Using Automated Scheduling to Assess Coverage for Europa Clipper and JUpiter ICy moons Explorer

Martina Troesch and **Steve Chien** and **Eric Ferguson**
Jet Propulsion Laboratory
California Institute of Technology
USA
firstname.lastname@jpl.nasa.gov

Abstract

We describe the use of an automated scheduling system to assess mapping coverage for space missions. This tool uses a gridded representation of target body surface regions of interest to calculate surface coverage based on science objectives, such as distance to target and lighting conditions, and spacecraft constraints, such as data volume. The science objectives and constraints are modelled in a greedy optimization, scheduling algorithm that generates observation schedules. We demonstrate the application of this tool to evaluating achievement of mission science criteria for the planned Europa Clipper mission and the JUpiter ICy moons Explorer (JUICE).

Introduction

The development of space missions involves the coordination of many complex parts, including the interaction between trajectory design and the opportunity to observe desired science targets. The use of automated scheduling tools aids in analyzing the achievement of science goals against different spacecraft configurations and trajectories by using models of the instruments, spacecraft, and targets. One useful application for automated scheduling is for flyby missions, where the limited opportunities for scientific observations with specific geometric constraints and restrictions on trajectory design result in a need for multiple analyses and optimizations.

Two such future missions include NASA's planned Europa Clipper mission [National Aeronautics and Space Administration] and ESA's JUpiter ICy moons Explorer (JUICE) [European Space Agency]. The Europa Clipper mission aims to determine if the Jovian moon Europa has conditions amenable for life. To achieve this objective, the spacecraft would use nine instruments and perform 45 flybys of Europa at altitudes between 25 km and 2,700 km.

The JUICE mission would also visit the Jovian system, but with the intent to understand its formation and development, as well as to investigate there exists the possibility of sustaining life. To complete this investigation, JUICE would use its ten instruments and perform tours and flybys of Jupiter and its moons Ganymede, Callisto, and Europa, with an emphasis on Ganymede.

We use these two missions as case studies to describe the use of an adaptation of the Compressed Large-scale Activity Scheduler and Planner (CLASP) [Knight and Chien 2006], an automated scheduling tool which uses Squeaky Wheel Optimization [Joslin and Clements 1999]. To use CLASP, we must first define the spacecraft and instrument geometries and constraints, such as when instruments can take data, data volume limits, and fields-of-view. Next, the science goals need to be specified, including constraints such as geometry and instrument modes, as well as the priority for the goal. From these definitions, CLASP can generate an observation schedule that maximizes the science priorities while maintaining all spacecraft and scientific constraints.

The remainder of this paper is organized as follows. First we describe the problem formulation, followed by an explanation of the scheduling algorithm. Next we present the results of the case studies, and provide discussion and related works. Finally, we present our conclusions.

This work is a continuation of work in [Rabideau, Chien, and Ferguson 2015], thus the structure of this paper follows the structure of that work, using similar text and information throughout as well as replicating text in the following sections: "Observation Selection" and "Discussion and Related Work". This paper uses updated instrument models and a new trajectory for Europa Clipper, and also evaluates coverage for the JUICE mission using regions of interest instead of solely global coverage.

Spacecraft and Instrument Modelling

For both the Europa Clipper mission and the JUICE mission, the position and orientation of the target bodies, the position of the spacecraft, and the instruments are modelled. The target body for Europa Clipper is Europa, whereas there are multiple target bodies for the JUICE mission, which for this analysis include Europa, Io, Callisto, and Ganymede.

Although the Europa Clipper spacecraft would carry a variety of scientific instruments, our analysis is focused on the remote sensing instruments. These include the Europa Imaging System (EIS) with wide-angle (WAC) and narrow-angle (NAC) cameras each with a pushbroom and framing mode, the Europa THermal Emission Imaging System (E-THEMIS) with a single mode, the Mapping Imaging Spectrometer for Europa (MISE) with a single mode, the Radar for Europa Assessment and Sounding: Ocean to Nearsurface (REASON) with a low and high PRF mode, and the Ultraviolet Spectrograph/Europa (UVS) with one mode, ESAA stare, out of many analyzed.

