High-Agile Spacecraft with Flexible Appendages: Limits, Risks, and Mitigation

Andrea Coco
Veoware
Leuven, Belgium
andrea.coco@veowarespace.com

Simon Debois
Veoware
Leuven, Belgium
simon.debois@veowarespace.com

***Abstract*— Modern Earth-observation missions increasingly rely on high-agility retargeting to maximise target acquisitions. At the same time, new-generation spacecraft are adopting larger, lightweight, and lightly damped appendages such as high-power solar arrays and deployable SAR antennas, which introduce low-frequency flexible dynamics and tighter slew-and-settle constraints. The combined trend raises a central question for mission design: how far can agility be increased without incurring residual vibrations that degrade pointing performance and structural margins? This work investigates an agile spacecraft equipped with a Control Moment Gyroscope (CMG) attitude control system and a reduced-order flexible dynamics model capturing the dominant appendage modes. Using an Earth-observation retargeting scenario as a motivating use case, we show that increased agility can improve target acquisition by up to 500% compared with a conventional reaction-wheel cluster. At the same time, this benefit can be limited by residual vibrations that extend manoeuvre settling time, potentially eroding the net advantage of CMG agility and reinforcing the perception that high agility and large flexible appendages are incompatible on the same platform. The first part of the study determines the natural-frequency regime in which residual vibrations become non-negligible following aggressive (bang-bang) slew manoeuvres. The second part evaluates practical mitigation strategies compatible with classical control architectures (e.g., PD), including jerk-limited reference trajectories and input-shaping techniques (e.g., ZVD). Sensitivity to modelling uncertainty and time delays is also assessed. Results indicate that CMG-actuated agile spacecraft can preserve rapid retargeting capability while satisfying pointing and structural requirements, demonstrating that high agility and flexible appendages can be successfully combined within a single satellite.**

Keywords*— *Agile Satellite, Control Moment Gyroscope, Flexible Appendages, Vibration

I. Introduction

Modern satellites are increasingly equipped with power-intensive payloads, such as Synthetic Aperture Radar (SAR) antennas, high data-rate communication systems, and high-resolution optical instruments. To meet these growing power demands, spacecraft are typically designed with large solar arrays. However, the coupling between these large, flexible, and lightly damped appendages and the rigid spacecraft bus can introduce persistent structural oscillations, ultimately degrading pointing performance.

At the same time, agility has become a key requirement for many missions. It is often assumed that large flexible spacecraft are inherently incompatible with agile manoeuvres; however, as highlighted by S. Pellegrino in [1], agility is not primarily limited by structural flexibility concerns, but rather by the available control torque and momentum capacity.

In this context, Control Moment Gyroscopes (CMGs) and Variable-Speed CMGs (VSCMGs) offer an interesting solution, as they can deliver significantly higher torque levels than conventional reaction wheels. While this enables highly agile manoeuvres, the associated high torque authority can potentially also lead to increased excitation of flexible modes in lightly damped structures if not properly managed.

The first goal of this work is to identify how flexible a spacecraft must be in order to be affected by a typical agile slew manoeuvre. A baseline time-optimal bang–bang maneuver is used to evaluate residual oscillations and to understand when flexible dynamics start to degrade performance, potentially reducing the benefit of high-torque actuators. In those cases, mitigation strategies are explored to enable those flexible spacecraft to still take advantage of agile actuation.

Vibration suppression can be addressed in two ways: Avoid exciting flexible modes (e.g., low-bandwidth control, smooth references such as jerk-limited profiles); Allow excitation and then actively suppress the resulting oscillations.

Prior work has explored both directions, including jerk-related vibration effects [2], input shaping techniques [3][5] that consist in the principle of destructive interferences, and active damping with CMGs [6]. In this work, the following methods were investigated: smooth reference tracking, input shaping, and active damping strategies. They are compared in terms of performance and robustness through Monte Carlo analysis under modelling uncertainties. The spacecraft is modelled as a 1-DOF rigid body with one dominant flexible appendage mode (representative of solar-array dynamics). A representative Earth-observation mission analysis is then used to quantify the coverage benefit enabled by these mitigation strategies.

