18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025.

SpaceOps-2025, ID # 452

MAPS, an automatic and robust multi-agent scheduler for large observation constellation

Vincent Debout[a]*, Nicolas Precigout[a], Fabien Scelles[a], Javier Pelegrin Garcia[a], Guillaume Durand[a], Geoffroy Periot[a], David Bourrel[a], Guillaume Renard[a], Emilie Converset[a]

[a] *CS-Group, France*, vincent.debout@cs-soprasteria.com,
* Corresponding Author

Abstract

MAPS is a mission scheduler developed by CS-Group, designed to accommodate current and future constellation missions. Built on a microservice architecture, MAPS is scalable to meet the growing demands of constellation missions. The operator's activity is reduced to the asynchronous submission of mission requests, with multi-agent scheduling processes running periodically and automatically to update the mission plans of each satellite.

The application of MAPS for a commercial constellation dedicated to Earth observation is presented. Clients submit customized requests that may include specifications like repetition frequency, acquisition geometry constraints, area shapes, or geographic lines. MAPS splits these requests into atomic acquisitions, which are then concurrently scheduled across the constellation's satellites. A comprehensive metric is designed to evaluate the global plan, considering the request parameters (such as priority and due date) and the constellation's resources management like memory load and energy levels. Planning constraints include forbidden time slots (e.g., station keeping, housekeeping, calibrations) and slew duration between consecutive acquisitions. Additionally, X-band downlink antenna passes must be scheduled to ensure data download, possibly concurrently with nearby acquisition targets.

A significant analysis phase is conducted to reduce the problem size, which includes partitioning the problem into smaller sub-problems by identifying groups of acquisitions with minimal external competition. Resource interaction still must be considered.

The plan optimization can be achieved by different approaches. The main method is a greedy algorithm that prioritizes acquisitions based on their individual score, while remaining within their accessible time windows, considering the constraints imposed by the preceding and subsequent acquisitions. This method is robust and adapts well to various constraints with low computational complexity. The resulting solution may be further refined using local search algorithms, such as simulated annealing, genetic algorithms, or local permutations. Additionally, we present a constraint programming approach and its effectiveness on sub-problems.

Keywords: Mission Planning, Mission Centre, Scheduling

Acronyms/Abbreviations
Mission and Planning Software (MAPS)
Ground Operation System and Mission Intelligence Centre (GOSMIC)
Data Take Opportunity (DTO)
Earth Observation (EO)
TeleCommand (TC)
TeleMetry (TM)
Flight Dynamic System (FDS)
Orekit Flight Dynamic System (Oreflids)
Key Performance Indicator (KPI)
Service-Level Agreement (SLA)
Constraint Programming (CP)
Mixed-Integer Linear Programming (MILP)
Iterative Local Search (ILS)

1. Introduction

Satellite constellations offer improved geographic and temporal coverage for space missions. However, managing a potentially large number of satellites requires an efficient approach to schedule multi-agent coordination for completing common tasks. Additionally, operating a constellation can be highly demanding, as each satellite has individual tasks to perform. Automating satellite operations significantly reduces this workload, as demonstrated by [1]. Recent research has extended this automation to entire constellations [2, 3].

To address these challenges, CS Group develops MAPS, an automated Mission Centre designed for large satellite constellations. Built on a scalable microservice architecture, MAPS enables the automatic computation of constellation-wide mission plans, which are triggered by the Ground Segment scheduler. As a result, operators' tasks are minimized, focusing primarily on analysis and asynchronous modifications of the mission tasks list.

The mission planning algorithms in MAPS rely on the low-level space dynamics library Orekit [4]. This Mission Centre is part of CS Group's GOSMIC product line [5], which shares a similar and compatible architecture. It integrates MAPS with the FDS Oreflids and the Control Center CSNano.

MAPS is a versatile solution applicable to most EO constellations with homogeneous or heterogeneous satellite roles, also potentially extendable to telecommunication and in-orbit services constellations.