The JUICE instruments that are included for analysis are the GAnymede Laser Altimeter (GALA), the camera system (JANUS), the Moons and Jupiter Imaging Spectrometer (MAJIS), the Radar for Icy Moons Exploration (RIME), the Sub-millimeter Wave Instrument (SWI), and the UV imaging Spectrograph (UVS). Each of these instruments is modelled with a single mode.

Each instrument is defined by its field-of-view, one or more modes, and an optional duty cycle. Since this analysis is focused on the maximum coverage, no data rate restrictions were imposed.

The data for the Clipper instrument models were extracted from instrument kernels produced for APGEN [Maldague et al. 2014]. The instrument data for JUICE was provided by ESA. This data was then written in a format that could be ingested by CLASP.

Science Campaigns

Scientists achieve research goals by acquiring data from specific instruments over particular regions of interest under desired geometric constraints. Science data acquisition in CLASP is defined through *campaigns*. Each campaign specifies one or more instruments with their operating modes (e.g. mono, stereo), the region of investigation on the target body (e.g. the north pole), required geometric conditions (e.g. emission angle, distance), as well as a priority chosen by the scientists. These campaigns are specified in Keyhole Markup Language (KML), which can be ingested by CLASP.

This study focuses on global coverage, therefore the surface region for investigation used by all campaigns is defined by multiple polygons covering the entire target body. Specifically, each target body is divided into 14 polygons with 4 at each of the poles ($\pm$ 60 degrees Latitude) and 6 around the equator ($\pm$ 30 degrees Latitude). The JUICE investigation also includes latitude, longitude point regions of interest.

Observation Selection

In order to assess areal coverage, CLASP uses a gridded representation of regions. In this representation, the planetary surface is represented by a set of roughly equidistant grid points with separation D. Specifically, grid points exist along lines of latitude that are spaced distance D apart. Along these lines, there are grid points spaced D apart, surrounding the globe.

This gridded representation allows CLASP to compute overlap between regions very efficiently. With this representation, rather than computing polygon overlap on a surface directly, the computation simply intersects grid point sets. Gridded overlap computation is bit set intersection and is $O(n)$ theoretically, where n is the number of points in the grid, but in practice these bit vector operations are effectively constant time. Polygon overlap computation is $O(n \log n)$ theoretically and in practice $O(n)$ where n is the number of points defining the polygons.

For this application, we use 800 grid points around the equator of Europa for the Europa Clipper mission. This results in 12.3 km between grid points and 200,000 potential targets of observation. For the JUICE mission, we use 1600 grid points around the equator of each target body.

First, CLASP computes the visibility of the Europa surface as seen from each instrument. The proposed trajectories for both the Europa Clipper mission and the JUICE mission have the spacecraft making a series of flybys around the target body. Each flyby results in a set of visibility "swaths" across the surface, where each swath represents observation opportunities for an instrument. The size, shape, and location of the swath depend on the position and orientation of the spacecraft, and the field-of-view of the instrument. To generate swaths, CLASP makes use of the CSPICE Toolkit provided by the Navigation and Ancillary Facility (NAIF) at JPL. CLASP was originally designed for nadir pointing coverage calculation, although adaptations have been developed that handle off nadir pointing by modelling this as different instrument modes. However, this problem formulation is not computationally efficient, since CLASP computes potential coverage for every alternate pointing for each timestep as a different instrument mode and then searches in the subset selection of possible instrument modes (e.g. instrument pointings).

Next, CLASP computes the intersection points between instrument swaths and campaigns. This is done by iterating through instrument swath points and creating a "potential observation" record for each such point, and for each campaign requiring a unique instrument mode. For example, if one campaign requires E-THEMIS and another campaign requires MISE, and a point is visible by both instruments, then two potential observation records are created. If both campaigns use the same instrument, only one observation record is created. Each observation record is then accorded the highest priority from each of its campaigns.

