Integrating Energy Optimal Control into the Planning Loop for Agile Satellites

Bjørn Andreas Kristiansen[1], **Simen Berg**[1], **Jan Tommy Gravdahl**[1], **Tor Arne Johansen**[1]
[1] Department of Engineering Cybernetics,
Norwegian University of Science and Technology (NTNU)
Trondheim, Norway
{bjorn.a.kristiansen, simen.berg, jan.tommy.gravdahl, tor.arne.johansen}@ntnu.no

Abstract

In this article, we integrate a recent net power optimal control algorithm into an agile satellite as a motion planning strategy where the calculations are performed in the ground segment. The net power optimal control strategy relies on solving a nonlinear program to produce an energy optimal attitude trajectory, which balances energy harvesting with the energy consumed when executing attitude control tasks. Optimal trajectories are tracked using a standard on board attitude controller. The main benefit of the net power optimal control algorithm is that it can yield a higher energy efficiency in the satellite than a standard controller, which is particularly a benefit in small and agile satellites. The optimal control problem is non-convex, and we present a strategy for producing initial guesses for this control problem that aims to find a local minimum that is better than the baseline solution. This strategy increases the robustness of the control strategy with respect to finding optimal trajectories for the net power optimal control problem. When the method is implemented on a small satellite in orbit, extra constraints are included in the problem formulation to ensure that the on board star tracker does not move into attitudes where it may become blinded by the Sun. In-orbit experiments show that the method produces expected behavior for both attitude and power.

Introduction

Energy harvesting is one of the main aspects of controlling a satellite. From a control perspective, the tasks of an Earth observation satellite can be split into pointing the satellite to a point on the ground, communicating with a ground station, and harvesting energy by pointing the satellite towards the Sun. Getting enough energy is important for any mission, so the control algorithm that is used to perform the maneuver can make a significant impact. Looking into optimal control methods is reasonable as even gaining a small amount of energy can make a difference for missions that are constrained by their energy budget. As was highlighted in Kristiansen, Gravdahl, and Johansen (2021), energy optimal attitude control algorithms have previously been focused on minimizing the norm of the control input (Wu and Han 2019) or power models (Schaub and Lappas 2009). This is unlike the energy harvesting tasks for satellites, as it focuses on using the actuators as little as possible, rather than ending up with as much energy as possible. Alternative control formulations that use time-optimal maneuvers such as in Bilimoria and Wie (1993) can be adapted to solve the problem, but they do not consider the entire energy harvesting task. This is done in Rigo et al. (2021) considering the model of solar power, and in Kristiansen, Gravdahl, and Johansen (2021), and Kristiansen et al. (2024), including the cost of the input.

It should be noted that solving optimal control problems such as these to perform satellite maneuvers is a feasible strategy, as NASA demonstrated on their Lunar Reconnaissance Orbiter in 2024 (Karpenko et al. 2024). In this article we use the method in Kristiansen et al. (2024) for motion planning for the HYPSO-1 satellite, where the novel contributions are as follows: we show how the energy optimal control scheme from Kristiansen et al. (2024) can be implemented for planning satellite maneuvers, in this case by integrating it into campaign planning software. The optimal control problem is extended and adapted to the operational setup of a small satellite mission, with the integration in a campaign planner pipeline and adding an extra constraint to enable safe satellite operations. The optimal control problem is solved as a nonlinear program, and because of its formulation, the problem is highly sensitive to initial guesses and the selection of several internal variables. Owing to this, the optimization module is designed in such a way that the optimal control problem will be rerun with a new configuration if it fails to find a solution, or if it fails to find a solution better than the proportional–derivative (PD) controller solution we use as a baseline. In the end, we verify through experimental results on a small satellite that a plan based on the energy optimal attitude control scheme yields the results which was previously only shown in simulation.

The remainder of the article follows the following structure: In the Method section, we discuss the place of our solution in the larger system structure known as the campaign planner. We show how this structure can be envisioned using a level of abstraction that allows for modularization, and where the energy optimal attitude planner module would fit into the system. The optimal control problem that is solved is then described. The description in this part is based on the work found in Kristiansen, Gravdahl, and Johansen (2021) and Kristiansen et al. (2024), but is included here to highlight the changes that are made to make the system ready for operations. The design of the optimization module, focusing on solving the optimal control problem required to find the energy optimal trajectory based on net power then follows in its own section. The section also describes the measures we take to make the system more robust due to the complexities that come with non-convex optimization. Then, the planner is tested in a maneuver as described in the Setup section. The Results section shows how the method performs on the small satellite HYPSO-1 (Grøtte et al. 2022). The Discussion section contains a brief discussion, before the paper is concluded in the Conclusion section. Additional information from Kristiansen et al. (2024) that is required to reproduce the results is included in the Appendix.

Method

The method that is implemented in this article is based on the optimal control problem in Kristiansen, Gravdahl, and Johansen (2021) and Kristiansen et al. (2024). More specifically, the objective is to implement the method for the method and complexity as presented in Kristiansen et al. (2024), but for the satellite considered in Kristiansen, Gravdahl, and Johansen (2021). The main results from the simulation studies show that the method can improve the energy gained by the satellite over a maneuver by exploiting the transient phases of the maneuver better than alternative schemes. A consequence of this result is that the method yields higher gains the more maneuvers the satellite performs. We want to emphasize that the method can be just as useful and needed for other attitude maneuvers, but we choose to display it in a pure energy harvesting setting here for clarity. The optimal control problem is singular, and the main consequence of this is that the problem is difficult to solve numerically. To smooth the implementation of the method in the operations of the satellite, we present the optimal control problem and how we try to mitigate some issues with local minima and regions of infeasibility in this section.