This paper is structured as follows: first, we present the planning request process, detailing how customer demands are decomposed into atomic sub-tasks and the access computations. Next, we describe the chronology of automatic scheduling and mission plan generation. Finally, we discuss the schedule optimization process and its compatibility with a wide range of optimization algorithms.

2. Task request.

A database maintains a record of all mission tasks assigned to the constellation. Asynchronous actions, with respect to the automatic planning, are managed via microservice to add, remove or modify them. Each task constitutes a full customer demand, with an observation mode, geometry constraints, a geographic area, a priority level (KPI rules), a temporal validity range and a potential repeat period.

MAPS supports most classical observation modes to align with instrument specifications and accommodate complex calibrations. In this paper, we focus on PushBroom observation, widely used in recent EO hyperspectral missions. This mode allows for various acquisition patterns, adjusting both the length and heading of the sweeping leg. Additionally, forward motion compensation can be applied, using pitch rate adjustments to extend acquisition duration over a local target, thereby enhancing the signal-to-noise ratio [6]. During acquisition, MAPS calculates the optimal guidance, including pre-observation stabilization. While the standard approach involves along-track compensation for Earth motion, customized headings can be used for specific requirements.

From instrument point of view, target visibility is evaluated considering all relevant geometry constraints. These include the target's relative position within the instrument’s frame (e.g., distance, pointing angle) and, conversely, the instrument’s position in the target’s ground frame (elevation, azimuth). MAPS can also handle illumination constraints like local solar time or Sun elevation or even instrument solar angle with respect to Sun specular reflectivity (glint angle [7]).

Each task is assigned a priority level and/or multiple KPIs, which influence the scheduling process. The task’s time horizon defines the period within which it must be completed. For recurring tasks, a repeat period can be configured.

2.1 Task decomposition

To efficiently distribute tasks among the constellation, each task is divided into sub-tasks that can be completed continuously by a single satellite. For large geographic areas using push-broom acquisition, a tiling process divides the region into atomic meshes oriented along selected heading (likely along-track). Each mesh has a width approximately equal to the instrument swath and a limited length. Meshing parameters include the heading, a maximum length limit and overlap settings along-track and across-track to ensure potential image reconstruction. A similar decomposition process is applied for other types of mission. MAPS also supports collaborative task execution, such as stereoscopic observations.

The task decomposition is stored in the database and remains unchanged thereafter, allowing the system to track the completion of each sub-task. MAPS provides the operator with various monitoring tools to assess the overall completion of the different tasks (both in real time and as expected after the current mission plan). It may be helpful to modify task priorities to prevent them from being unachieved. Task constraints and time horizon are declined to the sub-tasks which can be scheduled, into the mission plan, in continuous time slots.

For tasks requiring the observation of linear geographic features (e.g., borders, rivers, roads) with a width smaller than the instrument swath, the default along-track heading may result in many short-length meshes, especially for orthogonal lines. Instead, using a custom heading aligned with the feature’s mean direction can significantly reduce the number of meshes, as illustrated in Fig. 1. However, this approach may also lead to fewer access opportunities and increased image distortion. MAPS includes an analysis tool to assist in selecting the optimal trade-off for such cases (see section 4.3).

