SpaceOps-2021,13,x1383

Using Automated Scheduling for Mission Design: A Case Study for EMIT

Amruta Yelamanchili[a]*, Christopher Wells[a], Steve Chien[a], Joseph Russino[a], Robert Green[a], Bogdan Oaida[a], David R. Thompson[a]

[a] *Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109*
{firstname.latstname@jpl.nasa.gov}
* Corresponding Author

Abstract

The Earth Surface Mineral Dust Source InvesTigation (EMIT) is an Earth Ventures-Instrument (EVI-4) mission to map the surface mineralogy of arid dust source regions. EMIT used automated scheduling technology to analyze aspects of the mission design. The automated scheduling technology was used to construct schedules which were then automatically analyzed with respect to science acquired. These analyses can be performed for a range of spacecraft hardware configurations, observation strategies, and science requirements. By studying the effects of changes on the above inputs, better hardware configurations, observation strategies, and science requirements can be formulated. The use of a pointing mirror on EMIT was under consideration, and this analysis aided in determining whether or not to keep it as part of the design of the instrument. Clouds will also have a large impact on the coverage of science targets achievable by the mission. Analysis was done on how clouds could impact the coverage achievable as well as the data volume. This analysis with clouds also aided in determining the coverage criteria for the mission. It was necessary to find a criteria that was achievable with some margin as well as satisfies the science goals of the mission.
Keywords: planning, scheduling

Acronyms/Abbreviations
Compressed Large-scale Activity Scheduling and Planner (CLASP)
International Space Station (ISS)
Keyhole Markup Language (KML)
Solar Zenith Angle (SZA)
Visible to Short Wavelength Infrared (VSWIR)

1. Introduction

EMIT seeks to map the composition of dust sources on the Earth's surface to better understand how surface dust effects radiative forcing. Radiative forcing is the difference between the energy reaching the Earth from the sun and the energy that is reflected back out into space. The maps created by EMIT will be used to better constrain Earth system models to understand and predict dust cycles and their effect on heating and cooling the Earth [1].

The instrument consists of a VSWIR imaging spectrometer. EMIT is scheduled to be launched in the early 2020's and will be installed on an exterior facility of the ISS.

Automated scheduling has been used by the EMIT mission to produce schedules which can then be automatically analysed with respect to science value attained. These analyses can be performed for a range of spacecraft hardware configurations, observation strategies, and science requirements. By studying the effects of changes on the above inputs, better hardware configurations, observation strategies, and science requirements can be formulated.

The automated scheduling technology used by EMIT is CLASP [2]. CLASP is a long-range scheduler which addresses the problem of choosing the orientation and on/off times of a space-based pushbroom instrument such that the schedule covers as many target points as possible, but without oversubscribing memory and violating any other spacecraft constraints.

CLASP has been used on a variety of other missions, for scheduling operations as well as mission design analysis. It is currently used in the operations of the Orbiting Carbon Observatory-3 [3] and ECOSTRESS missions [4], and was previously used in the operations of the Intelligent Payload Experiment Cubesat Mission [5]. CLASP is being used in the mission design of the NASA-ISRO Synthetic Aperture Radar Mission [6] and will be used for scheduling its operations after launch. CLASP was also used in coverage analysis for the Europa Clipper and Jupiter Icy Moons Explorer missions [7]. CLASP was also evaluated as a scheduler for the Thermal Emission Imaging System on the Mars Odyssey spacecraft [8].

The remainder of this paper describes the CLASP adaptation for EMIT and how it has been used in mission analysis. Section 2 describes the CLASP scheduling system. Section 3 describes the initial strategy to prioritize scheduling the highest quality observations. Section 4 describes how the effects of clouds were considered in predicting coverage achievable. Section 5 describes how a coverage criteria was defined with the effects of clouds. Section 6 describes the analysis of the effect of removing the pointing capability from the instrument. Section 7 describes the current work underway for using CLASP for operations of EMIT. Section 8 presents the conclusions.

2. CLASP for Scheduling

CLASP [1] is a scheduler for space-based instruments that can be modelled as pushbrooms – one-dimensional line sensors dragged across the surface of the body being observed. It uses the SPICE toolkit [9] for geometric computations. It addresses the problem of choosing the on/off times and the orientation of one or more instruments to observe as many target points as possible without violating any constraints, such as data volume.