The observation selection problem is the following:

Given
  a set of potential observation records $O = \{o_1 ... o_n\}$
  a set of regions of interest $R = \{r_1 ... r_n\}$
  a set of instrument swaths $I = \{i_1 ... i_n\}$
    where $\forall o_i \in O \exists (r_i, i_i)$ such that
    $(grid(o_i) \in grid(r_i)) \wedge (grid(o_i) \in grid(i_i))$
  a scoring function $U(r_i) \rightarrow \mathbb{R}$
  a constraint function $C(S) \rightarrow T, F$
    where $S \subseteq O$ and $C(S)$ is T
    if S satisfies spacecraft constraints
Select a set of observation records $A \subseteq O$
  To maximize $U(r_i) \forall r_i \in R$
  subject to $C(A) \rightarrow T$

CLASP for the Europa Clipper and JUICE missions currently validates several operations constraints when selecting observations:

- Instrument field-of-view: each instrument is defined by a different field-of-view, which results in different visibilities for each instrument during the flybys.

- Distance from target: since the resolution, and thus quality, of data is dependent on the distance from the surface, and since flybys result in a large amount of time spent very far the target body, the campaigns are all defined with maximum distance constraints.
- Lighting conditions: lighting conditions affect the quality of the data, therefore the campaigns impose constraints on the solar zenith angle, emission angle, phase angle, and occassionaly a range of target body local times.

CLASP uses squeaky wheel optimization (SWO), an iterative heuristic approach to optimization. SWO uses a simple, priority-based, greedy scheduler as an inner loop with an outer loop that iteratively tweaks inputs to the inner loop. Each iteration is a call to SWO_INNER below and consists of iterating through the potential observation records in order of decreasing priority. The observation record is added to the schedule if it can be performed without violating any spacecraft operations constraints. Otherwise, the observation record is discarded and the next observation record is considered.

Whenever an observation record is added to the schedule, CLASP must compute which additional observation records are also implied to be in the schedule (the Propagate function below). This propagation occurs based on two checks. The instrument swath polygon associated with the selected observation record may include multiple grid points. For any of these grid points, other observation records with the same instrument mode will also be added to the schedule. The result of SWO_INNER is a set of observation records A such that $C(A)$ is satisfied.

The outer loop of SWO consists of first initializing the observation record priorities to the priorities of the parent science campaigns. Then, SWO_OUTER repeatedly calls SWO_INNER to produce a set of selected observation records A. Note that the propagation in SWO_INNER only adds observation records that are logically entailed by selected observations. For example, imaging an area (or grid point) with instrument I at 4m spatial resolution subsumes imaging with I and 8m spatial resolution. Or imaging an area with instrument I with spectral bands 2, 4, and 6 subsumes imaging with instrument I with spectral bands 2 and 4. Therefore there is no search (and no backtracking) in the propagation step. As long as a progress metric is satisfied, we increment the priority of all observation records that did not make it into the current schedule A, and re-run. This rerunning proceeds a number of iterations, which is specified by the user, and the best schedule (scored by initial priorities) is returned.

```
procedure SWO_OUTER
    Initialize priorities
    while progress made do
        SWO_INNER() → A
        for each o in (O − A) do
            increase the priority of o
        end for
    end while
end procedure

procedure SWO_INNER
    O = all candidate observation records
    B = {}
    for each o in O in decreasing priority order do
        if C(B + o + PROPAGATE(o)) = True then
            B := B + o + PROPAGATE(o)
        end if
    end for
end procedure
```