II. Modeling

A simplified model is adopted to isolate and analyze the effects of flexible appendages, such as solar arrays or booms. The spacecraft is modeled as a free-floating, two degree-of-freedom system, capturing the dominant coupling between the rigid-body motion and the flexible dynamics. The corresponding state-space representation is given as follows, derived from [1]:

$$\begin{bmatrix} J_{sat} & 0 \\ 0 & J_f \end{bmatrix}\begin{bmatrix} \ddot{\theta}_{sat} \\ \ddot{\theta}_f \end{bmatrix} + \begin{bmatrix} c & -c \\ -c & c \end{bmatrix}\begin{bmatrix} \dot{\theta}_{sat} \\ \dot{\theta}_f \end{bmatrix} + \begin{bmatrix} k & -k \\ -k & k \end{bmatrix}\begin{bmatrix} \theta_{sat} \\ \theta_f \end{bmatrix} = \\ = \begin{bmatrix} \tau \\ 0 \end{bmatrix}$$

Where J_f is the inertia of the flexible appendage and J_{sat} is the inertia of the spacecraft while θ_{sat} and θ_f denotes the satellite and flexible appendage angle. The parameters c and k are respectively the viscous damping coefficient and the spring stiffness. To isolate the perturbation induced by

structural flexibility, the coupled equations of motion can be decoupled, allowing the flexible dynamics to be expressed as a disturbance acting on the rigid-body motion. In this framework, the spacecraft attitude is decomposed as $\theta_{sat} = \theta_{sat,r} + \theta_{sat,f}$, separating the rigid-body double-integrator response from the flexible second-order contribution. Accordingly, the input torque is written as $u' = \frac{\tau}{J_{sat}+J_f}$

$$\ddot{\theta}_{sat,f} + 2\left(1+\frac{J_f}{J_{sat}}\right)\zeta\omega_n\dot{\theta}_{sat,f} + \left(1+\frac{J_f}{J_{sat}}\right)\omega_n^2\theta_{sat,f} = \frac{J_f}{J_{sat}}\,u'$$

With

$$\omega_n = \sqrt{\frac{k}{J_f}} \text{ and } \zeta = \frac{c}{2\sqrt{kJ_f}}$$

This decomposition is useful because it separates rigid tracking from flexible excitation, allowing direct analysis of residual vibration induced by different command profiles.

From the previous equation, the free-free natural frequency can be retrieved $\omega_{FF} = \omega_n\sqrt{1+\frac{J_f}{J_{sat}}}$ and damping ratio $\zeta_{FF} = \zeta\sqrt{1+\frac{J_f}{J_{sat}}}$.

Solar panels are modelled and represent the source of the disturbance due to flexible modes. The size of the panels was computed following standard system engineering design formulas [4], providing at End Of Life 450 W. The inertia of the panels were modelled as

$$J_f = 2\left(\frac{1}{12}\,m_{sp}l_{sp}^2 + m_{sp}l_{sat}^2\right)$$

Where 2 Solar panels are considered and the transport theorem is applied to account the mounting of the panel at the side of the satellite.

A. Control Moment Gyroscopes

A Control Moment Gyroscope is a momentum exchange device that acts differently than reaction wheels, it exploits the gyroscopic torque obtained by tilting the angular momentum axis rather than by varying its magnitude.

The current CMGS in production and in development at veoware are reported in Table 1,

Table 1 Veoware Control Moment Gyroscope

	MicroCMG	***MiniCMG***	***MediumCMG***
Angular Momentum (Nms)	0.7	2.8	6
Max Torque (Nm)	1.1	4.4	9
Mass (kg)	2.4	3.2	4
Power (W)	15	20	30

A cluster of four CMGs is typically used due to its ability to provide full three-axis controllability while offering a degree of redundancy. Although the achievable angular momentum and torque capabilities vary slightly depending on the cluster configuration, the Veoware CMGs are designed to support spacecraft in the 50–500 kg class, enabling angular rates up to 5 deg/s and 5 deg/s^2

.

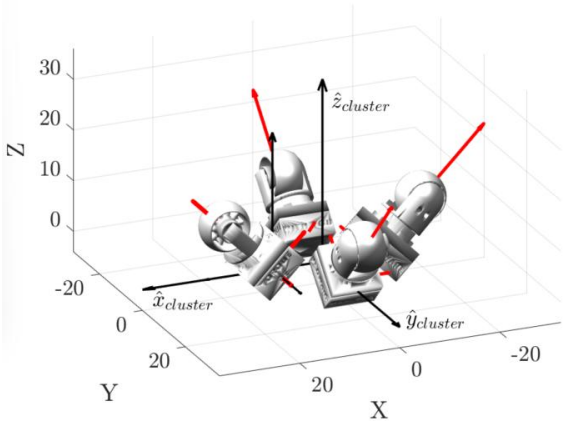

Figure 1 Pyramid Cluster of CMGs

III. Strategies

The focus of this paper is to analyze strategies for mitigating, or ideally cancelling, residual oscillations in flexible spacecraft maneuvers. In general, two approaches are considered. The first is to minimize excitation from the beginning, by avoiding control inputs with sharp steps or spikes that strongly excite flexible modes. The second is to intentionally allow excitation and then apply properly timed commands to cancel the induced vibration.