2.1 Inputs

CLASP is able to schedule observations with multiple spacecraft, each of which can have multiple instruments, each of which can have multiple modes. As CLASP is used for scheduling non-agile instruments, roll capability can be defined at the spacecraft level. As early designs of EMIT included a pointing mirror, the pointing was modelled with the spacecraft rolling. Spacecraft are given roll bounds and an angular rate at which they can rotate. Instruments are given a swath size, as well as a time delta at which to split the planning horizon into potential observations. Instruments can then have multiple sensor modes, each of which can have a unique data rate. In the case of EMIT, there is only one spacecraft, one instrument, and one sensor mode.

Science campaigns are specified in the form of KML. This is the file format used with Google Earth and allows for simple visualization of the coverage achieved by a schedule. Polygons or placemarks are given with an associated description tag that defines the campaign. The following fields define a campaign:

- Name
- Priority
- Mode to observe
- Illumination Constraints
- Number of observations requested

In the analysis done for EMIT, there was no maximum number of observations desired of a particular target. For the entire target region, observations are desired to be scheduled whenever the illumination conditions are met. For number of observations, this field was given a value that exceeded the upper bound on the number of times any target would be viewable by the spacecraft in the planning horizon. The target region used initially was a dust source map previously created [10]. It has been under continual refinement by the EMIT science team. Spacecraft orbits are given in the form of SPICE kernels. For EMIT, a spice kernel containing a prediction of the ISS orbit over a year has been used. A gridded approximation of polygons is used, and the grid spacing to use is specified as an input. These gridpoints are generated over the entirety of the target body, and gridpoints inside of polygons are considered targetpoints.

Memory constraints can also be considered. Storage available onboard and downlink capability can be given as inputs, and more complex data models can be written into the mission-specific adaptation to better reflect the actual system. For the analysis presented here, data volume was not constrained at the time of scheduling. There is ongoing work into incorporating the instrument data model into the EMIT CLASP adaptation.

2.2 Observation Generation

All possible visibility swaths are first created. For each instrument defined, the planning horizon is split into visibility windows according to the time delta specified in the instrument definition. Then the targetpoints and visibility windows are intersected to create a mapping of targetpoints to the visibility windows that cover them. It is from these visibility windows that observations will be selected and added to the schedule.

2.3 Observation Selection

Targetpoints are considered in decreasing priority order for scheduling observations that cover them. Observations for a single targetpoint are considered for scheduling in time order from the start of the planning horizon. If an observation does not violate any constraints, such as data volume or illumination, it is added to the schedule.

There are additional factors considered when scheduling an observation for a spacecraft that has roll capability. The roll angle of the observation must be determined, as well as checked for whether it is compatible with observations that have already been placed in the schedule. An observation is given bounds for its roll angle that entail the set of all angles such that the targetpoint under consideration would be in view. As more targetpoints are scheduled, these bounds are refined further to maintain enough time to roll between the scheduled bounds. An observation may also be able to cover multiple targetpoints, and the roll angle bounds are refined to reflect this as well. The final roll angle could be taken as any angle in between the output bounds, but the angle midway between the bounds is generally used.

3. SZA and Observation Angle Analysis

Initial analyses considered how to best define CLASP science campaigns to maximize the quality of the science data while also reaching coverage goals and respecting data volume constraints. In these initial analyses, the instrument was planned to have a pointing mirror that could look up to 30° off nadir. It would be preferential to observe with pointing angles closer to nadir, as this gives the best resolution in the resulting data, as well as gives the best signal to noise ratio. Closer to nadir there is a stronger spectral reflectance due to superior viewing geometry, and the measurement has the least atmospheric distortion as the instrument is viewing through the least amount of air mass compared to a pointing that is farther away from nadir.

Furthermore, taking observations when the target has an SZA of less than 45° was preferred. As part of this analysis, it was explored whether coverage of the target region could be attained with this illumination constraint. It was decided observations could be taken with SZA between 45° and 60°, if that was necessary to achieve full coverage of the target region.

