Appears in Proc. Intl. Workshop on Planning & Scheduling for Space, Moffett Field, CA, March 2013.

Mission Design Evaluation Using Automated Planning for High Resolution Imaging of Dynamic Surface Processes from the ISS

Russell Knight, Andrea Donnellan, and Joseph J. Green
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
Russell.L.Knight@jpl.nasa.gov
Andrea.Donnellan@jpl.nasa.gov
Joseph.J.Green@jpl.nasa.gov

Abstract

A challenge for any proposed mission is to demonstrate convincingly that the proposed systems will in fact deliver the science promised. Funding agencies and mission design personnel are becoming ever more skeptical of the abstractions that form the basis of the current state of the practice with respect to approximating science return. To address this, we have been using automated planning and scheduling technology to provide actual coverage campaigns that provide better predictive performance with respect to science return for a given mission design and set of mission objectives given implementation uncertainties. Specifically, we have applied an adaptation of ASPEN and SPICE to the Eagle-Eye domain that demonstrates the performance of the mission design with respect to coverage of science imaging targets that address climate change and disaster response. Eagle-Eye is an Earth-imaging telescope that has been proposed to fly aboard the International Space Station (ISS).

Introduction

This abstract is organized as such:

1. brief Eagle-Eye domain description
2. planning problem description
3. planning and scheduling architecture
    a. role of ASPEN
    b. role of SPICE
4. key transformation of the point coverage problem into a constraint-based timeline
5. results
6. related work

Portions of this work were carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Eagle-Eye Domain Description

The primary goal of the proposed Eagle Eye mission is to understand how Earth's vulnerable systems reflect changes in climate [Donnellan et al 2013]. This is accomplished by measuring the change over time of various features on the earth, e.g., glaciers and sand dunes. Another role of the proposed Eagle Eye mission is to increase our understanding of natural hazards. This includes monitoring various features over time, e.g., volcanoes [Chien et al 2011[1]], fire [Chien et al 2011[2]] landslides, dynamics near faults due to earthquakes, flooding [Chien et al 2012], and coastal change, in concert with the earth-observing sensor web [Chien et al 2005].

To satisfy these goals, we intend to image the Earth using a telescope aboard the ISS. The telescope consists of a 0.5 m primary mirror with integral active control that provides on-demand figure control and telescope alignment. This maintains high quality imaging through the lifetime of the instrument. The area visible from the ISS is from 1 to 2 km^2 at a resolution similar to current commercial satellites. The field of regard (area to which we can slew the telescope) is approximately 45 degrees. An integral pointing/isolation system compensates for ISS disturbances, and external mounting on a nadir truss provides excellent access to Earth-science targets.

Figure 1 Active Mirror and ISS Telescope Concept

The various goals for science are described as prioritized points on the Earth. For this domain, there were approximately 100,000 targets.

Planning Problem Description

The goal of the planning system is to model the constraints of the telescope including the ISS ephemerides and produce a plan that services as many points as possible while respecting the priorities given by the scientists. Specifically, produce a series of slews and settles that cover targets, and ensure that no higher priority target can be scheduled by removing a lower priority target, all while respecting data pipeline constraints, instrument duty cycle and maintenance periods, and respecting pointing, slew, and settle time constraints. What follows is a detailed list of the modeled constraints.

Communications and Memory

The number of ground sites is an adjustable parameter. For each ground site, we are given a name, a communication rate, a minimum elevation angle, a position (which includes latitude, longitude, and altitude), and a style (a string indicating the KML style to use when rendering "communication windows", e.g., pale_green). One example set that we use (which is highly unlikely to be used in practice) is the set of three deep space network antenna complexes. We use a minimum elevation of 30 degrees (which is pessimistic) and a communication rate of 1.2 Gigabits per second.

Total onboard memory is modeled, as well as the compression rate from the instrument and total bandwidth for the instrument bus and communications bus. Onboard we have 1 Terabit of memory and assume a 3:1 compression factor (which is highly compressed, but attainable).

Instrument

The instrument FOV is set to 2 kilometers at closest nadir. The angular FOV value is derived from this and used throughout. The total field of view (area that we can slew to) is 45 degrees.

The number of frames per second is at most 5 (for this model, but this might be optimistic and a value of 4 might be closer to ground truth). An image is 1 Megabit.

Spacecraft

SPICE is used to produce a timeline (a series of states over time) of the location of the spacecraft.