The first approach is typically more conservative: it often requires smoother and therefore slower profiles than bang-bang time-optimal control. Its advantage is lower structural stress, but it does not guarantee zero residual vibration. The second approach can preserve high agility and shorter maneuver times, but it generally induces higher structural loads; its performance also depends strongly on model accuracy, since is effective only if the flexible dynamics are identified correctly.

A. Avoid Modes Excitation

It is well known that the optimal time slew maneuver is a bang-bang profile. Although the acceleration step, might excite structural natural modes. There exist different ways to smooth this profile by using versine or jerk limited references.

Smoothing the bang–bang profile effectively consists in shaping the acceleration profile by acting on its derivative. This introduces jerk as an explicit design parameter in the maneuver definition, resulting in a continuous and more structurally compatible trajectory of the following form:

$$\theta(t) = \dddot{\theta}\,t^3 + \ddot{\theta}t^2 + \dot{\theta}t + \theta_0$$

The latter equation represent the slew angle parametrization, where the acceleration is no more instantaneous but modelled up to its first derivative.

Another way consists of using a versine, which belongs to the C^∞ class of functions and ensures continuity of all derivatives, effectively extending the zero-derivative condition at the boundaries to all orders.

Both methods appear feasible and technically relevant. However, the versine profile is slower than the jerk-based maneuver for the same maximum jerk value. In this analysis,

the two profiles are instead constrained to have the same duration; as a result, the versine profile requires a higher peak jerk, as can be seen in Figure 2. A comparative analysis will be presented to quantify the resulting differences in residual oscillations.

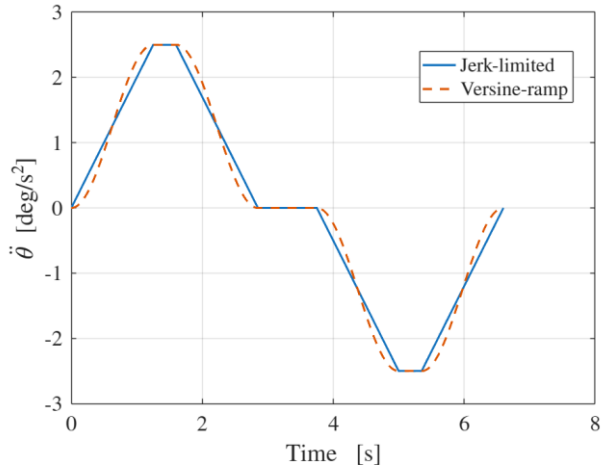

Figure 2 Jerk Limited versus Versine

The acceleration ramps of the jerk limited manoeuvre are replaced by versine profile with the following law:

$$versin(\theta) = 1 - \cos(\theta)$$

The constant jerk acceleration ramp profile, will follow the following rule:

$$\ddot{\theta}(t) = \frac{1}{2}\ddot{\theta}_{max} versin\left(\frac{\pi t}{t_j}\right)$$

Where t_j is the jerk time. It is important to point that the versine is designed to last the same time as the jerk limited profile. For this reason, the peak jerk reached by the versine manoeuvre is higher than the constant one used in the previous method.

B. Excite and Cancel Oscillations

Input shaping techniques are widely used to suppress residual oscillations in flexible systems. The key idea is to filter a command by convolving it with a sequence of delayed impulses. The impulse amplitudes and delays are chosen so that the vibration generated by each impulse interferes destructively, while the rigid-body objective is preserved.

For a natural mode with natural frequency ω, damping ratio ζ, and damped frequency $\omega_d = \omega\sqrt{1-\zeta^2}$, the normalized residual vibration can be written as

$$V(\omega,\zeta) = e^{\zeta\omega t_n}\sqrt{C(\omega,\zeta)^2 + S(\omega,\zeta)^2}$$

With

$$C(\omega,\zeta) = \sum_{i=1}^{n} A_i\, e^{\zeta\omega t_i}\cos(\omega_d t_i)$$

$$S(\omega,\zeta) = \sum_{i=1}^{n} A_i\, e^{\zeta\omega t_i}\sin(\omega_d t_i)$$

Here, A_i and t_i are impulse amplitudes and times, and n is the number of impulses. To obtain a zero-vibration shaper, impose $V = 0$, which requires $C = 0$ and $S = 0$, together with practical constraints such as $\sum_i A_i = 1$ and $A_i > 0$.