The Campaign Planner

The campaign planning system is based on Berg et al. (2023). The campaign planner considers cloud coverage, communication constraints, and many other aspects that are required for the satellite to function. For this article, we only look at the attitude planning capabilities of the system. An abstraction of the planner for our purposes can be seen in Figure 1.

The information flows as follows: The campaign planner uses external data as input, such as weather forecasts to detect cloud coverage. The campaign planner decides what maneuvers should be performed in what order, and sends the information to Mission Control Software (MCS), which then sends the commands to the satellite. Telemetry is then returned from the satellite to the MCS, which is displayed in Grafana, and then fed back into the campaign planner. Grafana is, for clarity, a piece of software for displaying data.

The optimization module as designed in this article would fit into the planner as shown in Figure 2. The optimization module is only to be used for one type of reference generation, during solar energy harvesting, so the planner is still to carry out the majority of attitude planning, as well as the scheduling layer on top of this module.

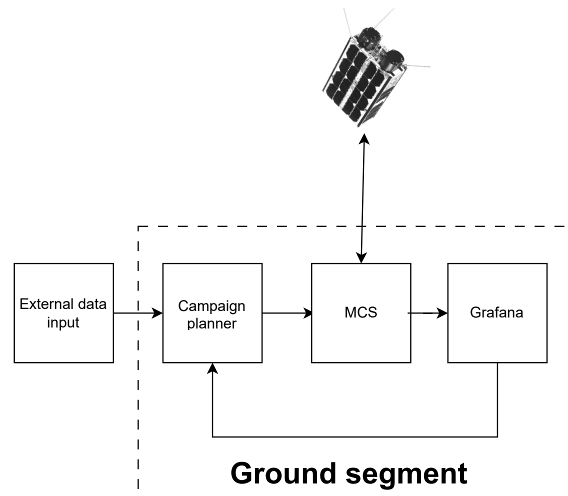

Figure 1: Simplification of the campaign planner in the ground segment and its interaction with the satellite, inspired by Berg et al. (2023).

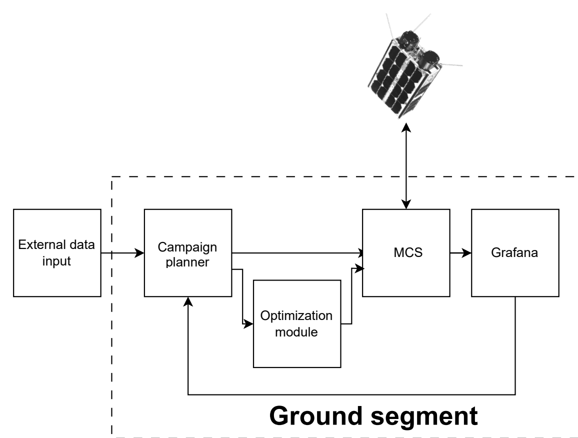

Figure 2: The campaign planner integration of the optimization module.

The tasks that each of the blocks performs when the optimization module is included, in Figure 2, are shown in Table 1.

Table 1: Functionality of each block

Block	**Functionality**
Optimization module	Calculate quaternion references
Campaign planner	Schedule, plan
Mission Control Software (MCS)	Communicate with the satellite Store data
Satellite (on board)	Track quaternion references Collect and send telemetry
Grafana	Store and display telemetry
External data	Produce data for planning

Optimal control problem

As in Kristiansen et al. (2024), the optimal control problem is defined by

$$\min_{\mathbf{x},\mathbf{u}} \quad J_{\mathrm{F}} - \int_0^T P(\mathbf{x},\mathbf{u})\,\mathrm{d}t \tag{1a}$$

$$\text{s.t. } \dot{\mathbf{q}}_b^i = \frac{1}{2}\mathbf{T}(\mathbf{q}_b^i)\boldsymbol{\omega}_{ib}^b + \frac{1}{2}\rho\mathbf{q}_b^i\left(\left(\left(\mathbf{q}_b^i\right)^\mathsf{T}\mathbf{q}_b^i\right)^{-1} - 1\right) \tag{1b}$$

$$\frac{{}^b d}{dt}\boldsymbol{\omega}_{ib}^b = \mathbf{J}_s^{-1}\big(-\mathbf{S}\left(\boldsymbol{\omega}_{ib}^b\right)\left(\mathbf{J}\boldsymbol{\omega}_{ib}^b + \mathbf{A}\mathbf{J}_w\boldsymbol{\omega}_{\mathrm{RW}}^w\right) - \mathbf{A}\boldsymbol{\tau}_{\mathrm{RW}}^w + \boldsymbol{\tau}_{\mathrm{mtq}}^b + \boldsymbol{\tau}_{\mathrm{ext}}^b\big) \tag{1c}$$

$$\frac{{}^w d}{dt}\boldsymbol{\omega}_{\mathrm{RW}}^w = \mathbf{J}_w^{-1}\boldsymbol{\tau}_{\mathrm{RW}}^w - \mathbf{A}^\mathsf{T}\frac{{}^b d}{dt}\boldsymbol{\omega}_{ib}^b \tag{1d}$$

