Satellite scheduling optimization for BepiColombo STC channel target-phase

Alberto Avallone
University of Modena and Reggio Emilia - Department of Sciences and Methods for Engineering (DISMI)
Reggio Emilia, Italy
alberto.avallone@unimore.it

Francesco Sala
University of Modena and Reggio Emilia - Department of Sciences and Methods for Engineering (DISMI)
Reggio Emilia, Italy
253589@studenti.unimore.it

Maxence Delorme
Tilburg University - Tilburg School of Economics and Management (TiSEM)
Tilburg, The Netherlands
m.delorme@tilburguniversity.edu

Cristina Re
Italian National Institute of Astrophisics (INAF) - Astronomic Observatory of Padua (OAPd)
Padua, Italy
cristina.re@inaf.it

Manuel Iori
University of Modena and Reggio Emilia - Department of Sciences and Methods for Engineering (DISMI)
Reggio Emilia, Italy
manuel.iori@unimore.it

ABSTRACT

The ESA-JAXA *BepiColombo* mission is a European-Japanese effort aimed at exploring Mercury, scheduled to reach the planet in 2027. Among its multiple instruments is the Stereo Imaging Channel (STC), a dual-camera that, after a global mapping of the planet, will acquire medium-resolution color images of different targets on the hermean surface. Planning STC target observations is challenging due to the limitations in data volume and operations. It requires a careful optimization of surface observations, considering the scientific priority of each target. The scheduling problem is formulated as a mixed-integer linear programming model and addressed using greedy heuristic algorithms. The results demonstrate that optimal, constraint-compliant solutions can be generated with the model, whereas the heuristics produce moderate-quality solutions within minutes of computation. The present work represents an initial phase to develop an autonomous agent for managing STC acquisition.

KEYWORDS

Scheduling, Optimization, BepiColombo

ACM Reference Format:

Alberto Avallone, Francesco Sala, Maxence Delorme, Cristina Re, and Manuel Iori. 2026. Satellite scheduling optimization for BepiColombo STC channel target-phase. In *Appears at the International Workshop on Autonomous Agents and Multi-Agent Systems for Space Applications (MASSpace-26). Held as part of the Workshops at the 25th International Conference on Autonomous Agents and Multiagent Systems., Paphos, Cyprus, May 2026*, IFAAMAS, 4 pages.

1 INTRODUCTION

The BepiColombo mission [2] has been developed as a joint effort between the European Space Agency (ESA) and the Japan Aerospace eXploration Agency (JAXA). Its nominal science phase is scheduled to begin in 2027 and its primary objective is to study the planetary surface of Mercury.

The European suite, called *Mercury planetary orbiter*, is equipped with multiple instruments such as the Spectrometer and Imagers for Mercury planetary orbiter BepiColombo Integrated Observatory SYStem (SIMBIO-SYS), an integrated imaging system composed of three complementary scientific channels: the Visible Infrared Hyperspectral Imager Channel (VIHI), the High Resolution Channel (HRIC), and the Stereo Imagining Channel (STC) [3]. The STC mission is divided into two phases: a global mapping of Mercury's surface over two aphelions (here, an *aphelion* is defined as a predetermined period during a hermean year when the planet is farthest from the Sun), followed by a target acquisition of specific areas over two additional aphelions. The STC is a camera system composed of two optical heads, called *channels*. Each channel has two color filters (for a total of four), used during the target acquisition phase, and one panchromatic filter used during the global mapping phase.

The present work focuses on the second phase of STC operations. We modeled the task of acquiring specific areas as an image scheduling problem, a well known problem in the literature [4], and addressed it using a range of optimization techniques.

In Section 2, the problem is described from an operational perspective. Section 3 presents a complete planning approach, including the preprocessing of input data. Sections 4 and 5 describe the proposed solution techniques, while preliminary results are presented in Section 6. Finally, conclusions are given in Section 7.