The Zero Vibration (ZV) shaper uses two impulses of amplitude A_1 and A_2 to be applied at t_1 and t_2:

$$A_{1_{ZV}} = \frac{1}{1+K},\quad A_{2_{ZV}} = \frac{K}{1+K} \quad \text{where } K = e^{-\frac{\zeta\pi}{\sqrt{1-\zeta^2}}}$$

$$t_1 = 0,\quad t_2 = \frac{\pi}{\omega_d},\quad \omega_d = \omega_{FF}\sqrt{1-\zeta^2}$$

The Zero Vibration Derivative (ZVD) shaper adds one impulse to improve robustness to modeling and timing errors:

$$A_{1_{ZVD}} = \frac{1}{(1+K)^2},\quad A_{2_{ZVD}} = \frac{2K}{(1+K)^2},\quad A_{3_{ZVD}} = \frac{K^2}{(1+K)^2}$$

$$\text{where } K = e^{-\frac{\zeta\pi}{\sqrt{1-\zeta^2}}},\ \omega_d = \omega_{FF}\sqrt{1-\zeta^2}$$

$$t_1 = 0,\quad t_2 = \frac{\pi}{\omega_d},\quad t_3 = \frac{2\pi}{\omega_d}$$

In practice, ZV is shorter and faster, while ZVD is generally more robust when modal uncertainty is non-negligible. The full derivation of the formulas presented above is presented in [3] .

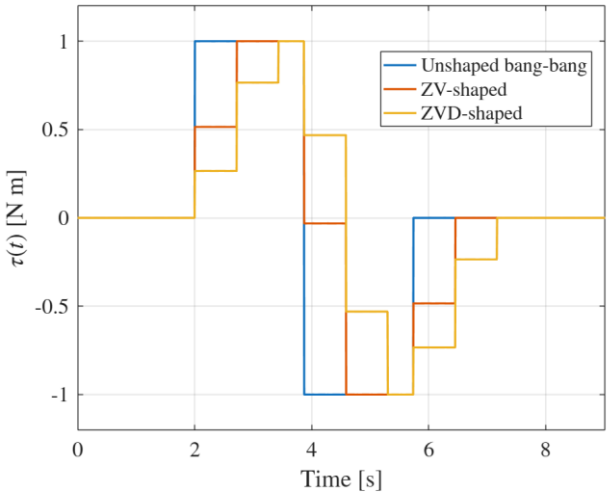

Figure 3 Comparison ZV and ZVD

Figure 3 compares ZV and ZVD. Both methods reduce residual vibration, but ZVD is generally more robust to modeling uncertainty (e.g., frequency or damping mismatch) and timing errors, at the cost of a longer shaping duration. This technique can be applied both in open-loop command shaping and in closed-loop architectures, where shaping is combined with feedback control [4].

Figure 4 presents an analysis that includes modeling uncertainties in filter design. In particular, uncertainty in the natural frequency is introduced, and this makes the difference between the two methods clear: both ZV and ZVD are effective under ideal tuning, but once frequency mismatch is present, ZVD is noticeably more robust. Thus, although both methods work well at the nominal point, ZVD provides

superior tolerance to natural-frequency uncertainty, at the cost of a typically longer shaping sequence.