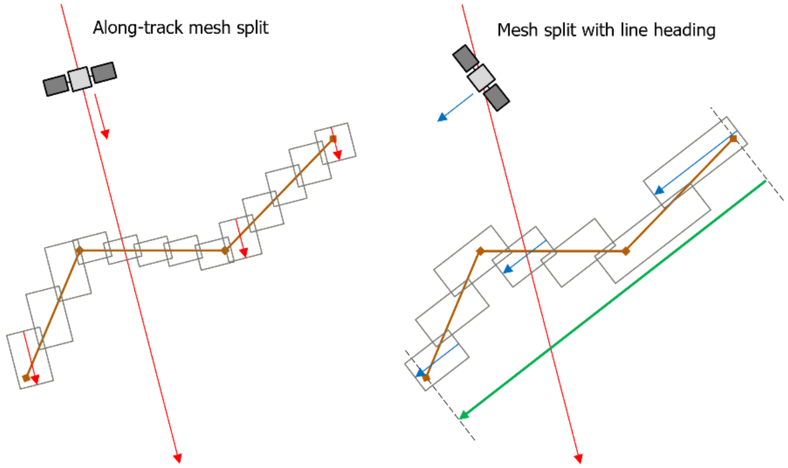

Fig.1 Headings of the mesh splitting

2.2 Access computation

Data Take Opportunities (DTOs) define the time windows when a satellite can execute a sub-task while respecting all constraints and the task's time horizon. Each satellite's orbital ephemerides are provided by an external agent, such as Oreflids, the FDS of GOSMIC. When the sub-tasks of a new submitted task are created, their DTOs are computed and stored in the database. This computation is resource-intensive but can be easily parallelized. Once this asynchronous step is completed, the task is activated in the automatic mission plan computation.

Calibrations and station flybys (for TM and TC) are treated similarly to sub-tasks, with their access windows precomputed and stored in the database.

3. Timeline

The full constellation mission plan is automatically computed at a customizable rate, up to once per orbit period (approximately 100 minutes). The process begins with a census phase, during which all DTOs are retrieved from the database, considering only tasks and ephemeris updates with fully completed computations. A scoring phase is processed to evaluate the KPIs and the score of each DTO. Next, a full constellation schedule is generated using the optimization process detailed in Section 4. This schedule is then split into individual mission plans for each satellite, generating the required TCs. At the next TC opportunity, the latest satellite mission plan is uploaded onboard. The overall process timeline is illustrated in Fig. 2.

DTOs updates are scheduled whenever new ephemeris data is provided. In the absence of manoeuvres, orbital shifts remain small, typically only a few meters for sub-daily updates. In this case, the DTO time-shift is searched in the vicinity of the previous one, reducing significantly the computation time. For this reason, embryo of DTO, local minimum slightly violating a constraint are kept in base, as they may transform to acceptable DTO. On the other hand, the added time slot at the end of the horizon is fully computed.

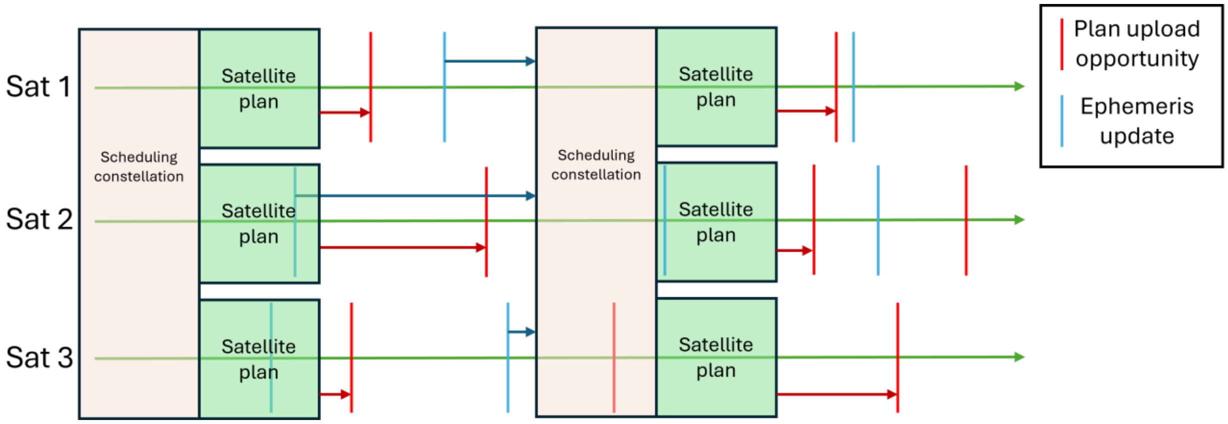

Fig.2 Timeline of the automated constellation scheduling

While high-priority tasks ensure that certain observation requests are included in the next planning cycle, emergency requests may require even faster data acquisition and customer availability. To accommodate such requests, MAPS allows for manual emergency sub-planning, which must be executed within a reduced computation time.