2 PROBLEM DESCRIPTION

The objective of the second phase of the STC mission is to cover all possible target areas, ranked according to their scientific relevance. The targets can be acquired by the STC during one of two distinct periods when Mercury is the furthest from the Sun, known as *aphelions*. The STC mission will take place during the third and the fourth aphelions, each lasting approximately thirty earth-days (the first two aphelions being used for a global mapping phase). A target is considered covered when its entire area is acquired with all four color filters. Each STC channel image (also referred to as a *frame*) has fixed dimensions determined prior to acquisition. Because the field of view of a single frame is typically smaller than the target area, multiple frames must be acquired. To merge them into a single image, a certain degree of overlap between consecutive frames is required. In this study, it was agreed with the SIMBIO-SYS team to set the minimum overlap to 15%. A contiguous sequence of frames captured within the same orbit that meets this overlap requirement is defined as an *observation.* For larger targets, multiple observations may be required. In such cases, an inter-observation overlap constraint is enforced: for each pair of (selected) consecutive observations covering the target, the intersection must cover at least 15% of the area of the smaller observation. A *pattern* is defined as a set of observations that collectively cover a target area while satisfying all internal (along-track) and external (cross-track) overlap constraints.

There are two main constraints on the BepiColombo spacecraft. The first is the data volume generated by the acquisitions, which is limited both per earth-day (given the average estimated downlink data rate) and over the entire mission, as agreed with the SIMBIO-SYS team. In this work, we assume that this daily data limit remains constant. The second constraint limits the total number of spacecraft operations, each of which is referred to as a telecommand (TC). Hosting numerous instruments, BepiColombo must balance energy consumption by limiting the number of TCs per orbit for each instrument. Each STC observation requires a fixed set of TCs, which is the same regardless of the number of frames in the observation. Our goal is to maximize the scientific value of all acquired targets while respecting the spacecraft's physical and operational constraints.

Furthermore, the SIMBIO-SYS team has defined a strategic operation, called *fusion*, to save TCs in a given orbit. Indeed, two observations can be merged into one by keeping the STC channels continuously switched on. This approach halves the number of required TCs, since each observation has the same TC cost regardless of its length, but increases the total data volume, as some frames acquired during a fusion may cover areas that are not scientifically valuable.

3 PLANNING PIPELINE

This work presents the initial phase in the development of an autonomous agent for managing the STC mission acquisitions, which consists of evaluating various strategies to identify the most effective approach for integration into the agent.

The framework takes as input the shapefiles generated by the Simulator for Operation of Imaging Missions (SOIM) [7], a tool developed within the SIMBIO-SYS team to simulate the expected instrument pointing at specific time instances for each filter, with a minimum time interval of one second. Each frame is represented as a shapefile, generated with a time step chosen to ensure at least 15% overlap along-track. The output consists of the set of TCs to be sent to the spacecraft to initiate its operations, along with the shapefiles of the resulting coverage. These can be visualized in a geographic information system visualization tool using the SOIM outputs.

In our pipeline, we utilize the convex-hull area of the target, as shown in Figure 1. This approach allows for a fast identification of the intersection area between frames and targets, ensuring that only one (the most complete) sequence of frames is generated per target per orbit, thereby avoiding an increase in the number of TCs for irregular targets.

The framework follows a four-step pipeline: preprocessing, observation generation, single-channel integration, and solution techniques. In the following, we describe the first, second, and third steps, while the fourth step, which involves the development of dedicated optimization algorithms, is detailed in the next section.

3.1 Preprocessing

During the preprocessing phase, all frames that cannot be acquired due to the Sun's influence are excluded. If the Sun elevation angle over the horizon is outside the range 0°–180° (i.e., outside a hermean day), spacecraft operations cannot be performed due to the lack of light. To select admissible observations, we matched the possible frames with the convex hull of each target. The set of admissible frames is then divided into two groups: frames that cover at least part of a target and frames that do not cover any target. The latter group is used only when fusing observations, if necessary.