With the above preferences, a set of tiered science campaigns was created, with each campaign differing in its lighting constraint, its viewing constraint, and its priority (Fig. 2). CLASP does not have pointing angle built in as a constraint that can be expressed through a science campaign, so the emission angle was used as a surrogate constraint for this. The pointing angle and emission angle are roughly the same, with a small difference coming from the curvature of the Earth.

Priority	SZA	Emission Angle
1	$0 \leq x \leq 45°$	$0 \leq x \leq 10°$
2	$0 \leq x \leq 45°$	$10 < x \leq 20°$
3	$0 \leq x \leq 45°$	$20 < x \leq 30°$
4	$45 < x \leq 60°$	$0 < x \leq 10°$
5	$45 < x \leq 60°$	$10 < x \leq 20°$
6	$45 < x \leq 60°$	$20 < x \leq 30°$

Fig. 2. Table showing priorities of observational constraints

For EMIT, the scientists expressed that having a more optimal SZA for an observation was more important than having a more optimal pointing angle. This is why all the campaigns with the lower SZA constraint have a higher priority than the campaigns with a higher SZA constraint. Should the scientists have expressed an optimal pointing angle would have been preferential to an optimal SZA, the campaigns with a lower emission angle would have strictly higher priority than those with a greater pointing angle.

It is also possible to create more campaigns with a finer granularity of constraints. However, the more campaigns there are for the same set of targets, the longer it may take to produce a schedule. Targetpoints are generated from each campaign separately. A single point in the target region would correspond to six targetpoints, each of which need to go through the scheduling process.

4. Cloud Statistics

Observations taken by EMIT can be obscured by clouds, reducing the quality of the data gathered. It is the goal of the mission to achieve some level of coverage of global dust sources, and observations of a point on the Earth would not be considered successful if the observation is considered too cloudy. Thus, the likelihood of regions on the Earth being cloudy when they are observed by the instrument have a large impact on the coverage able to be attained.

The instrument will have cloud screening software onboard [11, 12]. This software can excise cloudy data, reducing the amount of data that is stored onboard as well as downlinked to the ground that has less science value. The use of this software has a large impact on data volume as well.

To analyze the effects of clouds on coverage achievable and data volume, a cloud probability mask was used. The mask used assigns to each 1° x 1° cell of the Earth that contains land a probability that the cell would be cloudy. This mask does not account for seasonal changes in cloud likelihoods. This mask is based off of the MODIS cloud mask [13], which specifies confidence levels that a point on the Earth is unobstructed by clouds.

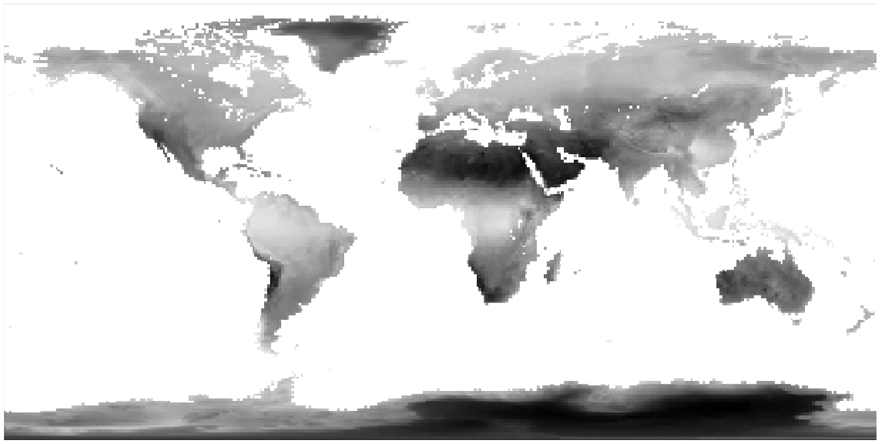

Fig. 2. Cloud Probability Mask – Darker color shows lower probability of clouds

Two approaches of incorporating the cloud statistics have been used. They have been used at the observation level, as well as the more fine-grain level of individual target points.

4.1 Cloudy Observations

The approach at the observation level has been used to determine the effect of clouds on data volume. An observation is deemed cloudy or not cloudy by determining the center point of the observation, determining its probability of being cloudy using the cloud probability mask, and generating a random number between 0 and 1 with a uniform distribution. If the random number is less than the probability, the observation is deemed cloudy. Otherwise it is deemed not cloudy. This method is less accurate than the method described in Section 4.2, but is used for its ease of determining the impact on data volume.