For each affected satellite, the time horizon starts at the earliest possible station uplink pass, potentially saving several hours with respect to the next nominal plan update. A sub-plan is then generated, considering only the emergency request and any available downlink options. If the sub-plan is feasible and validated, it overwrites the regular mission plan within the affected time slots. To meet resource constraints, some scheduled acquisitions may need to be cancelled (see relaxation in section 4.2). The updated mission plans are then uploaded to the satellites as soon as possible.

4. Schedule optimization

The scheduling problem for a single satellite is generally NP-hard [8]. Extending the problem to a constellation introduces a combinatorial multi-agent dimension. Compared to the classical job-shop scheduling problem, each satellite is subject to several resource constraints, which are typically nonlinear (e.g., AOCS, on-board memory, power, thermal management, etc.).

4.1 Greedy scheduling

For each mesh, every DTO is evaluated and given a score built with a weighted sum of its KPIs. These KPIs may characterize the customer demand (priority ...), the mesh (geographic position, landscape type ...), the DTO (forecast cloud coverage, time before expiration, observation roll ...) or even possibly the agent (instrument, satellite resources …).

MAPS includes a generic greedy algorithm for generating the constellation mission plan. This algorithm considers every DTO in decreasing score order, inserting it into the plan if possible. For each satellite, only the ordered sequence is tracked, allowing freedom within the access window for each acquisition. The insertion request is evaluated in linear time, resulting in quadratic complexity when constructing the full plan. Each insertion potentially reduces the flexibility windows of previously inserted DTOs, taking into account any necessary delays (e.g., rally time, tranquillisation time …). Once all DTOs are either inserted or rejected, each inserted one is fixed within its remaining freedom window, resulting in a scheduled acquisition.

This classical approach [9] generates a fast and robust plan and is compatible with most specific mission constraints. Resources of the different satellites are monitored as the plan is growing and add an additional constraint to consider at each insertion. Moreover, it does not need an explicit metric for the plan.

4.2 Generic optimization problem

If a metric for the global plan can be defined, MAPS can extract the mathematical optimisation problem and apply other optimisation processes. Defining such a metric may require a specific mission analysis, as it should reflect the overall long-term mission efficiency. The main component of this metric is expected to be the weighted sum of achieved acquisitions.

The scheduling problem can then be solved using custom algorithms or more standard MILP methods, significantly improving the optimality of the solution. The greedy approach can still be used to initialise these approaches or ILS algorithms.

Some of these methods are not relevant for the full problem but adapted to size-limited sub-problems. In lot of missions, due to large areas without acquisition request (i.e. oceans for EO on emerged ground), the meshes ensemble may be cut into several subgroups with minimum interaction. By “minimum interaction”, it is meant a sub-group of meshes which DTO have no time overlap with any DTO of any mesh external to this sub-group. In other words, it is a sub-group of meshes with no external time concurrence, analogous to graph connected component. These components constitute sub-problems that may be optimized more efficiently by ILS or mathematical programming, possibly using parallelized computation [10]. As resources are renewable and subgroups are temporally distant by design, resource interaction between subgroups still needs to be considered.

Resources generally have a very specific behaviour and are possibly challenging to transcript to mathematical expression. A simplified expression can be provided to MAPS to make efficient combinatorial exploration. While it may bound the constraint, an additional relaxation loop may be necessary. It consists of removing some of the scheduled acquisition until the plan becomes feasible. This may be done greedily or with a more comprehensive algorithm.

4.3 Analysis

MAPS provides an analysis service to support operations. It estimates the efficiency of the constellation in achieving a potential task. Various parameters of the request can be comprehensively customised to optimise task decomposition (see Section 2.1) and completion likelihood. In the case of an emergency plan, the system also provides an overview of the consequences on the current plan before validation.