Europa Clipper Coverage Analysis

The Europa Clipper campaigns were run in CLASP to produce a schedule of observations from a single squeaky wheel outer loop. The coverage analysis was performed for global coverage using a single trajectory of the first five flybys of Europa, with only nadir pointing, and allowing all instruments to be on simultaneously without data limits. The trajectory used was 15F10_DIR_L220614_A250305_V1_scpse and the flybys were approximately every two weeks starting in February 2026. Each flyby was run as a separate simulation by changing the start time and simulation duration. A single flyby takes approximately forty minutes of wall clock time to simulate. If desired, it is possible to perform post-processing to dermine how much additional coverage is achieved by each flyby or the consequences of choosing only a subset of flybys. Figure 1 reports the coverage results for each of the instruments by flyby, where each flyby coverage is computed independently. As an example, Fig. 2 shows the swath covered by REASON on the second flyby.

JUICE Coverage Analysis

CLASP was run for the JUICE campaigns and generated an observation schedule from one squeaky wheel outer loop. Only nadir pointing was considered and all instruments were allowed to be taking data simultaneously with no data limits. First, the coverage analysis for the JUICE mission was performed for the cumulative global coverage of the ten closest flybys for each target body: Europa, Io, Callisto, and Ganymede. Like the Europa Clipper coverage analysis, each flyby was simulated separately. These coverage results are shown in Fig. 3.

Next, at the request of ESA scientists, the coverage for regions of interest on Ganymede were determined. Each region of interest is defined as a single latitude, longitude point on the surface of Ganymede. The coverage is calculated like the global coverage problem, but instead of having an even grid covering the entire surface of the target body, the only grid points are those that make up the regions of interest. A total of 92 regions of interest were considered. Table 1 summarizes the number of regions of interest that were observed by any instrument for each of the Ganymede flybys. If the same region of interest is seen for multiple flybys, it is counted for each flyby.

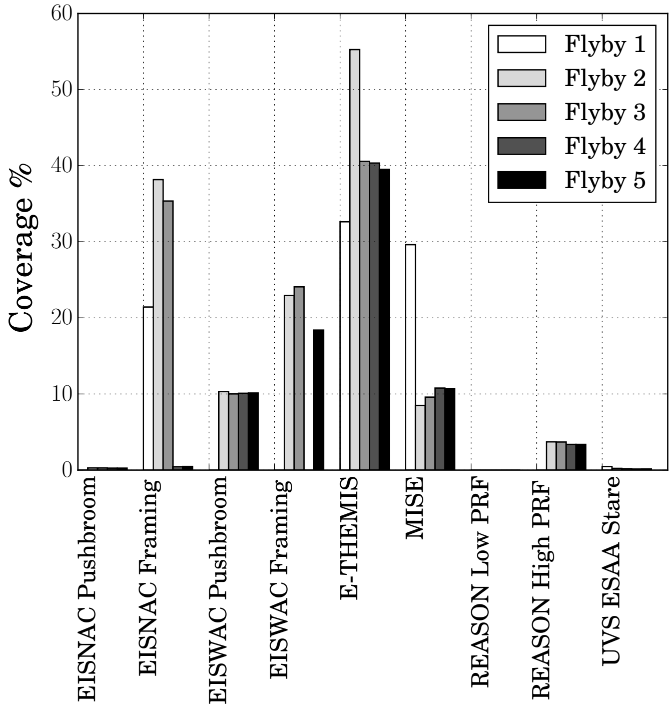

Figure 1: The percentage coverage of Europa for each instrument and flyby of the Europa Clipper spacecraft. The coverage for each flyby is computed independently.