A CSV containing all of the observation times over a year, including a field that specified each observation as cloudy or not cloudy, was sent to the mission systems engineer to simulate the data volume over the course of the mission. It was determined that with the cloud screening software onboard, the data volume acquired would not exceed the available capacity with some margin.

EMIT observations have been between 5 and 15 seconds long, and the cloud screening software works at a much smaller scale than this. The method discussed in Section 4.2 seeks to capture the effect of clouds at a smaller scale.

4.2 Cloudy Targetpoints

To determine the effect of clouds on the level of coverage achievable, the cloud statistics were used on a finer scale. The target region EMIT seeks to cover is abstracted into targetpoints that have a 5 km spacing between each other. A single observation could cover multiple of these targetpoints, and targetpoints in different locations of the area covered by the observation could be considered cloudy or not cloudy. Rather than observations being probabilistically declared as cloudy or not cloudy, individual targetpoints are declared cloudy or not cloudy, each time they are observed. A targetpoint can be considered covered if it is declared not cloudy at least once. This targetpoint method is discussed further in Section 5.2.

5. Coverage Criteria with Cloud Statistics

Further analysis was done to aid in better defining the mission requirements for coverage. This involved two related analyses – one to determine how to define the target region, and another to determine what percentage of the target region needs to be successfully observed to meet science goals and is also feasible.

5.1 Target Definition

It was decided that the target region that the instrument would cover would be a set of 100 km x 100 km “bins”. The proposed target map would need to be abstracted to this set of bins. Because processed science data will be at a 5 km resolution, the grid spacing used for abstracting the target map to targetpoints was 5 km. Thus, the set of target bins would be chosen from the set of 100 km x 100 km bins that contain a targetpoint (Fig. 4).

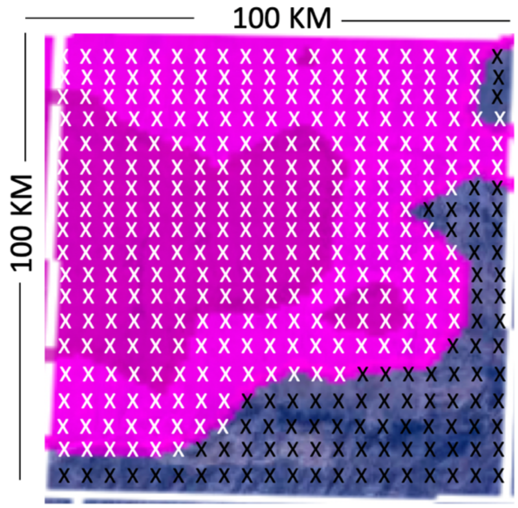

Fig. 4. Diagram showing 5 km gridpoints (x’s), targetpoints (white x’s) in the context of a 100 km bin

However, not all of these bins are of high scientific value. There may be bins that contain only one or a few targetpoints (Fig. 5). Thus, there needed to be a threshold of how many targetpoints a bin would need to contain to be considered of high enough science value to be included as part of the requirements. The chosen bin criteria would also need to meet the coverage thresholds defined in Section 5.2.

