, Calibrating` and `Take-Image`. The SSR transitions between `Recording, Playback` and `Idle`. The time required for `Pointing, Calibrating, Recording` and `Playback` activities are constrained by the parameters of those activities. In addition, `Take-Image` and `Recording` activities must be simultaneous, and whenever a `Playback` occurs on the SSR the instrument must be `Idle`.

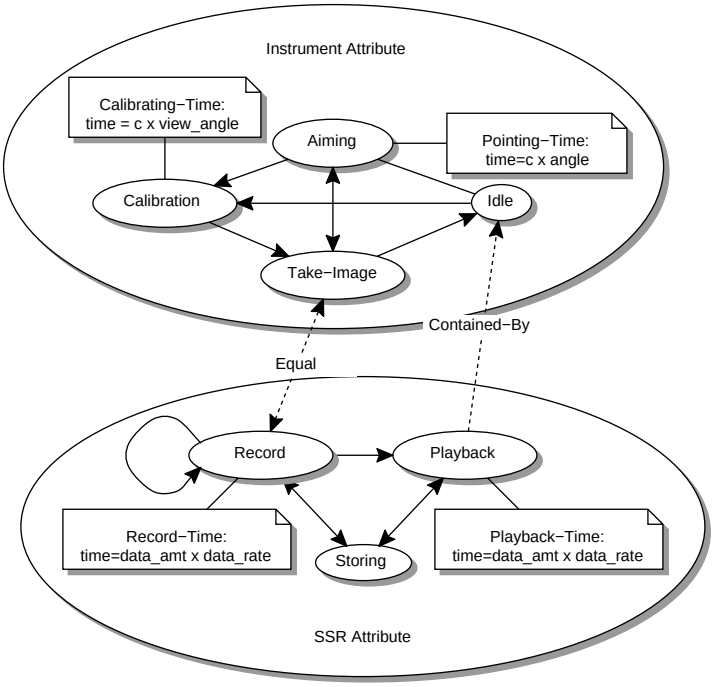

Figure 1: A simplified model showing the interaction of instrument and SSR attributes. Ovals represent the states permitted for each attribute. Solid lines indicate possible state transitions within an attribute, dashed lines indicate temporal constraints required between attributes, and boxes indicate constraints on the parameters of certain state.

In the EUROPA planning approach, the world is described in terms of a set of timelines, or *state variables*. The values for a state variable are the possible actions or states of that variable. Thus, the action values for an SSR timeline represent the actions of recording, playing back data, or idle. The model also represents set-up events such as warming up an instrument, or slewing for antennae or pointable sensing instruments. Figure 2 illustrates a small EUROPA plan involving two satellites, and a TDRSS communication satellite for downlinking data. The figure indicates that every `Take-Image` activity is synchronous with a `Record(`N`)` activity on the associated SSR, where N is a parameter standing for the amount of data added to storage. Similarly, every `Playback(`N`)` activity for a satellite is synchronous with a `Contact` activity when TDRSS is in contact with that satellite. Activities such as `Aiming` the antenna are also shown.

In the course of the work reported here, a number of different EOS scheduling models were developed, that differed in the number and types of satellite state variables, and their associated values, that were introduced. The most detailed model contained, for each satellite, state variables for multiple slewable imaging devices, SSR utilization, antenna, and ground station availability. The detailed model had the ad-

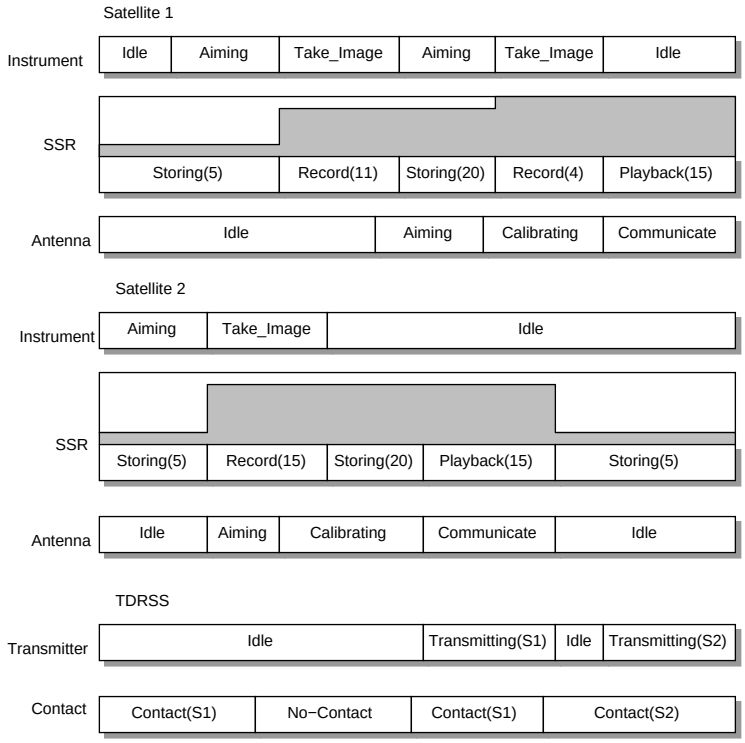

Figure 2: A complete EUROPA plan with two EOSs and one TDRSS. Each satellite has 3 attributes: the instrument, SSR, and communication antenna. The TDRSS has two attributes: the contact and transmitter.