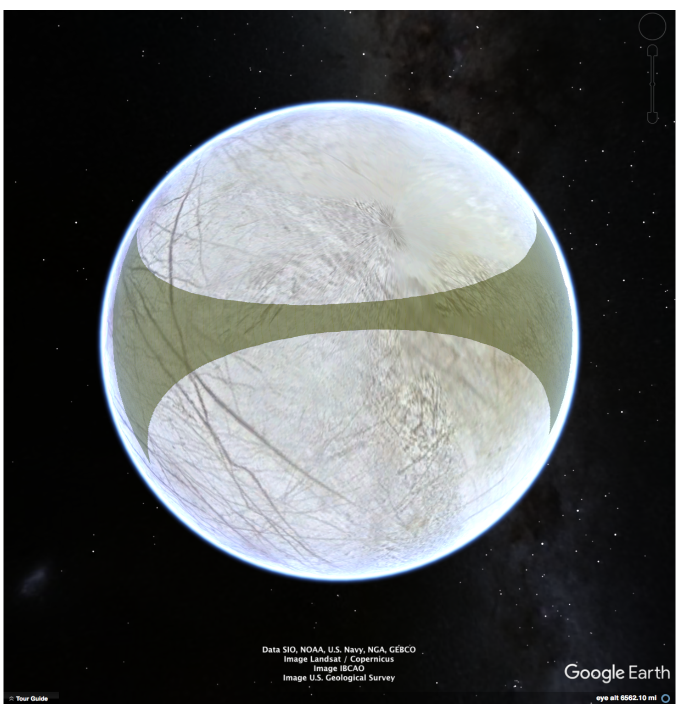

Figure 2: Coverage of Europa from the Europa Clipper spacecraft using REASON on the second flyby.

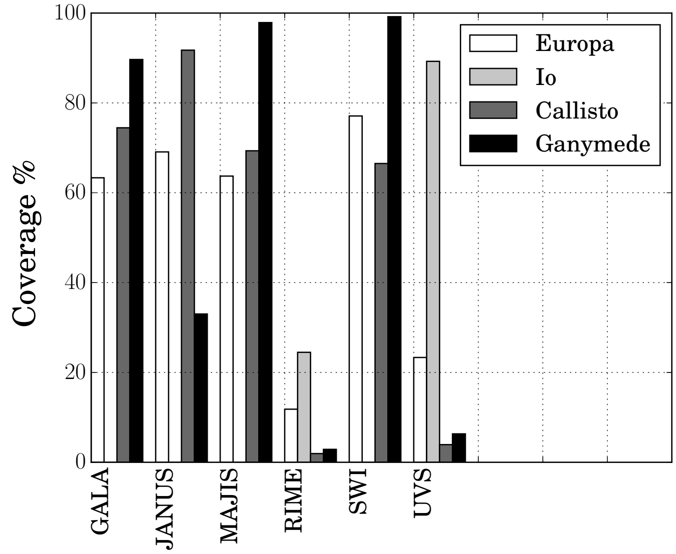

Figure 3: The cumulative percentage coverage of the four target bodies for the ten closest flybys of each instrument of the JUICE spacecraft.

Table 1: The number of regions of interest covered by each of the Ganymede flybys by any of the instruments. If the same region of interest is covered by multiple flybys, it is counted for each flyby.

	Ganymede Flyby									
	1	2	3	4	5	6	7	8	9	10
ROI Count	19	18	22	19	8	6	7	6	6	3

Europa Clipper Evaluation with Other Tools

We performed two additional analyses for the Europa Clipper mission. The first exercises CLASP's ability to import an observation schedule and to compare the coverage achieved between the two different scheduling softwares for the same campaign, using CLASP's coverage analysis capability. The second analysis compares CLASP's coverage analysis to another coverage analysis tool.

CLASP is capable not only of generating observation times, but also of ingesting a schedule. We leveraged this capability in order to perform coverage analysis for an observation schedule generated by APGEN [Maldague et al. 2014] that used the same campaign definitions. The results of this study showed that the observations for most of the campaigns resulted in negligible differences in coverage percentage between the two schedulers. Those campaigns that had greater differences in coverage were mainly due to the different treatment of the timing of duty cycles and fixed observation durations between the two schedulers.