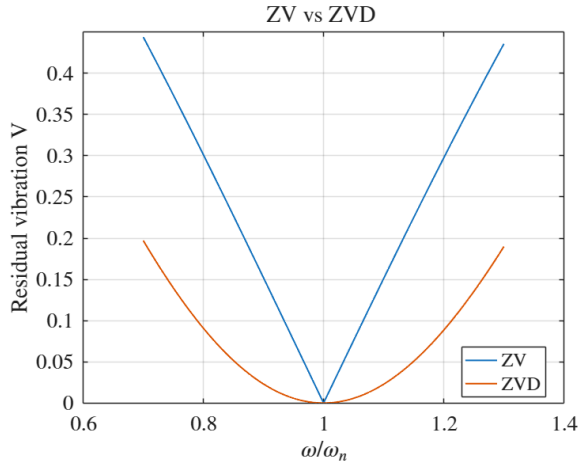

Figure 4 Robustness to Natural Frequency Modeling Errors

IV. Results

A. Residual Oscillation Characterization

The first objective is to identify which spacecraft classes are more affected by agile manoeuvres when flexible appendages are present. The analysis uses the 1-DOF flexible free-floating satellite model introduced in Section II and compares two time-optimal bang-bang profiles: an agile case and a non-agile case. The agile profile is characterized by a maximum acceleration of 4 deg/s^2 , while the slow profile is limited to 0.2 deg/s^2 and both profiles span an angle of 30 deg. For spacecraft in the 100-300 kg class, such high-acceleration manoeuvres are typically achievable only with CMGs, VSCMGs, or thrusters.

For this comparison, a spacecraft with inertia $J_{sat} = 130$ kgm^2 is considered, with two attached solar panels. The panel sizing is based on a required electrical power of 450 W, and the corresponding appendage inertia is computed accordingly [II] and is equal to 6 kgm^2.The total maneuver time comprehends also the settling time, that is defined as the time after which the relative oscillation remains within a 1 mdeg bound and lastly the damping ratio ζ is equal to 0.02.

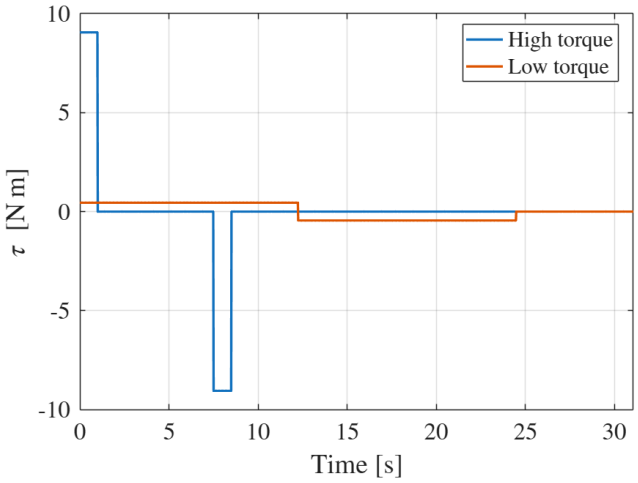

Figure 5 Optimal Time Bang-bang Profiles

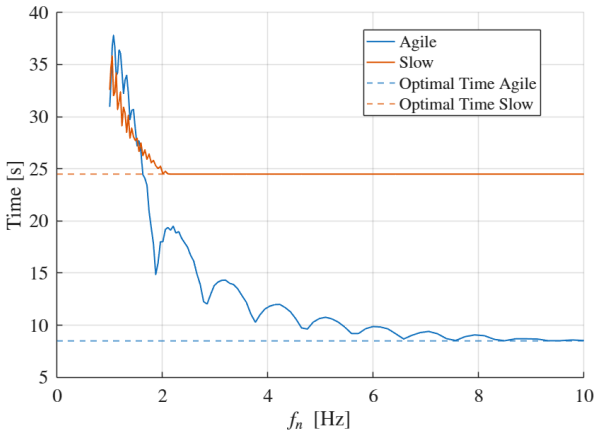

Figure 6 Optimal Bang–Bang Maneuver Comparison: Agile vs. Slow

Figure 6 illustrates that a more flexible spacecraft exhibits significantly higher residual vibrations when subjected to the same command profile as a stiffer one. For a spacecraft with a dominant flexible mode below 2 Hz, a bang–bang maneuver excites the structure to such an extent that it largely offsets the benefits of an agile actuator cluster. From 2 Hz to up to 5 Hz, performance degradation due to flexible excitation is still evident, although agile maneuvers remain beneficial overall. For higher natural frequencies, the induced residual oscillations become negligible.

An interesting feature of the Figure 6 is the presence of multiple local minima. This behavior is linked to command timing in two ways. First, the finite acceleration pulse duration t_{acc} acts as a spectral filter on the command input, producing repeated peaks and notches with characteristic spacing on the order of $f \approx k/t_{acc}$, When a structural mode aligns with a notch, the injected energy is reduced, leading to lower excitation. Second, the time separation Δ between acceleration and deceleration phases introduces a phase shift $2\pi f\Delta$, which at specific frequencies results in destructive interference and further attenuation of residual vibrations (pulse-separation effect).

In the following sections, a representative case with a natural frequency of 2 Hz is considered, and the different mitigation strategies are evaluated.