vantage of producing the most detailed plans, but slowed down the planning process to the extent that only relatively small problem sizes (less than 50 requests) could be solved effectively. A simpler model was proposed which contained state variables for only the imaging devices. This model leverages the assumption that the periods when the satellite can communicate with Earth are known before hand, and that the satellite is guaranteed to use those periods to empty the SSR. In this model, SSR utilization and duty cycle chceks were managed via special-purpose code in the planner. This allows for the generation of solutions to problems with up to 150 requests over a 10,000 second planning horizon in time on the order of minutes. Future work will document the performance of the planner in greater depth.

3 EOS Observation Planning

The size of a typical EOS observation scheduling problem, expressed in terms of the potential number of activities that need to be scheduled in order to solve it, renders solution techniques based on complete search inapplicable. Even single satellite instances of the problem such as the single day Landsat 7 scheduling problem tend to be comprised of of hundreds of candidate imaging activities, as well as associated activities for storing and downlinking the data. On the other hand, standard greedy heuristic approaches that do not perform complete search often suffer from mypoia due to the forced adherence to the advice provided by the heuristic evaluator. This myopia often results in the inability to find high quality schedules.

As a middle ground, we have chosen a planning algorithm that combines heuristics, stochastic search and constraint propagation. A sketch of the EOS planning algorithm appears in Figure 3. During the stochastic search phase, the algorithm repeatedly selects an observation that still has time windows available, then selects a time to schedule the observation. This assignment is added to the plan, and consequences of the newly scheduled observation are checked for consistency with the existing schedule. The inferences performed during this step include the following:

- A check to ensure that the sensing instrument can be slewed in time to capture the observation just scheduled, and the observation immediately following it. The reason is that the new observation may require a long instrument slew after the observation preceeding it, or may require a long slew to the observation following it; if there is insufficient time for either slew, the observation can't be scheduled at the current time.
- A check to ensure that instrument duty cycle constraints are not violated by the added observation. The duty cycle limits the amount of time the instrument may be continuously operating; if this duration is exceeded by inserting an observation at this time, then the observation can't be scheduled at the current time.
- A check to ensure that the spacecraft SSR has capacity remaining in the interval between downlinks. Since the times of downlink are known, and the spacecraft is assumed to empty the SSR at these times, the exact storage can be computed; if inserting the observation at the current time would exceed the capacity, then the observation can't be scheduled at the current time.
- Support activities that must be assigned as the result of the added observation are inserted into the plan, and the effects of these additions are propagated.

The resulting inferences may lead to the detection of an inconsistency, meaning that the scheduling of this particular observation in this time slot is not possible and must be undone. If other times are available for the observation, the process is repeated for these candidate slots; if there are no slots left, the observation is rejected from the plan.

When it is not possible to schedule any more observations, the planning process enters its second phase. The schedule generated during the first phase is further refined to ensure that all observations having subgoals (setup steps or other preconditions) have been completely expanded. If this process is successful, the resulting schedule is returned. The process of choosing timeslots for observations and completing the resulting schedule can be repeated many times, thereby randomly sampling from the space of possible schedules.

What distinguishes the HBSS algorithm from ordinary greedy search is the way in which observations and time slots are selected. The HBSS algorithm employs a heuristic to rank the possible alternatives. HBSS then chooses probabilistically from among the alternatives, weighted according

```
HBSS
  repeat for a fixed number of times
    while observations are still possible
      Randomly select an observation
      using heuristic as stochastic bias
      until a consistent time slot is found or no choices remain
        Randomly select a time
        Assign the observation to the time slot
        Propagate constraints and decisions until
        nothing left to do or plan inconsistent
        end until
      end while
    Expand any remaining subgoals
    Check for consistency
  end repeat
end
```

Figure 3: A sketch of the HBSS algorithm modified for the EOS Scheduling problem.

to their ranking or score. Thus, possibilities ranked highly by the heuristic have higher probability of being selected, but other lower ranked possibilities are sometimes selected. This means that several alternatives with roughly the same score will have roughly equal probability of being chosen. Because of the stochastic character of the selection steps, alternative schedules are likely to be explored with each successive restart of the algorithm.

4 Contention Heuristics

The success of greedy search method for observation scheduling depends on the heuristic used for selecting the observation to schedule next, and selecting the time slot for the observation. The heuristic evaluation function chosen for the EOS scheduling problem is a weighted sum of two measures of *contention*: contention for time slots and contention for the SSR. In this section, we formally define these two measures.