SIMPLEX, SIMulator for PLanetary EXploration, [Johns Hopkins Applied Physics Laboratory] is a tool developed at the Johns Hopkins Applied Physics Laboratory that is capable of calculating the surface coverage with respect to science requirements. Inputs to the instrument models and science constraints were entered into the SIMPLEX framework as similar as possible to those used for the CLASP campaigns. However, the coverage calculated by SIMPLEX revealed that the differences in modelling result in very different coverage percentage. Particularly, the different field-of-view shapes for some of the instruments and the inability to specify a duty cycle created discrepencies in the coverage results for the two different tools.

Related Work

Spacecraft operations have been a major area of application for automated planning and scheduling. Numerous space missions have used automated planning and scheduling on the ground to enable significant operational efficiencies including the Hubble Space Telescope [Johnston et al. 1993], space shuttle refurbishment [Deale et al. 1994], shuttle payload operations [Chien et al. 1999], The Modified Antarctic Mapping Mission [Smith, Engelhardt, and Mutz 2002], Mars Exploration Rovers [Bresina et al. 2005], Earth Observing One (EO-1) [Chien et al. 2005a], Mars Express [Cesta et al. 2007; Rabenau et al. 2008], Orbital Express [Chouinard et al. 2008], and Rosetta [Chien et al. 2015].

Automated planning has even flown as a technology demonstration on the Deep Space One (DS1) Mission [Muscettola et al. 1998] and as the primary operations system on 3CS [Chien et al. 2001], EO-1 [Chien et al. 2005b], and IPEX [Chien et al. 2016]. However, all of the above applications focused on the state, resource, and timing aspects of mission operations rather than automating both the spatial coverage as well the state and resource reasoning.

Notable exceptions are [Knight and Hu 2009; Doubleday and Knight 2014; Rabideau et al. 2010; Doubleday 2016] which also use the CLASP system. AEOS [Lemaître et al. 2002] also considered mapping regions but is generally focused on directing an extremely agile earth observing platform to cover small areas from a regular earth orbit generally or potentially within a single overflight. In contrast we examine a problem with a non-agile spacecraft covering large areas from a highly eccentric flyby typically in multiple overflights. In [Verfaillie et al. 2012], observation requests are expressed in terms of time instead of geometric locations, and the search explores the space of all possible instrument start and stop times. A heuristic is used to choose times that are expected to have the best impact on the plan, which includes short durations to keep resource use low. CLASP explicitly prunes times that will cause resource violations.

This paper is a continuation of work performed in [Rabideau, Chien, and Ferguson 2015], which also performed coverage analysis for the Europa Clipper mission. However, the work in [Rabideau, Chien, and Ferguson 2015] used notional payload instrument definitions, earlier trajectories, and focused on the comparison between choices of trajectories as well as imposed data volume limits. This paper revisited the coverage analysis of the Europa Clipper mission with updated instruments and trajectory and also explored coverage analysis of the JUICE mission, which has regions of interest.

Future Work

Possible future work for this application includes improving campaign definitions and updating the state and resource models to provide more accurate observation schedules that satisfy resource constraints and thus reflect more accurate coverage possibilities.

Future work with respect to the CLASP software includes representing articulated instruments such as in the Eagle Eye software in [Knight, Donnellan, and Green 2013], which describes planning and scheduling to maximize coverage by modelling both spacecraft and instrument pointing. This is particularly relevant to Europa Clipper as the EIS and MISE instruments are articulated, and additionally the E-THEMIS and UVS instruments can perform scans across the planetary disk to achieve greater coverage.

Another area where this tool could be used in the future is for mission operations, where the coverage plan could be recomputed rapidly in cases where the coverage request mix is dynamic.