B. Reference smoothing

1) Jerk Limited / Versine

The jerk-limited and versine approaches were tested for a spacecraft with first natural frequency of 2 Hz, taken from the analysis conducted in the previous paragraph. The slew angle is 30 deg and the maximum acceleration and angular velocity are respectively 4 deg/s^2 and 4 deg/s.

For a proper implementation of the jerk-limited maneuver, the key design parameter is the selection of the maximum allowable jerk. It must be low enough to limit the excitation of flexible modes, yet high enough to avoid an excessive increase in maneuver time.

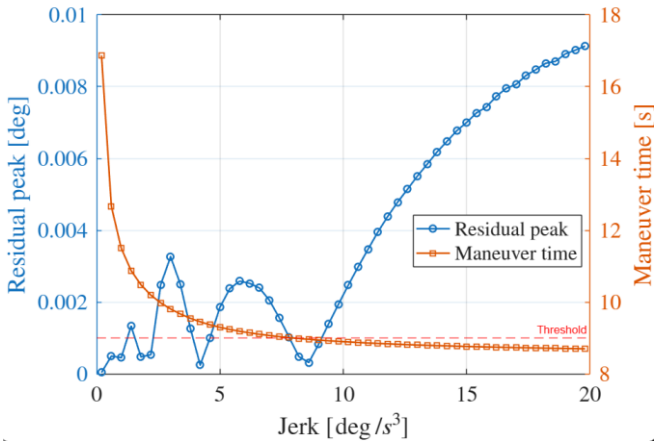

Figure 7 Effects of Maximum Jerk on Total Maneuvering Time

Figure 7 shows a non-monotonic trend of the jerk with respect to the induced residual oscillations. While the maneuver time decreases with increasing maximum jerk, as the input profile approaches the time-optimal bang–bang solution, the same behavior is not observed for the residual peak. This is due to the different torque input profiles generated, where specific values of the acceleration duration correspond to multiples of half the natural period, leading to destructive interference and a reduction in the oscillation amplitude.

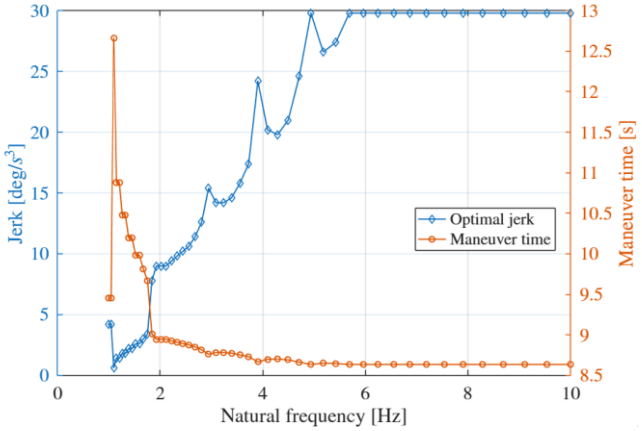

Figure 8 Optimal Jerk for Jerk Limited Across Frequencies

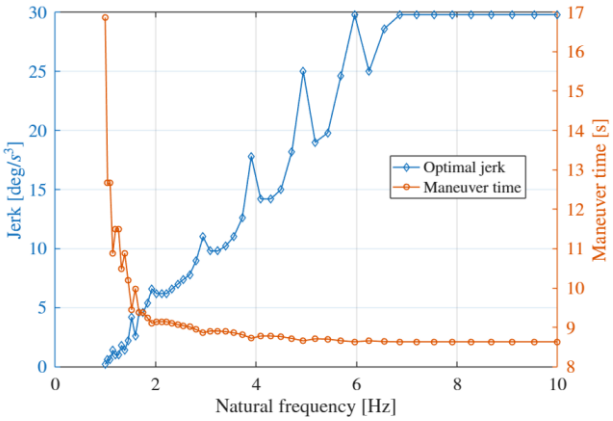

Figure 9 Optimal Jerk for Versine Across Frequencies

Figure 8 and Figure 9 generalize the analysis presented in Figure 7 by identifying, for each case, the jerk value that minimizes the maneuver time while keeping the induced residual oscillation below 1 mdeg For natural frequencies below 5 Hz, the optimal jerk tends to increase with frequency, as lower values are needed to limit the excitation of flexible modes. Above this threshold, the jerk reaches the imposed limit of 30 deg/s³, and the maneuver progressively resembles a time-optimal bang–bang profile without residual oscillations.