The slew and settle times are modeled using a simple controller that had a maximum slew rate and a maximum slew acceleration/deceleration. The settle time is a constant of 20 seconds. The longest possible slew requires less than a second, and we book-keep at least a second per slew, so the transition time between frames is a constant factor of 21 seconds. The reason for this is that we are gimbaled to the ISS. This allows us to slew very quickly as we have a large amount of available power and the telescope mass is a tiny fraction of the ISS mass.

Power and thermal considerations are minimal (and thusly ignored) as we are using the ISS as our platform, thus we have enough power and the ability to thermally sink our assembly.

Targets

For science targets, we adapted SPICE to produce a timeline of valid times for each point to be observed, resulting in over 100,000 visibility timelines. Each target consists of a location, a name, a minimum elevation angle to the instrument, a minimum solar elevation angle, a duration, a frame rate, and a priority. Note that we are enforcing a total ordering on priority, but in practice we realize it will likely be the case that there are ties in priority and possibly incomparable classes of targets. Also, it is required that we image these targets at the maximum quality, which in practice is a black-box function of solar angle, instrument angle, and instrument distance. This results in grounded times for each target. This black-box function was easily integrated into the SPICE library as it contains a single peak in the space of any legal interval with respect to sun angle and instrument angle. This peak is quickly found using simple binary search techniques in $\log_2 n$ time, where n is the number of bits of accuracy required.

Maintenance Activities

To ensure the quality of the figure of the mirror, we occasionally need to compute the error in the figure of the mirror and adjust the actuators accordingly. The amount of time between measuring the error and correcting the error needs to be minimized, but we need to compute the changes to the mirror on the ground. The amount of data to be transferred is minimal. It requires approximately 1 hour to compute the change and 20 minutes to either measure the error in the mirror or adjust the mirror.

Planning and Scheduling Architecture

We clearly wanted the flexibility of modeling the data flow and timing constraints, and we wanted to guarantee a certain level of performance on the part of the planner with respect to quality. We also had to include a high level of accuracy with respect to modeling of pointing constraints and ephemerides. To achieve this, we combined two systems: the ASPEN planning system (Fukunaga et al, 1997), and the SPICE library (Acton, 1992).

ASPEN Adaptation

The ASPEN planning and scheduling system was used to model the slew, settle, imaging, data transfer, and maintenance activities. ASPEN provides a rich modeling language that was sufficient for most of this modeling, with the exception of computing durations for slews and valid times for imaging activities. These needed to be provided by an ephemeris propagation and geometry toolkit: SPICE.

Spacecraft location consists of an ASPEN parameter indicating the kernel file name. This is passed to SPICE, which in turn generates an ASPEN double-valued state variable timeline for latitude and longitude.

Communications ground sites consisted of parameters in ASPEN for minimum elevation, name, latitude, longitude, and altitude. These were passed to SPICE, which in turn generated a timeline of visibilities that was reflected in ASPEN as a string-valued state variable with values of "in view" or "not in view", accordingly.

Observations are modeled as ASPEN activities with effects on the frame rate timeline. This is modeled as a non-depletable resource with effects on the raw data rate timeline. The effect from the frame rate timeline on the raw data rate timeline is simply the product of the rate (or value) and the memory per image, which is 1 Megabit.

The raw data rate is modeled as an ASPEN non-depletable resource timeline with effects on the compressed data rate timeline. The effect is simply rate (or value) times the compression rate (or 1/3).

The compressed data rate is modeled as an ASPEN non-depletable resource timeline with effects on the onboard memory timeline. The effect is simply the rate.

The onboard memory timeline is double-valued ASPEN integral timeline. This timeline can be constrained by levying rates of usage, and it accumulates or drains according to the current rate. Constraints can be placed on the minimum and maximum fill levels, as well as clamps on the value that cause the cessation of filling or draining, e.g., a minimum clamp of 0 is common as once memory is empty; it isn't a constraint violation to try to empty it some more (although it might be a waste of time). The maximum capacity is 1 Terabit.

Our downlink operations are to downlink whenever a site is in view. So, each site has an associated state timeline with effects on the ground sites in view timeline that, if the value of the in view timeline is "in view", add 1 to the ground sites in view timeline.

The ground sites in view timeline is a non-depletable resource with effects on the onboard memory timeline. The effect is the maximum rate of the visible ground stations (or actually, the product of this value and -1 as we are releasing memory).

SPICE Adaptation

As previously stated, SPICE was used to produce a timeline (a series of states over time) of the location of the spacecraft. For science targets, we adapted SPICE to produce a timeline of valid times for each point to be observed. To produce the location timeline of the spacecraft, we needed to generate a SPICE kernel.