$$(\hat{\mathbf{p}}_i^b)^\mathsf{T}\hat{\mathbf{k}}_i^b \leq \cos\theta_i \quad \forall i \tag{1e}$$

$$\boldsymbol{\tau}_{\mathrm{RW,lb}}^w \leq \boldsymbol{\tau}_{\mathrm{RW}}^w \leq \boldsymbol{\tau}_{\mathrm{RW,ub}}^w \tag{1f}$$

$$\mathbf{x}(0) = \mathbf{x}_0, \tag{1g}$$

where the kinematics and kinetics are in (1b) and (1c), respectively, and (1d) describe the dynamics of the reaction wheels Kristiansen, Grøtte, and Gravdahl (2020). The attitude is represented by the unit quaternion, $\mathbf{q}_b^i = [\eta_b^i,\ (\boldsymbol{\epsilon}_b^i)^\mathsf{T}]^\mathsf{T}$, where η_b^i is its scalar and $\boldsymbol{\epsilon}_b^i$ is its vector part. The $\mathbf{T}(\mathbf{q})$ matrix is given by (Egeland and Gravdahl 2002)

$$\mathbf{T}(\mathbf{q}) = \begin{bmatrix} -\boldsymbol{\epsilon}^\mathsf{T} \\ \eta\mathbf{I}_{3\mathrm{x}3} + \mathbf{S}(\boldsymbol{\epsilon}) \end{bmatrix}, \tag{2}$$

where $\mathbf{I}_{3\mathrm{x}3}$ is the three-dimensional identity matrix and $\mathbf{S}(\cdot)$ is a skew-symmetric matrix equivalent to the vector cross product in three dimensions. The norm of the unit quaternion is not constrained by a normalization constraint other than that the initial attitude should be a unit quaternion due to the necessity to maintain constraint qualifications, but it is rather preserved through the accuracy of the numerical integrator and the Baumgarte stabilization term $\frac{1}{2}\rho\mathbf{q}_b^i\left(\left(\left(\mathbf{q}_b^i\right)^\mathsf{T}\mathbf{q}_b^i\right)^{-1} - 1\right)$, where ρ is a small positive constant (Gros, Zanon, and Diehl 2015). The state variable is defined $\mathbf{x} = \left[\left(\mathbf{q}_b^i\right)^\mathsf{T},\ \left(\boldsymbol{\omega}_{ib}^b\right)^\mathsf{T},\ (\boldsymbol{\omega}_{\mathrm{RW}}^w)^\mathsf{T}\right]^\mathsf{T}$ and the input $\mathbf{u} = \left[(\boldsymbol{\tau}_{\mathrm{RW}}^w)^\mathsf{T},\ \left(\boldsymbol{\tau}_{\mathrm{mtq}}^b\right)^\mathsf{T}\right]^\mathsf{T}$, where $\boldsymbol{\omega}_{ib}^b$ is the angular velocity of the body frame $\{b\}$ about the inertial frame $\{i\}$ represented in the body frame. $\frac{{}^b d}{dt}$ is the derivative with respect to time in the body frame. $\boldsymbol{\omega}_{\mathrm{RW}}^w$ is a vector in a wheel frame $\{w\}$, meaning that it has one element for each of the four reaction wheels, which torque is defined by $\boldsymbol{\tau}_{\mathrm{RW}}^w$. $\boldsymbol{\tau}_{\mathrm{RW}}^w$ is here the only decision variable, constrained by upper and lower bounds, $\boldsymbol{\tau}_{\mathrm{RW,ub}}^w$ and $\boldsymbol{\tau}_{\mathrm{RW,lb}}^w$. The translational motion variables are calculated before the optimal control problem is solved and used as parameters in the calculation of both the magnetorquer torque $\boldsymbol{\tau}_{\mathrm{mtq}}^b$ and the external disturbance torques, $\boldsymbol{\tau}_{\mathrm{ext}}^b$ (see the Appendix for details). The underlying assumption for this choice is that the changes to the translational variables depending on the rotational motion will be negligible over the optimization horizon. Due to the momentum storage capabilities of the reaction wheels, there are three inertia matrices: $\mathbf{J}_s$, $\mathbf{J}_w$, and $\mathbf{J} = \mathbf{J}_s + \mathbf{A}\mathbf{J}_w\mathbf{A}^\mathsf{T}$ represent the inertia of the spacecraft only, the reaction wheels about their axis of rotation, and the total inertia of the complete spacecraft, respectively. The torque distribution matrix $\mathbf{A}$ gives the mapping between the wheel frame and the body frame for the four reaction wheels in the following way,

$$\mathbf{A} = \frac{1}{3}\begin{bmatrix} 1 & 0 & 0 & \sqrt{3} \\ 0 & 1 & 0 & \sqrt{3} \\ 0 & 0 & 1 & \sqrt{3} \end{bmatrix}. \tag{3}$$

The constraint in (1e) was not included in Kristiansen et al. (2024). The purpose of the constraint is to keep one given vector $\hat{\mathbf{p}}_i^b$ away from another vector $\hat{\mathbf{k}}_i^b$ by some angle θ_i for all i such constraints. In this article, one of these constraints is implemented for the star tracker, represented by the directional vector $\hat{\mathbf{p}}_{\text{Star tracker}}^b$ to keep it out of view of the Sun, represented by the sun vector $\hat{\mathbf{s}}^b$. In equation form, this becomes

$$(\hat{\mathbf{p}}_{\text{Star tracker}}^b)^\mathsf{T}\hat{\mathbf{s}}^b \leq \cos\theta. \tag{4}$$

The hat operator (^) indicates that it is a normalized vector. The constraint aims to keep the Sun out of a cone around the star tracker vector, defined by the sun exclusion half-angle θ. This type of constraint has been investigated for attitude problems in Garcia and How (2005) and Wu and Han (2019).