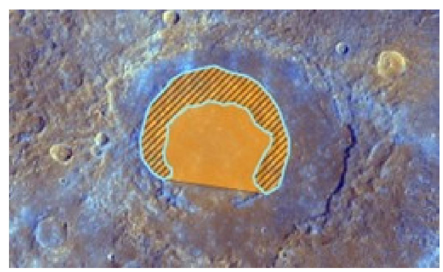

Figure 1: The original target (striped area) and the added area (not striped) which creates the convex hull

3.2 Observation generation

The observation generation phase focuses on identifying observations by grouping consecutive frames from previously selected sequences of filtered frames, as shown in Figure 2. Note that an observation may include consecutive frames covering two different targets simultaneously (see, e.g., the third observation from the right in the figure). Considering the variation in the equivalence between pixels and meters while moving towards the poles, we developed a dynamic filtering mechanism to remove unneeded frames for each observation. This mechanism determines the best possible time step while ensuring that the overlap constraint is respected.

3.3 Single-channel integration

Each STC channel acquires observations with both filters simultaneously, therefore, observations are merged in pairs of filters belonging to the same channel.

Table 1: Sets definitions

Name	Description
G	Set of all targets that can be acquired by STC
C	Set containing the two channels of STC
O	All the orbits performed by BepiColombo
D	Day of the year as recorded on earth
P_{gc}	Set of patterns covering the target $g \in G$ for the channel $c \in C$
Q_p	All the observations which belongs to a specific pattern $p \in P_{gc}$
Q_{oc}	Observations acquired by channel $c \in C$ in orbit $o \in O$
E_{oc}	Set of all fusion segments $e = (q_1, q_2)$ where $q_1, q_2 \in Q_{oc}$ are subsequent observations

Table 2: Parameters and variables definitions

Name	Description
s_g	Scientific value of the target $g \in G$
B	Maximum number of TCs for every orbit $o \in O$
V_d	Maximum data volume for each day $d \in D$
V_{tot}	Maximum data volume over the entire mission
x_q	Decision variable taking the value 1 if the observation $q \in Q_p$ is acquired, 0 otherwise
y_p	Decision variable taking the value 1 if the pattern $p \in P_{gc}$ is acquired, 0 otherwise
z_g	Decision variable taking the value 1 if the target $g \in G$ is acquired, 0 otherwise
x'_e	Decision variable taking the value 1 if the fusion observation $e \in E$ is acquired, 0 otherwise

4 SOLUTION TECHNIQUES

After the preprocessing and observation generation phases, the resolution step is applied. The task can be modeled as a weighted maximal coverage problem, which is a variation of the classical set coverage problem [6]. In our previous studies on satellite scheduling problems [1, 5], a mathematical model was defined and several heuristics were introduced. Building on these earlier works, we introduce here a mixed-integer linear programming (MILP) model and a greedy heuristic algorithm. A summary of the mathematical notation used in this section is presented in Tables 1 and 2.

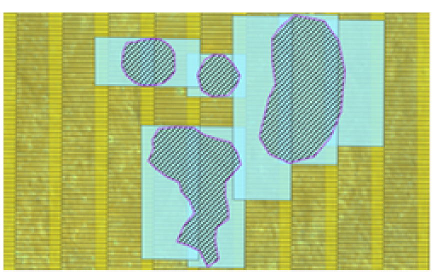

Figure 2: Generated feasible observations (cyan) over the raw sequences of frames (yellow) covering the targets (areas within the purple lines)

4.1 Mathematical model

The problem is defined as an MILP model that aims to maximize the accumulated scientific rank s_g of the acquired targets over all $g \in G$ possible targets. A target g is considered acquired ($z_g = 1$) only if the model selects a pattern $p \in P_{gc}$ that covers g with each channel $c \in C$. A pattern p can be selected ($y_p = 1$) only if all its observations $q \in Q_p$ are selected ($x_q = 1$). As defined in Section 2, the model must satisfy specific operational constraints. For each orbit $o \in O$, the total number of selected observations must not exceed the maximum number of TCs, B. Furthermore, for each day $d \in D$ and over the entire mission, the data volume of the selected observations must not exceed, respectively, the maximum daily capacity V_d and the cumulative data limit V_{tot}. To speed up the model resolution, we introduce a filtering mechanism that removes *dominated patterns*.