We generate the kernel (a file describing the path of the spacecraft over time) using the mkspk command and a configuration file that included the two line element (TLE) for the ISS. Note that we had to modify mkspk to allow for longer duration propagation using a TLE as usually kernels generated this way were limited to 2 days of duration, but we needed up to a year. The reason for the limited duration of TLE-based kernels is that the orbital propagation algorithm for TLEs is quite limited with respect to accuracy, and in practice JPL uses higher fidelity modeling techniques for ephemerides in actual flight projects.

This still left us with the challenge of how to actually schedule the observations.

Point Coverage Constraint Formulation

First, many targets can be observed at the same time as other targets by positioning the telescope such that both are in the field of view simultaneously. If we imagine each target point in the x/y plane, and we imagine that the projected telescope field of view is a square, then we could try to cover as many points as possible with a single square, thereby reducing our overall imaging requirements. Unfortunately, the projection of the instrument on the surface of the Earth is not square, so we instead projected the points to the instrument.

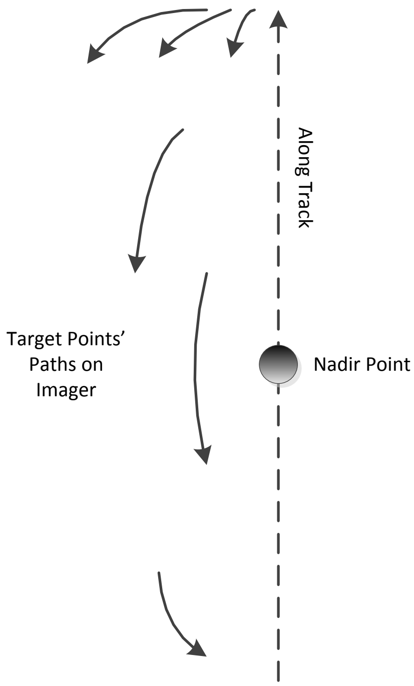

Figure 2 Paths of Points on the imager

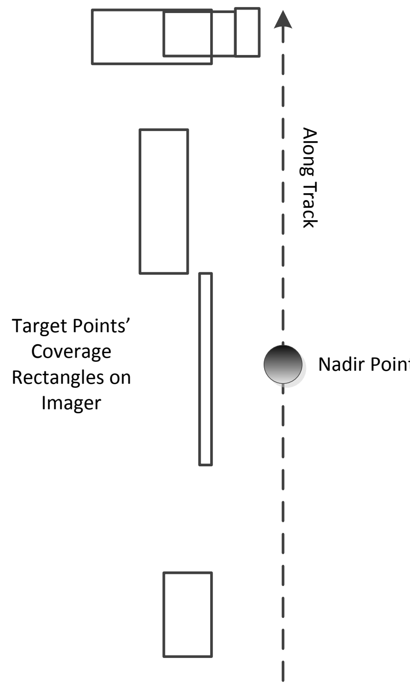

Figure 3 Rectangles on the imager

Now, for any given duration, a single point would be a curved line in the instrument space. If we expand the instrument space to include the whole slew space of the telescope, then we have the projected path of all possibly viewable points. Since each observation is required to occur at a certain point in time, then we have a fixed curve for each observation. These curves can be mapped to rectangular boxes where the goal is to collect each of the boxes.

Now, if we consider sorting all of the points in priority order, and then include them in a solution, but do not commit to the actual positioning of the telescope, we get a range of places that any observation can possibly be, based on previous observations. The problem is, if we have ungrounded previous observations, how do we control the combinatorial explosion as we propagate forward in time.

We manage this by first making the observation that there are only quadratically many meaningful observation orientations for a given set of points, thus there are only quadratically many possible positions of the telescope at any time.

Proof outline of tractability

Given a set of rectangles $r = r.x$, $r.y$, $r.width$, $r.height \in R$ in 2-d space, an aperture $a = a{,}x$ (variable), $a.y$ (variable), $a.width$, $a.height$ in the same 2-d space, then there exist at most $O(n^2)$ simultaneous assignments to $a.x$ and $a.y$ that contain unique sets of rectangles $r \in R$.