The cost function in (1a), $J_{\mathrm{F}} - \int_0^T P(\mathbf{x},\mathbf{u})\,\mathrm{d}t$, contains two terms: J_{F}, which contains the soft constraints for the attitude and the angular velocity reference, and the net energy term, $E = \int_0^T P(\mathbf{x},\mathbf{u})\,\mathrm{d}t$. The net energy is the integral over time of the net power $P(\mathbf{x},\mathbf{u})$,

$$P(\mathbf{x},\mathbf{u}) = s(\mathbf{x}) - m(\mathbf{x},\mathbf{u}), \tag{5}$$

where $s(\mathbf{x})$ represents the solar power harvested and $m(\mathbf{x},\mathbf{u})$ is the power used for actuation. The solar power function $s(\mathbf{x})$, defined as the sum of the power supplied by the solar panels, is given as

$$s(\mathbf{x}) = \eta_{\mathrm{in}}\kappa(\mathbf{q}_b^i,\mathbf{r}_{\mathrm{Sun}})\delta(\mathbf{r}_{\mathrm{Sun}},\mathbf{r}_{\mathrm{Earth}}), \tag{6}$$

where η_{in} is the efficiency of the batteries, $\kappa(\mathbf{q}_b^i,\mathbf{r}_{\mathrm{Sun}})$ is the instantaneous solar power, and $\delta(\cdot,\cdot)$ is a value between 0 and 1 depending on whether the satellite is in eclipse or not. For this article, the experiment is performed entirely in the Sun, so $\delta(\cdot,\cdot)$ is set to 1. $\mathbf{r}_{\text{Earth}}$ and $\mathbf{r}_{\text{Sun}}$ are vectors containing the distance from the satellite to the Earth's and the Sun's center, respectively.

For a CubeSat such as HYPSO-1, $\kappa(\cdot,\cdot)$ is defined as

$$\kappa(\mathbf{q}_b^i, \mathbf{r}_{\text{Sun}}) = \sum_{j=1}^{n_s} \max\left(\left(\hat{\mathbf{s}}^b\right)^\intercal \hat{\mathbf{n}}_j^b, 0\right)\left(\left(\hat{\mathbf{s}}^b\right)^\intercal \hat{\mathbf{n}}_j^b\right) c_{s,j} A_j, \tag{7}$$

where n_s is the number of solar panels, $\hat{\mathbf{n}}_j^b$ is the normal vector of each solar panel, $c_{s,j}$ is the product of the solar panel efficiency and solar irradiance, and A_j is the area of the solar panel. The max$(\cdot)$ function is implemented as

$$\max\left(x_1, x_2\right) = \frac{1}{2}\left(x_1 + x_2 + \sqrt{\left(x_1 - x_2\right)^2 + \alpha}\right), \tag{8}$$

where α is a small positive constant ensuring a smooth max function.

The actuation power is given by the sum of the power spent on reaction wheels, $P_{\text{RW}}(\cdot)$, and the power spent on magnetorquers, $P_{\text{mtq}}(\cdot)$ (Kristiansen et al. 2024),

$$\begin{aligned} m(\mathbf{x},\mathbf{u}) &= P_{\text{RW}}(\mathbf{x},\mathbf{u}) + P_{\text{mtq}}(\mathbf{u}) \\ &= \frac{1}{\eta_{rw}}\left|\left(\mathbf{A}\boldsymbol{\tau}_{\text{RW}}^w\right)^\intercal \mathbf{A}\boldsymbol{\omega}_{\text{RW}}^w\right| + \frac{\sum_{n=1}^{\infty}\left|m_{\text{mtq},i}^b\right|}{3m_{\text{mtq, ub}}} P_{\text{mtq, max}}, \end{aligned} \tag{9}$$

where η_{rw} is the efficiency of the reaction wheels, $m_{\text{mtq},i}^b$ is the magnetic moment of a magnetorquer on one axis with the upper bound $m_{\text{mtq, ub}}$, and $P_{\text{mtq, max}}$ is the maximum power value of the magnetorquers, where the output efficiency of the batteries is included.

The soft constraint term, J_{F}, in (1a) is chosen to penalize the final states

$$\begin{aligned} J_{\text{F}} &= k_1 J_{\text{path,ref}} + k_2 J_{\text{velocity}} \\ &= k_1\left(1 - \left|\left(\mathbf{q}_b^i\right)^\intercal \mathbf{q}_{\text{ref}}\right|\right) \\ &\quad + k_2\left(\boldsymbol{\omega}_{\text{ref}}^b - \boldsymbol{\omega}_{ib}^b\right)^\intercal\left(\boldsymbol{\omega}_{\text{ref}}^b - \boldsymbol{\omega}_{ib}^b\right), \end{aligned} \tag{10}$$

where k_1 and k_2 are constants, and $J_{\text{path,ref}}$ is a metric on SO(3) (Huynh 2009) that denotes the cost of not reaching the desired attitude. J_{velocity} introduces a cost for not reaching the desired final angular velocity. $\mathbf{q}_{\text{ref}}$ and $\boldsymbol{\omega}_{\text{ref}}^b$ are the reference values for the attitude quaternion and the angular velocity, respectively. $J_{\text{path,ref}}$ is implemented using the same smooth max function (8) as was used in the solar power part of the cost function as $\left|\left(\mathbf{q}_b^i\right)^\intercal \mathbf{q}_{\text{ref}}\right| = \max\left(\left(\mathbf{q}_b^i\right)^\intercal \mathbf{q}_{\text{ref}}, -\left(\mathbf{q}_b^i\right)^\intercal \mathbf{q}_{\text{ref}}\right)$.