4.1.1 Dominant pattern selection . The dominant pattern selection method eliminates all patterns that are *dominated* by another valid pattern. A pattern $p' \in P_{gc}$ is dominated by a pattern $p'' \in P_{gc}$ if p'' is a subset of p' for the same target. Such dominated patterns should never be selected, as they increase data consumption without providing additional coverage compared to their dominating counterparts. By identifying and removing these patterns, the number of generated patterns for each target is reduced, which in turn lowers the computational complexity. This selection process is illustrated in Figure 3.

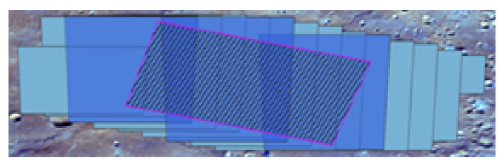

Figure 3: The selected coverage with 3 observations (dark blue) compared to 15 possible (light blue) covering the target (area within the purple line)

4.2 Greedy heuristic observation selection

The general framework of greedy heuristic algorithms can be described as follows. First, the available data volume and number of TCs are initialized, and the targets are sorted according to their scientific value. Then, starting with the highest-ranking target, the framework attempts to generate a feasible coverage pattern that minimizes cross-track overlap among observations. To construct the pattern, observations are first sorted and then added iteratively, provided no constraints are violated, until the target is fully covered. We implemented two sorting methods for the observations:

- *Orbit order search:* Observation are sorted by ephemeris time of execution. The process begins by selecting the last observation that covers the first edge of the target, and then iteratively selects subsequent observations that minimize overlap while adhering to the 15% threshold.
- *Point area search:* Observations are sorted by the proportion of the target area they cover. The process begins by selecting the observation that covers the largest portion of the target (typically located at the center of the target), and then iteratively selects subsequent observations that minimize overlap while adhering to the 15% threshold, proceeding first to the left and then to the right of the initial observation.

Note that, since the solution is constructed iteratively, it cannot be guaranteed that a pattern will be fully acquired, as the data and TC constraints may result in only partial coverage of some targets. The scientific value s_g of a target is collected only if the target is fully acquired.

5 FUSION

A particularly important consideration for the SIMBIO-SYS team is the trade-off between memory consumption and TCs, with memory being relatively abundant, while TCs are expected to be limited, as the satellite serves multiple scientific missions, each of which also requires TCs.

To allow fusion in the model, we enumerate every pair of consecutive observations (which, by construction, cover distinct targets) having the same orbit $o \in O$ and channel $c \in C$ and add the corresponding fusion segment into the set E_{oc}. If a fusion segment $e \in E_{oc}$ is selected, the TC count is reduced by two, and both the daily and total data volume are increased according to the size of the segment.

The fusion concept has not yet been tested for the MILP model, but we integrated it into the greedy heuristic as follows: when checking whether an observation satisfies all constraints (before adding it to a pattern), if it violates the TC constraint but not the memory constraint, we attempt to find the two observations in the same orbit whose fusion would minimize the increase in data volume, and fuse them.

6 RESULTS

The model-based ("Exact") and greedy heuristic ("Greedy") solution methods were tested to evaluate their operational performance, as summarized in Table 3, using the list of targets defined for the STC mission during the third aphelion. Both methods were implemented in Python and given a time limit (TL) of 3600 s. The model was solved to optimality using the commercial solver GUROBI. All computations were performed on a computer equipped with an AMD Ryzen 3700X CPU and 32GB of RAM, running Windows 11. The model was tested both with the dominant pattern selection approach ("Dominant") and without it ("No Filter"). The heuristic was tested in three configurations: orbit order search ("Orbit"), point area search ("Point"), and a combination of orbit order search and observation fusion ("Orbit-Fusion"). For each model, we report the number of variables ("#Var") to assess the effectiveness of the preprocessing, as well as the upper bound ("UB") value at which the solver terminates. For context, the third aphelion includes 195 targets (out of 404 in total) with a cumulative scientific value of 701.