2) Input Shaping

The ZV and ZVD input shapers exploit precise pulse timing to achieve complete cancellation of residual oscillations. Figure 10 presents a comparison where the shaped input profile from Figure 3 are applied and compared over the unshaped bang-bang input profile. The reference maneuver taken as example is a 30 deg slew maneuvers with 2 Hz first natural mode.

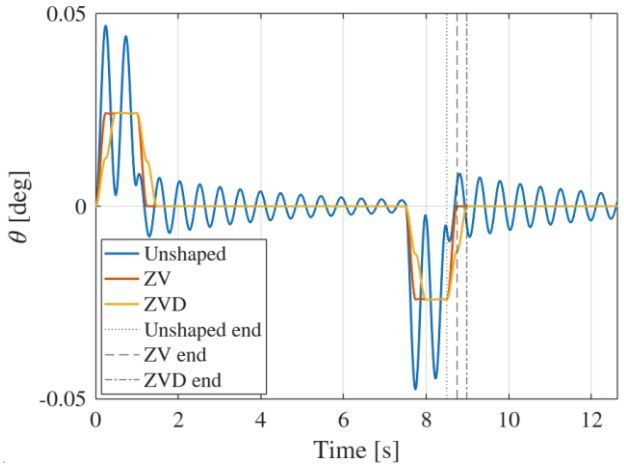

Figure 10 Comparison ZV and ZVD

The three-pulse decomposition of the ZVD results in a slower overall maneuver compared to both ZV and bang–bang profiles. Both methods, when designed using the exact natural frequency and with pulses applied at the precise timing, can theoretically achieve complete cancellation of residual oscillations. However, it must be noted that time delays, as well as the sampling rate at which commands are sent to the actuators and uncertainties about the natural frequency, can lead to pulses being applied at slightly incorrect times, thereby degrading the effectiveness of the methods.

In Table 2 are summarized the results for the 30 deg slew maneuver with 2 Hz as dominant flexible mode.

Table 2 Comparison Results

Comparison Results		
Methods	***Total Time (s)***	***Settling time (s)***
Bang Bang	16.89	8.39
Jerk Limited	8.94	0
Versine	9.14	0
ZV	8.74	0
ZVD	8.99	0

C. Methods Comparison

Figure 11 and Figure 12 report the adopted solutions, compared against the bang–bang reference maneuver. A significant improvement in performance can be observed at lower frequencies. The jerk-limited and versine reference trajectories exhibit the longest maneuver times in the 1–3 Hz range. The ZV method is the fastest among the considered approaches at lower frequencies, approaching the optimal bang–bang maneuver time highlighted by the dashed yellow line.

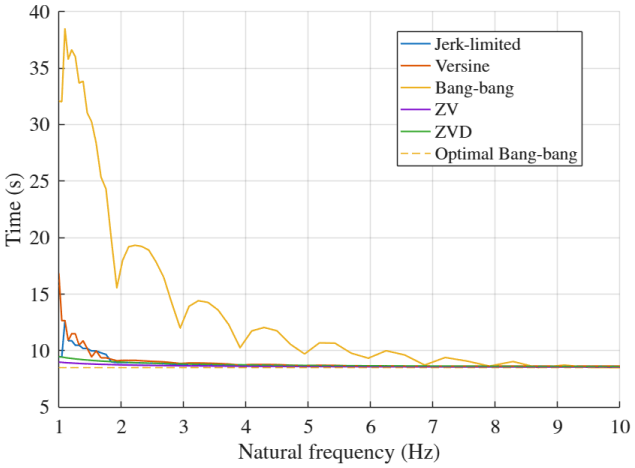

Figure 11 Mitigation Strategies versus Bang-Bang

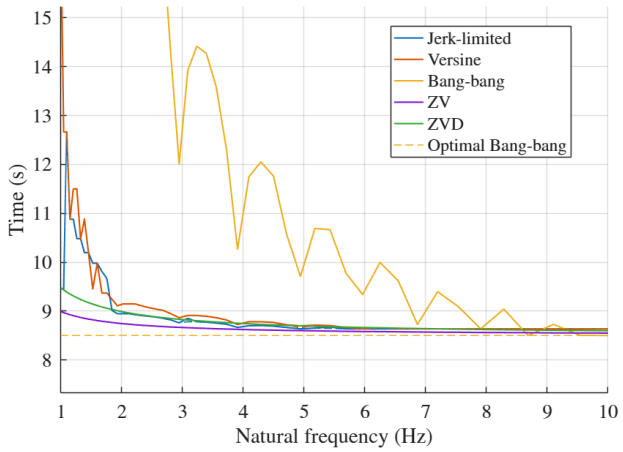

Figure 12 Zoom on Vibration Suppression Methods

As the system becomes progressively more rigid, the benefits of these methods decrease. For higher natural frequencies (e.g., 10 Hz), the bang–bang maneuver becomes the fastest solution, as residual oscillations are no longer induced, and the total maneuver time converges to the optimal reference time.