Design of the optimization module

A diagram describing the optimization module can be seen in Figure 3. Note from the figure that the module takes in certain information, such as the duration of the optimization horizon, the pre-calculated positions and velocities, and the Sun positions. The data gets downsampled so the data rate is in a reasonable rate for the optimal control problem to solve, and it is upsampled before any data is sent out of the block. The output is a list of quaternions that can be used as references for the onboard attitude tracking system.

The main challenge with the optimal control problem is that it is hard to solve. For this reason, the optimization procedure executes three steps before the solution is uploaded to the satellite:

1. An energy optimal attitude is found and a trajectory is produced using the PD controller by the method described in the appendix. The energy optimal attitude, which is an attitude where the satellite points towards the Sun, is used as the reference for the PD controller. The PD controller trajectory is used as the initial guesses for the optimal control problem.
2. Optimize the net power optimal control problem without the star tracker constraint, (4).
3. Add the star tracker constraint (4) to the optimal control problem. Use the output of the previous optimization as initial guesses to warm-start the solver, then solve the problem.

The motivation for separating points two and three is that the problem is difficult without the constraint keeping the star tracker away from the Sun. Introducing this extra constraint after finding an energy optimal trajectory makes it more likely that the solver will be able to solve the complete problem. The list above does not guarantee that the solver converges, that it finds an optimal solution, or that it finds a solution that is better than the baseline PD controller finds in the first point in the aforementioned list. Therefore, we may reiterate on the list several times using different quaternions, such as the start and end quaternion for the initial guesses, and other options to attempt to get the solver to move into other parts of the state space. A single iteration through the diagram in Figure 3 may take 15-20 minutes on a laptop, indicating how far in advance planning should take place. It is noted that the problem is sensitive to some of the parameters that require tuning, in particular k_1, k_2, ρ, and α.

Experimental setup

The method is being deployed as part of the campaign planner on the 6U CubeSat HYPSO-1. The satellite is currently in low-Earth orbit and is designed for ocean observation. The satellite itself is fully actuated through its four reaction wheels and it has magnetorquers on all three axes. The commands are calculated on ground before the reference trajectory is uplinked to the satellite. The optimization horizon is chosen to be between two separate images, covering the energy harvesting phase of the operations of the satellite. On board the satellite, the references are tracked using a linear controller. The goal is to verify that the maneuver yields the maximum amount of harvested energy, so the sampling rate of the telemetry measurements from the batteries is set higher than normal to capture the information from the maneuver. The battery telemetry sampling rate of the electric power system is set at 0.5 Hz for the experiment. Using a similar argument, the attitude control system is run at 4 Hz,

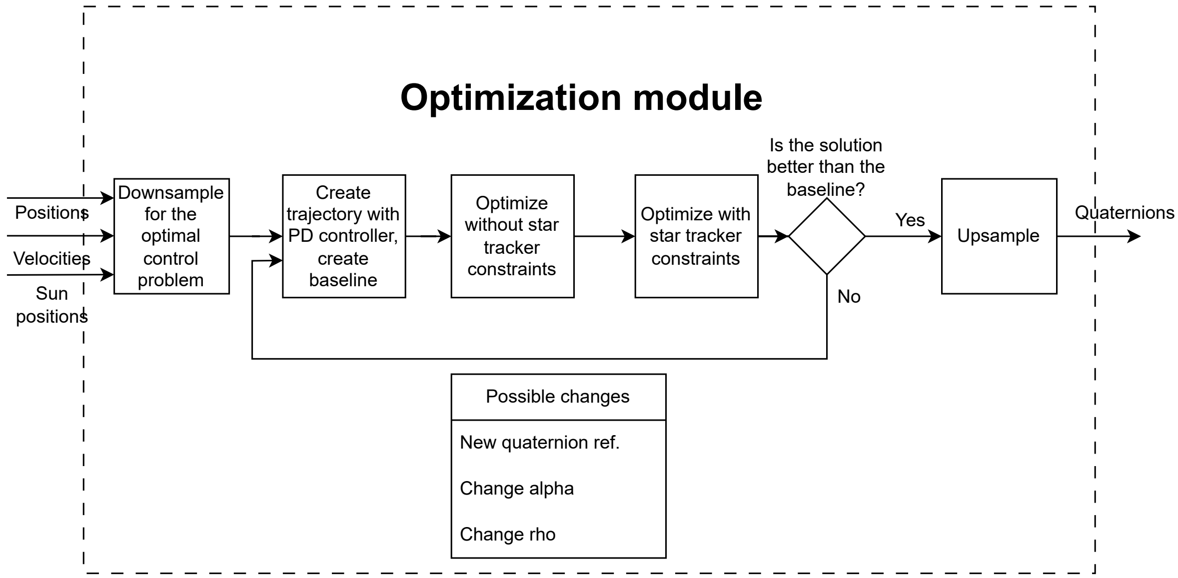

Figure 3: Inner workings of the optimization module. The module is adapted to the non-convex nature of the optimal control problem by creating initial guesses the optimal control problem in two steps.

far exceeding the 0.5 Hz rate that the references are generated at.