Let $\mathrm{Observations}(t)$ is the set of observations that could occur at time t, and $\mathrm{Opportunities}(o)$ is the set of discrete opportunities for observation o (noting that each discrete opportunity is exactly long enough to accommodate the observation.) The *need* of an observation can be defined as:

$$\mathrm{Need}(o) = \frac{\mathrm{Priority}(o)}{|\mathrm{Opportunities}(o)|}$$

The *contention* for a particular time slot can then be defined as:

$$\mathrm{SlotContention}(t) = \sum_{o \in \mathrm{Observations}(t)} \mathrm{Need}(o)$$

The *contention* for a particular observation can then be defined as:

$$\mathrm{SlotContention}(o) = \min_{t \in \mathrm{Opportunities}(o)} \mathrm{SlotContention}(t)$$

We take the minimum here because if there is a low-contention opportunity to schedule an observation, this should not be overshadowed by other higher contention opportunity. In other words, adding another opportunity for an observation should never increase the contention measure for that observation.

Measuring contention for a global resource like SSR capacity involves generalizing the above contention measure to consider the amount of the resource needed by an observation, the resource capacity, and the interval of time under consideration. Let $\mathrm{Requires}(o, r)$ be the amount of resource r required by observation o, and let $\mathrm{Capacity}(r, i)$ be the capacity of a resource r over a time interval i. Thus, an SSR with a capacity of 50 has a $\mathrm{Capacity}(r, i) = 50$. If a playback of 20 units occurs within the interval i, then $\mathrm{Capacity}(r, i) = 70$. We then generalize the above definitions to be:

$$\mathrm{Need}(o, r) = \mathrm{Requires}(o, r) \frac{\mathrm{Priority}(o)}{|\mathrm{Opportunities}(o)|}$$

$$\mathrm{SSRContention}(r, i) = \frac{\sum_{o \in \mathrm{Observations}(i)} \mathrm{Need}(o)}{\mathrm{Capacity}(r, i)}$$

$$\mathrm{SSRContention}(r, o) = \min_{i \in \mathrm{Opportunities}(o)} \mathrm{SSRContention}(r, i)$$

Again, note that these measures change as activities are scheduled. In particular, as activities that empty the SSR are scheduled, $\mathrm{Capacity}(r, i)$ may increase, and as observations are scheduled, $\mathrm{Capacity}(r, i)$ may decrease.

Intuitively, these contention measures provide a more accurate assessment of how hard it is to actually schedule an observation. Using these measures, our variable ordering heuristic is:

> *Schedule the observation of highest priority and highest overall contention*

where overall contention is a weighted sum of contention measures for time slots and SSR capacity. We speculate that employing this composite heuristic will obtain better schedules than an approach that uses either component heuristic taken alone. On the one hand, measuring time slot contention will result in fewer observations being rejected from a schedule, but may produce assignments that violate SSR capacity restrictions; on the other hand, the SSR contention heuristic will monitor SSR load, but in doing so might ignore time slots with lower contention, thus possibly rejecting observations that could have been scheduled.

Given an observation to schedule, we would prefer to put it in the place where it will compete with the fewest other observations. We can use the above contention measures to define a value ordering heuristic:

> *Schedule an observation in the opportunity with the least contention*

5 Conclusions and Future Work

The best tradeoff between the timeslot contention and the SSR contention is unknown. We have been experimenting with differ weight assignments to each component measure

in order to evaluate the relative importance of each in generating high quality schedules. These results will be published in future work.

This system described here has been proposed as part of a distributed architecture for science observation management that potentially includes an on-board component. The on-board system could be effective in providing a more accurate assessment of the utility of a scheduled observation, based on inputs from assets like on-board sensors, communication from other satellites, or real-time weather information. It could also allow for more reactivity to degraded capability of resources, whether consisting of loss of ground station availability, an observing instrument, or SSR deterioration. These inputs would be used for inserting new observations into the schedule, or discarding low priority stored data from previous observations. Future work will report results gathered from this aspect of the research.

References

[1] J. Frank, A. Jónsson, R. Morris, and D. Smith. Planning and scheduling for fleets of earth observing satellites. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2001.

[2] M Lemâitre, G. Verfaillie, F. Jouhaud, J. Lachiver, and N. Bataille. How to manage the new generation of agile earth observation satellites? In *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2000.

[3] H. Muraoka, R. Cohen, T. Ohno, and N. Doi. Aster observation scheduling algorithm. In *Proceedings of the International Symposium Space Mission Operations and Ground Data Systems*, 1998.

[4] J. Pemberton. Towards scheduling over-constrained remote sensing satellites. In *Proceedings of the 2d International Workshop on Planning and Scheduling for Space*, 2000.

[5] W. Potter and J. Gasch. A photo album of earth: Scheduling landsat 7 mission daily activities. In *Proceedings of the International Symposium Space Mission Operations and Ground Data Systems*, 1998.

[6] D. Smith, J. Frank, and A. Jónsson. Bridging the gap between planning and scheduling. *Knowledge Engineering Review*, 15(1), 2000.

[7] W. Wolfe and S. Sorensen. Three scheduling algorithms applied to the earth observing domain. *Management Science*, 46(1), 2000.