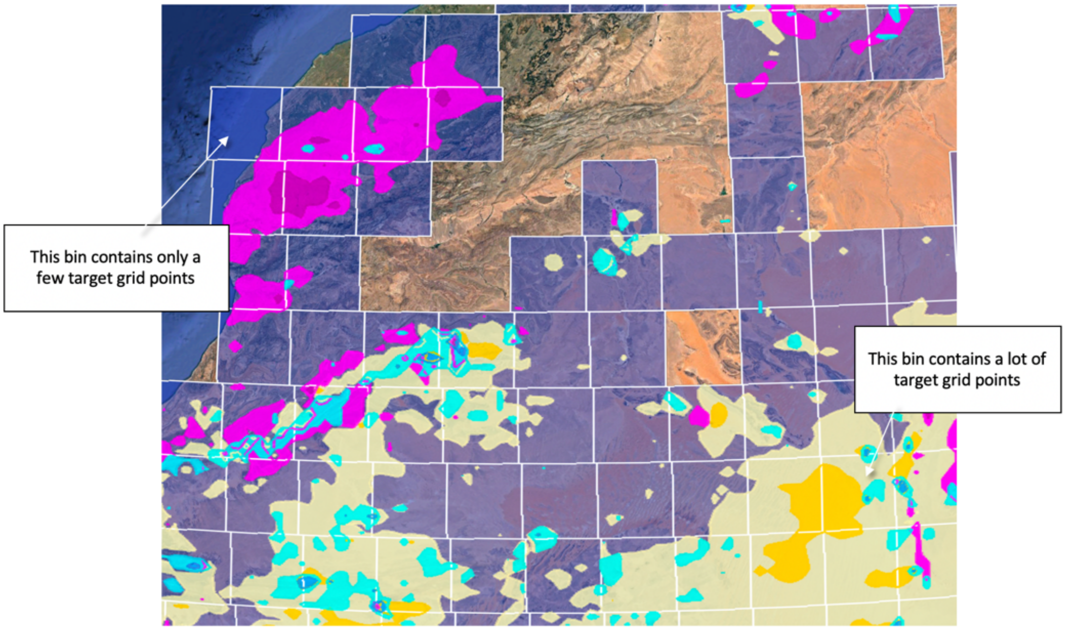

Fig. 5. Diagram showing 100 km bins in relation to original target map

Multiple schedules were created varying the target mask that was input. The input masks consisted of 100 km bins that had the following bin inclusion thresholds:

1) 1 or more targetpoints (any part of bin is within target region)
2) 40 or more targetpoints (≥ 10% of bin is contained within target region)
3) 200 or more targetpoints (≥ 50% of bin is contained within target region)
4) 380 or more targetpoints (≥ 95% of bin is contained within target region)

For each bin in the input masks, each 5 km spaced point within it is considered a targetpoint, regardless of whether it was in the original mask. The resulting schedules were then analyzed for coverage with clouds using the cloud realization method (Fig. 6, method discussed in Section 6.2).

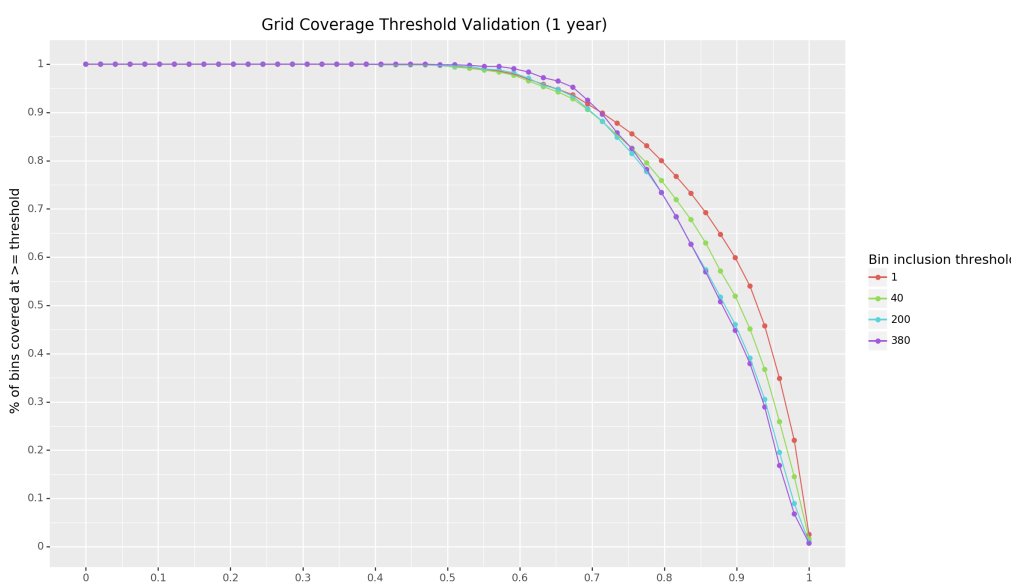

Fig. 6. Graph of coverage of target regions with different bin inclusion thresholds

It was determined that the bin inclusion threshold did not have a significant impact on the coverage able to be achieved. In subsequent analyses, the bin inclusion threshold used is 40, as this allows for most of the original target region to be included in the goal target mask, except for very small and isolated areas.