The optimization is performed as part of the campaign planner pipeline and thus performed in the ground segment. The optimization is solved using a multiple shooting discretization scheme using IPOPT (Wächter and Biegler 2006) and the automatic differentiation software package CasADi (Andersson et al. 2019). The dynamics are integrated using a fourth order Runge-Kutta method.

The constants used in the optimization problem are shown in Table 2. Upsampling is in the experiment in this article performed using zero-order hold, but there is not reason why interpolating techniques cannot be used instead.

Table 2: Optimization constants

Variable	Value	Unit
Step size (h)	2	s
Solar irradiance	1366	$\mathrm{W/m^2}$
Solar panel efficiency	20	%
c_s	272.2	$\mathrm{W/m^2}$
η_{in}	0.92	-
η_{rw}	0.85	-
α	10^{-5}	-
k_{max}	1	W^2
k_1	$T \cdot 1.75 \cdot 10^2$	W
k_2	$T \cdot 10^4$	$\mathrm{W} \cdot \mathrm{s}^2$
ρ	$1.1 \cdot 10^{-2}$	-
$\boldsymbol{\tau}^b_{\mathrm{RW,lb}}$	$-3.2 \cdot 10^{-3}$	N·m
$\boldsymbol{\tau}^b_{\mathrm{RW,ub}}$	$3.2 \cdot 10^{-3}$	N·m

Results

The data produced as telemetry by the satellite during the maneuver is shown in Figure 4 and Figure 5, while a comparative maneuver is shown in Figure 6. In Figure 4, the quaternion sequence generated by the optimization module, denoted as ”Target quaternion”, is shown relative to the estimated quaternions, the ”knowledge quaternion”. The trajectory from the optimization module starts when the reference starts swaying. The controller does not perfectly track the references in the beginning, as can be seen by comparing response in the beginning of the top two plots. After this phase, the knowledge quaternion plot follows the target quaternion rather well until shortly after 11:50, where there is a jump in the knowledge quaternion. This jump coincides with a spike in the estimated quaternion from the star tracker, which has been non-existent since the beginning of the maneuver. The star tracker immediately stops producing estimates when the maneuver starts and stays that way for the entire optimization horizon apart from two spikes at 11:51 and 11:54. Note that the quaternion is given in the LVLH-frame relative to the body frame instead of the inertial frame relative to the body frame as it is in the optimization problem due to the implementation of the on board system. In addition, note that the satellite rotates significantly during the maneuver. As can be seen in Figure 5, particularly in the solar panels' current plot, the generated power increases at the beginning of the maneuver and stays high the star tracker spike. Comparing this to Figure 6, which covers a maneuver of similar length the pass before, shows that the method implemented in this article produces higher solar current for the time the energy optimal plan manages to keep the solar current at its maximum.

Figure 4: The quaternion reference plotted as the ”target quaternion” and the estimated quaternion shown as the ”knowledge quaternion”. The estimated quaternion from the star tracker is shown in the bottom plot - a flat line indicates that there is no estimate from the star tracker.

Discussion

The significant rotation during the maneuver shown in the quaternion history, Figure 4, is not an issue with respect to getting the highest amount of solar power as the satellite is not constrained in all three axes when pointing a solar panel towards the Sun, and can thus rotate freely around the axis perpendicular to the solar panel if that impedes no other cost in the objective function. As can be seen in Figure 5, the power input goes straight to its maximum value and is kept there, which is the desired behavior.

The star tracker does not produce a quaternion estimate for most of the estimation period, and there are several potential reasons for this that may not be that it is blinded by the Sun. The star tracker on HYPSO-1 can lose track when the angular velocity is too high, for example. Some causes are yet to be discovered. The behaviors that are exhibited at 11:51 and 11:54 are easier to explain: The star tracker has lost track of the attitude, and at 11:51 and 11:54 two single estimates are produced. These estimates are valued higher than the IMU measurements by the onboard attitude estimation system, creating a discontinuous behavior in the knowledge quaternion. Such a spike in the knowledge quaternion results in a large deviation between the desired orientation, the target quaternion, and the estimated orientation, the knowledge quaternion, which in turn leads to the satellite abruptly moving away from an attitude which yields the largest amount of power.

Suboptimal tracking would be expected to lead to a small loss in energy in the transient phases of the maneuvers.

The power produced in the optimal maneuver, Figure 5, is higher than the alternative maneuver Figure 6. This is clearly because the green line that plots the current in panel 1 does not stay at the maximum value, just above 1A, as it does in our optimal solutions. This phenomenon is the topic of ongoing research.

Conclusion

The experimental data show that the method works as the simulation studies in Kristiansen, Gravdahl, and Johansen (2021) and Kristiansen et al. (2024) indicated, even if the tracked maneuver gets somewhat interrupted in the presented results by a sensor issue. Further work includes looking at the inclusion of more complex models for solar irradiance, including ways to incorporate the effect of albedo in the model of the sun vectors, a potentially significant part of the power radiated to the satellite (Vallado 2001). Furthermore, both investigating how well the PD-based tracking controller is tuned relative to the expected optimal performance and looking at alternative tracking controllers can potentially improve the output of the implemented method.