Conclusions

The work in this paper presented a tool for automated scheduling and coverage analysis, which was used to asses coverage for the Europa Clipper and JUICE missions. By using geometric knowledge of the position of the spacecraft, its instruments, and a gridded representation of the observation region on the target body, CLASP is able to generate a set of observable points. Then, campaign and spacecraft constraints are applied using the Squeaky Wheel greedy optimization algorithm. This produces possible observation schedules with the accompanying surface coverage, which can be used by scientists for mission analysis. A further comparison of the CLASP tool was performed using two different tools. The first investigation used the CLASP coverage analysis using an observation schedule generated by APGEN. The second comparison used the SIMPLEX tool to both generate the observation schedule and perfrom the coverage calculation. Differences in the coverage calculated revealed that the scheduling algorithms and thus coverage achieved are sensitive the exact modeling of the instruments and cadence of the observations.

Acknowledgements

We gratefully acknowledge the contributions of Nicolas Altobelli and Claire Vallat of the European Space Astronomy Centre (ESAC)/ European Space Agency (ESA) in Villanueva de la Cañada, Spain, for both the JUICE relevant data and many informative discussions.

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

Bresina, J. L.; Jónsson, A. K.; Morris, P. H.; and Rajan, K. 2005. Activity planning for the mars exploration rovers. In *International Conference on Planning & Scheduling*. Monterey, CA.

Cesta, A.; Cortellessa, G.; Denis, M.; Donati, A.; Fratini, S.; Oddi, A.; Policella, N.; Rabenau, E.; and Schulster, J. 2007. Mexar2: AI solves mission planner problems. *IEEE Intelligent Systems* 22(4).

Chien, S.; Rabideau, G.; Willis, J.; and Mann, T. 1999. Automating planning and scheduling of shuttle payload operations. *Artificial Intelligence* 114(1-2):239–255.

Chien, S.; Engelhardt, B.; Knight, R.; Rabideau, G.; Sherwood, R.; Hansen, E.; Ortiviz, A.; Wilklow, C.; and Wichman, S. 2001. Onboard autonomy on the three corner sat mission. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2001)*. Montreal, Canada.

Chien, S.; Cichy, B.; Davies, A.; Tran, D.; Rabideau, G.; Castano, R.; Sherwood, R.; Mandl, D.; Frye, S.; Shulman, S.; et al. 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.; Davis, A.; Mandl, D.; Trout, B.; Shulman, S.; et al. 2005b. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication* 2(4):196–216.

Chien, S. A.; Tran, D.; Rabideau, G.; Schaffer, S. R.; Mandl, D.; and Frye, S. 2010. Timeline-based space operations scheduling with external constraints. In *International Conference on Planning & Scheduling*. Toronto, Canada.

Chien, S. A.; Rabideau, G.; Tran, D.; Troesch, M.; Doubleday, J.; Nespoli, F.; Ayucar, M. P.; Sitja, M. C.; Vallat, C.; Geiger, B.; et al. 2015. Activity-based scheduling of science campaigns for the rosetta orbiter. In *IJCAI*.

Chien, S.; Doubleday, J.; Thompson, D. R.; Wagstaff, K. L.; Bellardo, J.; Francis, C.; Baumgarten, E.; Williams, A.; Yee, E.; Stanton, E.; et al. 2016. Onboard autonomy on the intelligent payload experiment cubesat mission. *Journal of Aerospace Information Systems* 1–9.

Chouinard, C.; Knight, R.; Jones, G.; Tran, D.; and Koblick, D. 2008. Automated and adaptive mission planning for orbital express. In *Space Operations Conference (SpaceOps 2008)*. Heidelberg, Germany.

Deale, M.; Yvanovich, M.; Schnitzuius, D.; Kautz, D.; Carpenter, M.; Zweben, M.; Davis, G.; and Daun, B. 1994. The space shuttle ground processing scheduling system. *Intelligent Scheduling* 423–449.

Doubleday, J., and Knight, R. 2014. Science mission planning for desdyni with clasp. In *Space Operations Conference (SpaceOps 2014)*. Pasadena, CA.

Doubleday, J. R. 2016. Three petabytes or bust: planning science observations for NISAR. In *SPIE Asia-Pacific Remote Sensing*. International Society for Optics and Photonics.