5.2 Coverage Threshold

It will likely be impossible for the entirety of the target region to be observed cloud-free in the time frame required by the mission. Thus, a lower threshold for coverage needed to be determined, and on a bin-by-bin basis. This threshold needed to be achievable in a certain time frame, while still acquiring an adequate level of science data.

A probabilistic method of determining coverage with clouds was developed that was validated by the realization method discussed in Section 4.2. For each targetpoint, the probability that there is one cloud-free observation of that targetpoint can be calculated. If there is probability c_t that targetpoint t will be cloudy at any given time, and there are n_t observation opportunities for t, the probability o_t that t will have one cloud free observation is $o_t = 1 - (c_t)^{n_t}$. Once the probability of each targetpoint having a cloud-free observation is calculated, the probability that a bin will have $x\%$ or more of its contained targetpoints have a cloud-free observation can be calculated, for a given value of x.

6. Pointing Mirror Analysis

Designs of certain hardware components have been in refinement as the mission progresses. One such component was the use of a pointing mirror. How far off-nadir an instrument can point and how large its swath is can both have bearing on the level of coverage the instrument is available to attain, as well as the time it takes to achieve a certain level of coverage.

To understand how removing the pointing capability would affect the instrument's ability to achieve coverage of the target region, three schedules were produced with the following parameters:

1) Small swath size with pointing capability
2) Small swath size with no pointing capability
3) Large swath size with no pointing capability

The third schedule served to show a theoretical upper bound on the level of coverage that could possibly be achieved. This better demonstrates how the coverage achieved with a small swath size with pointing capability compares to the coverage achieved with the same swath size but no pointing capability. The resulting schedules were analysed for coverage with the impact of clouds (Fig. 3) using the cloud method discussed in Section 4.2.

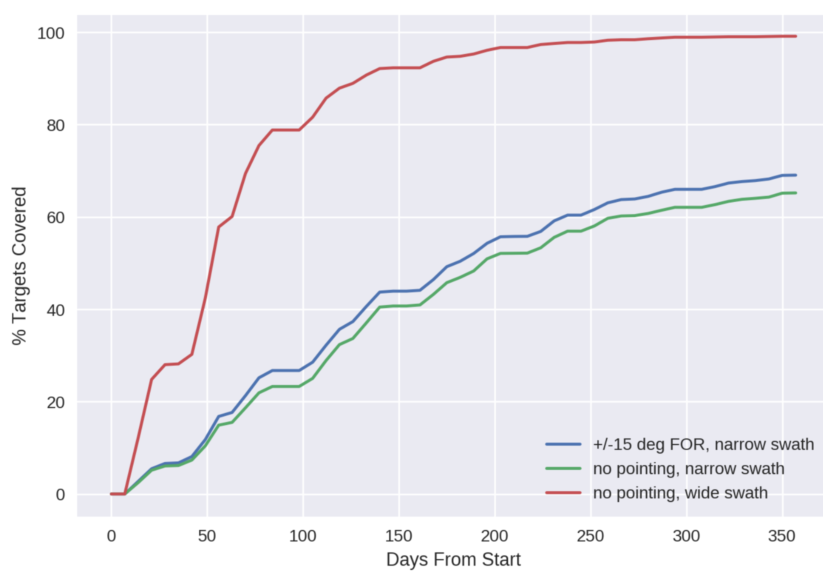

Fig. 3. Graph comparing coverage achieved with different pointing capabilities/swath sizes

The coverage attained with pointing capability is very close to the coverage attained with no pointing capability, compared to a theoretical upper bound. This analysis, along with other factors, resulted in the mission engineering team deciding to not go forward with the pointing mirror. This resulted in the mission being able to accommodate giving the instrument a larger swath size.

7. Scheduling for EMIT Operations

Development on the CLASP adaptation for scheduling operations for EMIT is underway. CLASP will schedule science observations as well as dark calibration observations. The dark calibration observations will be scheduled to optimize for a variety of factors. These observations should be temporally close to science observations, but should avoid bright areas of the Earth such as city lights and auroras. Science observations also require parameters to be calculated that will be input to the onboard cloud screening software. These parameters will be calculated as part of the scheduling pipeline.