Acknowledgments

This work was funded by the Research Council of Norway through project grant numbers 325961, 337495, and 353062, and Mission Mjøsa.

Appendix

Optimization and initial guesses

Solving non-convex optimization problems requires an informed choice of initial guesses, meaning the starting point for where the solver will look for solutions. This article uses a PD controller to produce the initial guesses for the optimal control problem in (1) by pointing the spacecraft toward an optimal solar power reference, like in Kristiansen et al. (2024). The PD controller formulation we use is given by (Wen and Kreutz-Delgado 1991)

$$\boldsymbol{\tau}^b_{\text{RW}} = \mathbf{K}_p \boldsymbol{\epsilon}_e - \mathbf{K}_d \left(\boldsymbol{\omega}^b_{ib,\text{ref}} - \boldsymbol{\omega}^b_{ib} \right), \tag{11}$$

where $\mathbf{K}_d > 0,\ \mathbf{K}_p > 0$ are constant controller gain matrices, $\boldsymbol{\omega}^b_{ib,\text{ref}}$ is the reference angular velocity. $\boldsymbol{\epsilon}_e$ is the error in the vector part of the quaternion, given as the final three elements of $\mathbf{q}_e = \mathbf{q}^{-1}_{\text{ref}} \otimes \mathbf{q}$, where $\mathbf{q}_{\text{ref}}$ is the reference quaternion and $\otimes$ is the Hamilton product.

The quaternion references for the PD controller are set to the attitude with the maximal incoming solar energy relative to the orientation given as input, then to the final reference quaternion towards the end. The maximal incoming solar power attitude is found using the following problem,

$$\min_{\mathbf{q}} \ -\kappa(\mathbf{q})^2 + k_{\text{max}} J_{\text{path}}(\mathbf{q}) \tag{12a}$$

$$\text{s.t. } \mathbf{q}^\intercal \mathbf{q} = 1, \tag{12b}$$

Figure 5: Power response during the maneuver by maneuver created by the optimization module.

Figure 6: Power response of a similar maneuver by the alternative control scheme on board the satellite in the previous pass.

where k_{max} is a positive constant, and J_{path} is a cost introduced to ensure that only one attitude would be optimal. J_{path} is defined as (Huynh 2009)

$$J_{\text{path}}(\mathbf{q}) = 1 - |\mathbf{q}^\intercal \mathbf{q}_{\text{ref}}|, \tag{13}$$

where $\mathbf{q}_{\text{ref}}$ is the reference attitude quaternion. The absolute value is implemented using the smooth max function (8), as $|\mathbf{q}^\intercal \mathbf{q}_{\text{ref}}| = \max(\mathbf{q}^\intercal \mathbf{q}_{\text{ref}}, \text{-}\mathbf{q}^\intercal \mathbf{q}_{\text{ref}})$. The choice of $\mathbf{q}_{\text{ref}}$ can vary to let the solver explore different trajectories in the space.

Magnetorquers control

The magnetorquers are only used for momentum management of the reaction wheels, with the control algorithm given as (Markley and Crassidis 2014)

$$\boldsymbol{\tau}_{\text{mtq}} = \mathbf{S}(\mathbf{m}^b)\mathbf{B}^b = \mathbf{S}\left(\frac{k_m}{\|\mathbf{B}^b\|_2}\left(\mathbf{S}\left(\mathbf{h}_e^b\right)\mathbf{B}^b\right)\right)\mathbf{B}^b, \tag{14}$$

where $\mathbf{m}^b$ is the magnetic moment produced by the magnetorquers and k_m is a positive constant. $\mathbf{h}_e^b$ is the error in angular momentum for the reaction wheels, given as

$$\mathbf{h}_e^b = \mathbf{A}\mathbf{J}_w(\boldsymbol{\omega}_{\text{RW,ref}}^w - \boldsymbol{\omega}_{\text{RW}}^w), \tag{15}$$

where $\boldsymbol{\omega}_{\text{RW,ref}}^w$ is the angular velocity reference for the reaction wheels.

Perturbing torques

The optimal control problem considers three types of perturbations: Gravity gradient, magnetic torque, and eddy current torque.

The gravity gradient torque is defined as (Hughes 1986)

$$\boldsymbol{\tau}_{\text{grav}}^b = 3\frac{\mu}{\|\mathbf{r}^i\|^3}\mathbf{S}\left(\frac{\mathbf{r}^i}{\|\mathbf{r}^i\|}\right)\mathbf{J}\frac{\mathbf{r}^i}{\|\mathbf{r}^i\|}. \tag{16}$$

where $\mathbf{r}^i$ and $\mathbf{v}^i$ are the position and velocity vectors given in the Earth-centered inertial frame, respectively. μ is the standard gravitational parameter of the Earth.

The magnetic torque is defined as

$$\boldsymbol{\tau}_{\text{mag}}^b = \mathbf{S}(\mathbf{m}_{\text{res}}^b)\mathbf{B}^b, \tag{17}$$

where $\mathbf{B}^b$ is the magnetic field of the Earth represented in body frame coordinates, and $\mathbf{m}_{\text{res}}^b$ is the residual magnetic dipole vector.