Lemma 1: any pair of rectangles that can be covered by an aperture assignment can be covered by a lower-left assignment. Let r_1 and r_2 be the rectangles to be covered. Let $a.x = \mathbf{min}(r_1.x,\ r_2.x)$ and $a.y = \mathbf{min}(r_1.y,\ r_2.y)$ be the lower-left assignment of a over r_1 and r_2. Every assignment that covers r1 and r2 satisfies the following equations:

1. $a.x \leq r_1.x$ and $a.x \leq r_2.x$
2. $a.y \leq r_1.y$ and $a.y \leq r_2.y$
3. $a.x + a.width \leq r_1.x + r_1.width$ and $a.x + a.width \leq r_2.x + r_2.width$

4. $a.y + a.height \leq r_1.y + r_1.width$ and $a.y + a.height \leq r_2.y + r_2.height$

$a.x$ is maximal in that increasing $a.x$ results in a violation of at least one of the equations in 1. The minimum legal assignment of $a.x$ is $minimal(a.x) = \mathbf{max}(r_1.x + r_1.width, r_2.x + r_2.width) - a.width$. If we reduce $a.x$ below this value then at least one of the equations in 3 would be violated. We need only show that the maximal assignment ($a.x$) covers the minimal assignment sufficiently to cover r_1 and r_2, which is clear as all corners in all rectangles have x values $\geq a.x$ and all corners in all rectangles have x values $\leq a.x + a.width$. The same argument holds for y values and *height*.

Lemma 2: for any set of coverable rectangles $C \subseteq R$, there exists a pair $r_1 \in C$, $r_2 \in C$ such that the maximal assignment to $a.x$ is the lower-left assignment over r_1 and r_2.

For our pair, we need only choose r_1 with the minimum $r_1.x$ value and choose r_2 with the minimum $r_2.y$ value. We then apply a pair-wise evaluation to prove that $a.x$ and $a.y$ serve as a covering assignment for all pairs that cover either r_1 or r_2 and any other $r \in C$.

Given Lemma 1 and Lemma 2, we see that the maximum number of useful simultaneous assignments to $a.x$ and $a.y$, given any $C \subseteq R$ is at most quadratic in $|C|$ as any coverable $C \subseteq R$ can be covered by covering a pair $r_1 \in C$ and $r_2 \in C$. (Note that $r_1 = r_2$ is allowed.)□

Of course, theoretical tractability (polynomial-time computation) does not explicitly imply practical tractability. To support the claim of practical tractability, we provide empirical results in our Results section.

Scheduling Algorithm

To schedule the observations, we first produce a timeline that consists of a series of time-stamped sets of assignments to a (the aperture) based on the observation times and durations, without respect to priority. Then, for each observation in priority order, we winnow down sets by propagating the constraints based on settling time and containment. Thus at any moment during scheduling, we have a disjoint set of aperture assignments to choose from. This is in essence a disjunctive state propagation technique, where we have at most n^2 disjoint states to propagate from that can continue to at most n^2 subsequent states.

To increase efficiency, we only propagate forward, and when the time comes to select a solution, we start from the end and work backwards. That is because forward propagation assumes some states in the future might be usable when in fact they are not.

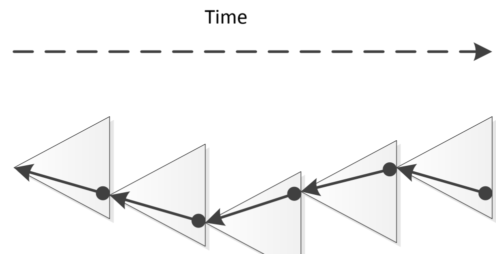

Figure 4 Extracting a Grounded Schedule from the disjoint schedule timeline

We need only compute the possible transitions from one state to the next to determine which states are infeasible. This propagating timeline in the expanded imaging field space actually enables us to provide provably optimal solutions in a reasonable amount of time.

Results

For the problems presented, we can cover all of the requested 100,000 targets over the course of a month. This requires approximately 12 minutes of compute time, including initialization. To aid the scientists in understanding the schedules, we output our results in KML format, thereby enabling display in Google Earth. Here we see the covered points during several passes, along with the ISS orbit.