8. Conclusion

This paper has presented how the EMIT mission has used the CLASP automated scheduling technology for analyzing and refining the mission design. Automated scheduling allows for a variety of factors to be easily changed, such as geometric constraints for observations and the capability of the instrument to point, and the automatic generation of observation schedules from these inputs. The resulting schedules can then be analyzed to see how those factors affect metrics of the schedule, such as coverage of the target area. Clouds will have a large impact on the coverage achievable by the instrument, and the ability to incorporate the likelihood of cloudy observations has given insight into how the mission can best be successful once it is in operations. This use of automated scheduling for mission analysis played a major role in the successful design of the EMIT mission.

Acknowledgements

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

References

[1] NASA, Objectives | Science – EMIT, https://earth.jpl.nasa.gov/emit/science/objectives/.
[2] R. Knight and S. Chien, Producing Large Observation Campaigns Using Compressed Problem Representations, International Workshop on Planning and Scheduling for Space (IWPSS), Pasadena, CA, 2009.
[3] A. Yelamanchili, C. Wells, S. Chien, A. Eldering, R. Pavlick, C. Cheng, R. Schneider and A. Moy, Scheduling and Operations of the Orbiting Carbon Observatory-3 Mission, International Conference on Space Operations (SpaceOps), Capetown, South Africa, 2021.
[4] A. Yelamanchili, S. Chien, K. Cawse-Nicholson, D. Freeborn, J. Padams, Scheduling and Operations of the ECOSTRESS Mission, International Conference on Space Operations (SpaceOps), Capetown, South Africa, 2021.
[5] S. Chien, J. Doubleday, D. R. Thompson, K. L. Wagstaff, J. Bellardo, C. Francis, E. Baumgarten, A. Williams, E. Yee, E. Stanton, J. Piug-Suari, Onboard Autonomy on the Intelligent Payload EXperiment (IPEX) CubeSat Mission, Journal of Aerospace Information Systems (JAIS), April 2016.
[6] J. Doubleday and R. Knight, Science Mission Planning for NISAR (Formerly DESDynI) with CLASP, International Conference on Space Operations (SpaceOps 2014), 2014.
[7] M. Troesch, S. Chien and E. Ferguson, Using Automated Scheduling to Assess Coverage for Europa Clipper and Jupiter Icy Moons Explorer, International Workshop on Planning and Scheduling for Space (IWPSS), 2017.
[8] G. Rabideau, S. Chien, D. Mclaren, D. Knight, S. Anwar, G. Mehall, P. Christensen, A Tool for Scheduling THEMIS Operations, International Symposium on Space Artificial Intelligence, Robotics, and Automation for Space (ISAIRAS), 2010.
[9] C.H. Acton, Ancillary Data Services of NASA's Navigation and Ancillary Information Facility, Planetary and Space Science, Vol. 44, No. 1, pp. 65-70, 1996.

[10] P. Ginoux, J.M. Prospero, T. E. Gill, N. C. Hsu, M. Zhao, Global-Scale Attribution of Anthropogenic and Natural Dust Sources and their Emission Rates Based on MODIS Deep Blue Aerosol Products, Reviews of Geophysics, 2012.
[11] D. R. Thompson, R. O. Green, D. Keymeulen, S. K. Lundeen, Y. Mouradi, D. C. Nunes, R. Castano and S. A. Chien, Rapid Spectral Cloud Screening Onboard Aircraft and Spacecraft, Geoscience and Remote Sensing, IEEE Transactions on. 52. 6779- 6792, 2014.
[12] M. W. Sandford, D. R. Thompson, R. O. Green, B. H. Kahn, R. Vitulli, S. A. Chien, A. Yelamanchili and W. Olson-Duvall, Global Cloud Property Models for Real Time Triage Onboard Visible-Shortwave Infrared Spectrometers, Atmospheric Techniques, 13, p. 7047-7057, 2020.
[13] S. A. Ackerman and R. Frey, MODIS Atmosphere L2 Cloud Mask Product (35_L2). NASA MODIS Adaptive Processing System, Goddard Space Flight Center, 2015.