The eddy current torque is given as (Hughes 1986)

$$\boldsymbol{\tau}_{\text{eddy}}^b = -k_{\text{eddy}}\mathbf{S}(\mathbf{B}^b)\mathbf{S}(\boldsymbol{\omega}_{ib}^b)\mathbf{B}^b, \tag{18}$$

where k_{eddy} is a positive constant.

The total environmental torque $\boldsymbol{\tau}_{\text{ext}}^b$ is the sum of the previously defined torques. In equation form, this becomes

$$\boldsymbol{\tau}_{\text{ext}}^b = \boldsymbol{\tau}_{\text{grav}}^b + \boldsymbol{\tau}_{\text{mag}}^b + \boldsymbol{\tau}_{\text{eddy}}^b. \tag{19a}$$

References

Andersson, J. A. E.; Gillis, J.; Horn, G.; Rawlings, J. B.; and Diehl, M. 2019. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1): 1–36.

Berg, S.; Bakken, S.; Birkeland, R.; Chiatante, C.; Garrett, J. L.; and Johansen, T. A. 2023. Ground systems software for automatic operation of the HYPSO-2 hyperspectral imaging satellite. In *Sensors, Systems, and Next-Generation Satellites XXVII*, volume 12729, 287–295. SPIE.

Bilimoria, K.; and Wie, B. 1993. Time-optimal three-axis reorientation of rigid spacecraft. *Journal of Guidance Control Dynamics*, 16: 446–452.

Egeland, O.; and Gravdahl, J. T. 2002. *Modeling and simulation for automatic control*. Marine Cybernetics Trondheim, Norway.

Garcia, I.; and How, J. 2005. Trajectory optimization for satellite reconfiguration maneuvers with position and attitude constraints. In *Proceedings of the 2005, American Control Conference, 2005.*, 889–894 vol. 2. ISSN: 2378-5861.

Gros, S.; Zanon, M.; and Diehl, M. 2015. Baumgarte stabilisation over the SO(3) rotation group for control. In *2015 54th IEEE Conference on Decision and Control (CDC)*, 620–625.

Grøtte, M. E.; Birkeland, R.; Honoré-Livermore, E.; Bakken, S.; Garrett, J. L.; Prentice, E. F.; Sigernes, F.; Orlandić, M.; Gravdahl, J. T.; and Johansen, T. A. 2022. Ocean Color Hyperspectral Remote Sensing With High Resolution and Low Latency—The HYPSO-1 CubeSat Mission. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–19.

Hughes, P. C. 1986. *Spacecraft Attitude Dynamics*. New York, NY: John Wiley & Sons. ISBN 978-0-486-14013-1.

Huynh, D. Q. 2009. Metrics for 3D Rotations: Comparison and Analysis. *Journal of Mathematical Imaging and Vision*, 35(2): 155–164.

Karpenko, M.; Halverson, J. K.; Tedla, Y.; Lippman, T.; Han, K.; Ross, I. M.; Barker, M. K.; Robinson, M. S.; Carter, W.; Boyd, A. K.; Patterson, G. W.; O'Shea, C. M.; Turner, F. S.; Poffenbarger, R. T.; Dennehy, C. J.; VanZwieten, T.; and Wolf, A. A. 2024. Fast Attitude Maneuvers for NASA's Lunar Reconnaissance Orbiter: Practical Flight Application of Attitude Guidance using Birkhoff Pseudospectral Theory and Hamiltonian Programming. *IEEE Control Systems Magazine*, 44(2): 26–54.

Kristiansen, B. A.; Gravdahl, J. T.; Gros, S.; and Johansen, T. A. 2024. Energy Optimal Attitude Control and Task Execution for a Solar-Powered Spacecraft. *IEEE Transactions on Control Systems Technology*, 1–14.

Kristiansen, B. A.; Gravdahl, J. T.; and Johansen, T. A. 2021. Energy optimal attitude control for a solar-powered spacecraft. *European Journal of Control*, 62: 192–197.

Kristiansen, B. A.; Grøtte, M. E.; and Gravdahl, J. T. 2020. Quaternion-Based Generalized Super-Twisting Algorithm for Spacecraft Attitude Control. *IFAC-PapersOnLine*, 53(2): 14811–14818.

Markley, F. L.; and Crassidis, J. L. 2014. *Fundamentals of spacecraft attitude determination and control*, volume 33. Springer.

Rigo, C. A.; Seman, L. O.; Camponogara, E.; Morsch Filho, E.; and Bezerra, E. A. 2021. Task scheduling for optimal power management and quality-of-service assurance in CubeSats. *Acta Astronautica*, 179: 550–560.

Schaub, H.; and Lappas, V. J. 2009. Redundant Reaction Wheel Torque Distribution Yielding Instantaneous L2 Power-Optimal Spacecraft Attitude Control. *Journal of Guidance, Control, and Dynamics*, 32(4): 1269–1276.

Vallado, D. A. 2001. *Fundamentals of astrodynamics and applications*, volume 12. Publisher: Springer Science & Business Media.

Wen, J.-Y.; and Kreutz-Delgado, K. 1991. The attitude control problem. *IEEE Transactions on Automatic Control*, 36(10): 1148–1162.

Wu, C.; and Han, X. 2019. Energy-optimal spacecraft attitude maneuver path-planning under complex constraints. *Acta Astronautica*, 157: 415–424.

Wächter, A.; and Biegler, L. T. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1): 25–57.