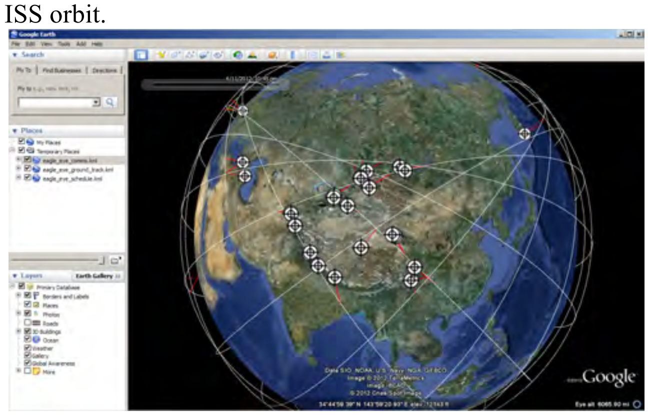

Figure 5 Schedule Fragment in KML

With these results in hand, we can provide convincing evidence that our planned system has the capability to handle the science load promised. We have the ability to perform rapid trade studies due to the short period of time required to run the simulation, and we have eased data entry by providing an Excel spreadsheet to our modelers.

Related Work

Much of the work in automated scheduling of coverage campaigns formulates the scheduling problem as one where the scheduler has more flexibility in time than the problem posed by this work, making the solving of the problems inherently NP-hard. [Knight and Chien 2006, Knight et al 2007, Rabideau et al 2010, Knight et al 2012, and Verfaillie et al 2012] In contrast, we have shown that the constraints imposed by mission designers have in fact made scheduling theoretically and practically tractable.

References

Acton, C. H. "The SPICE System: Observation Geometry Supporting Data Visualization." JPL TRS 1992, http://hdl.handle.net/2014/35980

Chien, S., Cichy, B., Davies, A., Tran, D., Rabideau, G., Castano, R., Sherwood, R., Mandl, D., Frye, S., Schulman, S., Jones, J. and Grosvenor, S., "An autonomous earth-observing sensorweb," IEEE Intel ligent Syst., vol. 20, no. 3, pp. 16–24, 2005.

Chien, S., Davies, A., Doubleday, J, Tran, D. Jones, S., Kjartansson, E., Vogfjörd, K., Gudmundsson, M.T., Thordarson, T., Mandl, D., "Integrating multiple space and ground sensors to track volcanic activity," ISRSE. Proceedings of the 34th International Symposium on Remote Sensing of Environment, 2011.

Chien, S., Doubleday, J., McLaren, D., Davies, A., Tran, D., Tanpipat, V., Akaakara, S., Ratanasuwan, A., and Mandl, D. "Space-based Sensorweb monitoring of wildfires in Thailand." In Geoscience and Remote Sensing Symposium (IGARSS), 2011 IEEE International, pp. 1906-1909. IEEE, 2011.

Chien, S., Mclaren, D., Doubleday, J., Tran, D., Tanpipat, R. Chitradon, R., Boonya-aroonnet, S., Thanapakpawin, P., and Mandl, D. "Automated space-based monitoring of flooding in Thailand," in Proc. Int. Symp. Artificial Intelligence, Robotics, and Automation for Space, Turin, Italy, Sep. 2012.

Chien, S., McLaren, D., Rabideau, G., Silverman, D., Mandl, D., Hengemihle, J. "Onboard Processing for Low-latency Science for the HyspIRI Mission", International Symposium on Space Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS 2010). Sapporo, Japan. August 2010

Donnellan, A., Green, J. J., DeJong, E. M., Knight, R. L., Bills, B., Arrowsmith, R.. "High Resolution Imaging of Dynamic Surface Processes from the ISS." IEEE Aerospace Conference, Big Sky Montana, Mar 2-Mar 9, 2013.

Fukunaga, A., Rabideau, G., Chien, S., Yan, D., "ASPEN: A Framework for Automated Planning and Operations." Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Oxnard, California, 1997.

Knight, R., McLaren, D., Hu, S., "Planning Coverage Campaigns for Mission Design and Analysis: CLASP for the Proposed DESDynI Mission," International Symposium on Artificial Intelligence, Robotics and Automation in Space, Turin, Italy, September 4-6, 2012.

Knight, R., and Chien, S., "Producing Large Observation Campaigns Using Compressed Problem Representations." International Workshop on Planning and Scheduling for Space, Space Telescope Science Institute, Maryland, October 23-26, 2006.

Rabideau, G., Chien, S., Mclaren, D., Knight, R. Anwar, S., Mehall, G. "A Tool for Scheduling THEMIS Observations," International Symposium on Artificial Intelligence, Robotics and Automation in Space, Sapporo, Japan, August 2010.

Verfaillie, G., Olive, X., Pralet, C., Rainjonneau, S., and Sebbag, I., "Planning Acquisitions for an Ocean Global Surveillance Mission," International Symposium on Artificial Intelligence, Robotics and Automation in Space, Turin, Italy, September 4-6, 2012.