European Space Agency. juice. `http://sci.esa.int/juice/`.

Fox, B. 1996. An algorithm for scheduling improvement by schedule shifting. Technical Report 96.5.1, McDonnell Douglas Aerospace, Houston. See also Scheduling Optimizer B. R. Fox US Patent 5,890,134 (1999).

Google. Google Earth. `https://www.google.com/earth/`.

Johns Hopkins Applied Physics Laboratory. SIMPLEX SIMulator for PLanetary EXploration. `http://simplex.jhuapl.edu/`. Build 6.

Johnston, M.; Henry, R.; Gerb, A.; Giuliano, M.; Ross, B.; Sanidas, N.; Wissler, S.; and Mainard, J. 1993. Improving the observing efficiency of the hubble space telescope. *AIAA Computing* 9.

Joslin, D. E., and Clements, D. P. 1999. Squeaky wheel optimization. *Journal of Artificial Intelligence Research* 10:353–373.

Knight, R., and Chien, S. 2006. Producing large observation campaigns using compressed problem representations. In *International Workshop on Planning and Scheduling for Space (IWPSS2006)*. Space Telescope Science Institute, Maryland.

Knight, R., and Hu, S. 2009. Compressed large-scale activity scheduling and planning (clasp) applied to desdyni. In *International Workshop on Planning and Scheduling for Space (IWPSS2009)*. Pasadena, California.

Knight, R., and Smith, B. 2005. Optimally solving nadir observation scheduling problems. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2005)*. Munich, Germany.

Knight, R.; Donnellan, A.; and Green, J. J. 2013. Mission Design Evaluation Using Automated Planning for High Resolution Imaging of Dynamic Surface Processes from the ISS. In *International Workshop on Planning and Scheduling for Space (IWPSS2013)*. San Jose, California.

Knight, R. L. 2005. *Solving the constrained coverage problem with flow networks as linear program approximations*. Ph.D. Dissertation, University of California at Los Angeles.

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(5):367–381.

Maldague, P. F.; Wissler, S. S.; Lenda, M. D.; and Finnerty, D. F. 2014. APGEN scheduling: 15 years of Experience in Planning Automation. In *Space Operations Conference (SpaceOps 2014)*. Pasadena, CA.

Muscettola, N.; Nayak, P. P.; Pell, B.; and Williams, B. C. 1998. Remote agent: To boldly go where no ai system has gone before. *Artificial Intelligence* 103(1-2):5–47.

National Aeronautics and Space Administration. Europa Clipper. `https://www.nasa.gov/europa`.

National Aeronautics and Space Administration, Navigation Ancillary Information Facility. NAIF/SPICE. `https://naif.jpl.nasa.gov/naif/`.

Rabenau, E.; Donati, A.; Denis, M.; Policella, N.; Schulster, J.; Cesta, A.; Cortellessa, G.; Oddi, A.; and Fratini, S. 2008. The raxem tool on mars express-uplink planning optimisation and scheduling using ai constraint resolution. In *SpaceOps 2008 Conference*, 3560.

Rabideau, G.; Chien, S.; Mclaren, D.; Knight, R.; Anwar, S.; Mehall, G.; and Christensen, P. 2010. A tool for scheduling themis observations. In *International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS 2010)*. European Space Agency Sapporo, Japan.

Rabideau, G.; Chien, S.; and Ferguson, E. 2015. Using automated scheduling for mission analysis and a case study using the europa clipper mission concept. In *International Workshop on Planning and Scheduling for Space (IWPSS2015)*. Buenos Aires, Argentina.

Smith, B. D.; Engelhardt, B. E.; and Mutz, D. H. 2002. The radarsat-mamm automated mission planner. *AI Magazine* 23(2):25.

Verfaillie, G.; Olive, X.; Pralet, C.; Rainjonneau, S.; and Sebbag, I. 2012. Planning acquisitions for an ocean global surveillance mission. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2012)*. Turin, Italy.