5. Conclusions

MAPS is a new Mission Centre developed to address the growing demand for commercial constellations. Its web-based, microservice architecture is highly parallelisable, enabling it to scale effectively large number of satellites. MAPS can manage a wide range of Earth Observation missions with possibly heterogeneous constellation. This paper presents an example of its application in Earth Observation, using PushBroom scanning. The customers' demands are systematically divided into atomic tasks, which can be efficiently completed by any agent.

MAPS automates the constellation planning process, minimising manual work and maximising the potential frequency of plan updates. Operator involvement is limited to the asynchronous modification of the requests pool. The Mission Centre also accommodates "emergency requests", allowing it to temporarily exceed the automatic scheduling rate to meet specific demands.

The scheduling process organises the activities of the various satellites to create an efficient global plan. MAPS enables operators to classify demands according to their relative importance (e.g., priority, expiry date …), and solves the multi-agent scheduling problem using a simple yet robust greedy algorithm. A more detailed metric can be worked out for the plan, enabling the use of complex algorithms, including metaheuristics or mathematical programming solvers. To facilitate the planning process, connected components of tasks are identified, enabling the resolution of smaller sub-problems in parallel.

References

[1] Chien, Steve & Johnston, Mark & Frank, Jeremy & Giuliano, M. & Kavelaars, Alicia & Lenzen, Christoph & Policella, Nicola. (2012). Generalized Timeline Representation, Services, and Interface for Automating Space Mission Operations.

[2] Topal E.; Yeniçerioğlu B. S.; Çelik S.; Sayali D.; Kizilkaya V.: A New Approach for Mission Planning of Satellite Constellations. Enabling Platform Technologies for Resilient Small Satellite Constellations for NATO Missions, NATO STO Review, Spring 2022.

[3] Yang, Q.; Song, B.; Chen, Y.; He, L.; Wang, P. A Distributed Autonomous Mission Planning Method for the Low-Orbit Imaging Constellation. Algorithms 2023, 16, 475. https://doi.org/10.3390/a16100475

[4] L. Maisonobe, B. Cazabonne, E. Ward, R. Serra, M. Journot, S. Dinot, G. Bonnefille, T. Neidhart, C. Jonglez, M. Rutten, yjeand, J. Hernanz, lirw1984, V. Cucchietti, A. Goetz, Guiuux, gaetanpierre0, L. Næsbye Christensen, P. Hyvönen. (2024). CS-SI/Orekit: 12.2 (12.2). Zenodo. https://doi.org/10.5281/zenodo.13950582

[5] N.Precigout, E.Converset, MAPS, multi mission centre, 2025, this issue #451

[6] A. Abdollahi, M. Dastranj, A. Riahi (2014). Satellite Attitude Tracking for Earth Pushbroom Imaginary with Forward Motion Compensation. International Journal of Control and Automation. 7. 437-446. 10.14257/ijca.2014.7.1.39.

[7] T. Várnai , A. Marshak, A. Kostinski (2021). Operational Detection of Sun Glints in DSCOVR EPIC Images, Frontiers in Remote Sensing, vol.2. 10.3389/frsen.2021.777806
[8] Qamar, A., Salah, E.E., Badran, K.M. and ElNashar, G.A. (2020) 'Mission planning and scheduling for Earth observation space system', Int. J. System of Systems Engineering, Vol. 10, No. 1, pp.24–38.
[9] J. Bonnet, M-P Gleizes, E. Kaddoum, S. Rainjonneau, G. Flandin. (2015). Multi-satellite Mission Planning Using a Self-Adaptive Multi-agent System. Ninth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO'15)}. Self-Adaptive and Self-Organizing Systems (SASO), 2015 IEEE 9th International Conference on. 11-20. 10.1109/SASO.2015.9.
[10] Robert, V., Hertz, A. (1996). How to decompose constrained course scheduling problems into easier assignment type subproblems. In: Burke, E., Ross, P. (eds) Practice and Theory of Automated Timetabling. PATAT 1995. Lecture Notes in Computer Science, vol 1153. Springer, Berlin, Heidelberg. https://doi.org/10.1007/3-540-61794-9_71