Table 3: Results for Aphelion 3 instance

Method	Type	#Var.	Objective	UB	Time(s)
Exact	No Filter	1 093 368	678	699	TL
Exact	Dominant	7156	678	678	2162*
Greedy	Orbit	-	533	-	15
Greedy	Point	-	530	-	17
Greedy	Orbit-Fusion	-	543	-	82

*Excluding filtering time (12.83 s)

We observe that, without pattern filtering, the model cannot be solved to optimality within the time limit because of the very large number of generated patterns (and, consequently, variables). Applying the filtering mechanism has a significant impact, reducing the number of variables by 99.35% and allowing the model to be solved within the time limit. Regarding the heuristics, their performance is mixed. Although they are relatively fast, they produce solutions of moderate quality, with an average optimality gap of around 20%. Among the three heuristics tested, "Orbit-Fusion" achieves slightly better results but also requires the longest computation time.

7 CONCLUSION

In the present work, the BepiColombo STC mission problem is defined, and both a mathematical model and a set of greedy heuristics have been developed. Whereas the greedy heuristic reached only moderate-quality solutions, their fast running times and the fact that they do not require a commercial solver make them especially useful tools for gauging the viability of planning strategies for the identified targets. They also allowed for an estimation of the potential gains offered by the fusion technique; although computationally more demanding, this approach appears to be highly beneficial. We are currently working to integrate this process into the SIMBIO-SYS planning pipeline for application across the entire mission as an autonomous planning tool. We also plan to develop a metaheuristic that is expected to be both faster than the exact model and more effective than the current heuristics, as well as to test an extension of the model incorporating the fusion technique. We will further apply our methods to an extended set of 650 targets recently provided by the SIMBIO-SYS team. Finally, we plan to analyze potential future deployment conditions to perform a sensitivity analysis and evaluate the operational capability of our algorithms.

ACKNOWLEDGMENTS

We gratefully acknowledge financial support by NEXT Ingegneria dei Sistemi S.p.A and by the Italian Ministry of University and Research, under project PRIN PNRR 2022, grant n. P2022XF72W, "Combining mAchine Learning and optImization for Planetary remote Sensing missiOns (CALIPSO)".

REFERENCES

[1] A. Avallone, B. Ferrari, A. Cicchetti, M. Delorme, M. Iori, M. Lippi, and R. Orosei. 2025. Stochastic Optimization of the MARSIS Observation Scheduling Problem using a Predict-then-Optimize Approach. In *Proceedings of the 14th International Workshop on Planning & Scheduling for Space (IWPSS 2025)*. Toulouse, France. Paper presented at IWPSS 2025; detailed program available online.

[2] J. Benkhoff and al. 2021. BepiColombo - Mission Overview and Science Goals. *Space Science Reviews* 217, 8 (2021), 90. https://doi.org/10.1007/s11214-021-00861-4

[3] G. Cremonese and al. 2020. SIMBIO-SYS: Scientific Cameras and Spectrometer for the BepiColombo Mission. *Space Science Reviews* 216, 5 (2020), 75. https://doi.org/10.1007/s11214-020-00704-8

[4] B. Ferrari, J.-F. Cordeau, M. Delorme, M. Iori, and R. Orosei. 2025. Satellite Scheduling Problems: A survey of applications in Earth and outer space observation. *Computers & Operations Research* 173 (2025), 106875. https://doi.org/10.1016/j.cor.2024.106875

[5] B. Ferrari, M. Delorme, M. Iori, M. Lippi, and R. Orosei. 2025. *A Predict-then-Optimize Approach for the Research of Underground Water on Mars.* Technical Report. University of Modena and Reggio Emilia (UNIMORE).

[6] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming* 14, 1 (1978), 265–294. https://doi.org/10.1007/BF01588971

[7] E. Simioni, M. Zusi, and R. Politi. 2022. *SOIM User Manual v6.0.* Technical Report 202. INAF. https://doi.org/10.20371/INAF/TechRep/202