D. Earth Observation Scenario

Using the VEO Mission Analysis Simulator [8], an earth observation mission scheduling task is simulated. The versine method was implemented in the guidance generation, replacing the optimal bang bang guidance. The target consist of a strip over Tokyo, and the satellite is placed on a 500km altitude Sun-Synchronous orbit. The satellite's inertia is 130 kgm^2 and two attitude control architectures are considered: an agile cluster, made of Mini CMG allows torque up to 10 Nm, while the non-agile made of 4 reaction wheels each with 4 Nms momentum storage and 0.1 Nm torque capability.. The targets must be acquired using a pushbroom techniques, and are located as illustrated in Figure 13, where the area of interest is discretized into acquisition strips.

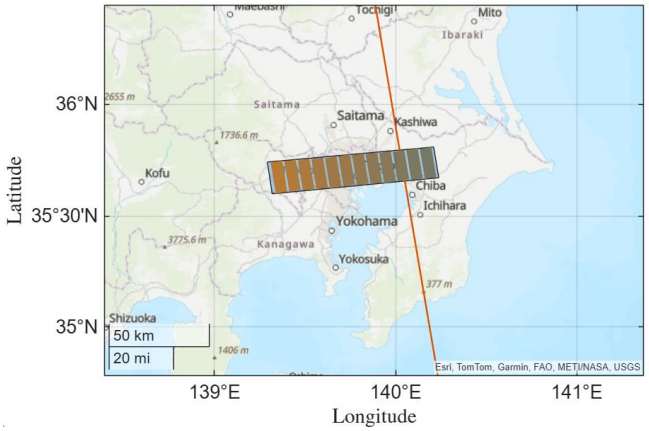

Figure 13 Target Discretization

The versine reference maneuver, the optimal bang–bang profile, and a solution based on a reaction wheel cluster are compared. The simulation is performed using a rigid-body spacecraft model. The versine profile is tuned assuming a dominant natural frequency of 2 Hz, consistent with the previous analysis. In the table below there are reported the results obtained with the three different architectures.

Table 3 Results Earth Observation Use Case

Methods	***n. of Targets collected***
Agile Bang Bang	11
Agile Versine Reference	9
RW Bang Bang	2

Even though the versine maneuver is tuned for a 2 Hz natural frequency, the slew angles required in this use case are relatively small. As a result, the associated time penalty is significantly reduced compared to what is observed in Figure 12 leading to only two fewer targets being collected.

.

V. Conclusion and Future Works

This work presented a comparative assessment of different strategies to mitigate residual oscillations that can significantly extend the duration of slew maneuvers. The proposed methods are compatible with existing control laws and should be regarded as a first-level approach to reducing structural excitation. Despite their simplicity, they prove to be highly effective, enabling high-agility Earth Observation missions to be conducted with reduced settling times and improved efficiency.

Future work will focus on extending the analysis to a higher-fidelity spacecraft model, including multiple flexible modes and more representative structural dynamics. In addition, the study will be expanded to consider clusters of Control Moment Gyroscopes and to evaluate performance

within a closed-loop simulation framework, allowing for a more comprehensive assessment of guidance and control interactions.

REFERENCES

[1] Marshall, M. A. and Pellegrino, S., “Slew Maneuver Constraints for Agile Flexible Spacecraft”, Journal of Guidance Control Dynamics, vol. 46, no. 12, pp. 2300–2314, 2023. doi:10.2514/1.G007430.

[2] Steiner, W. (2021), On the impact of the jerk on structural dynamics. Proc. Appl. Math. Mech., 21: e202100202. https://doi.org/10.1002/pamm.202100202

[3] Goubej, Martin & Skarda, Radek & Schlegel, Milos. (2009). INPUT SHAPING FILTERS FOR THE CONTROL OF ELECTRICAL DRIVE WITH FLEXIBLE LOAD.

[4] Wertz, J. R., Everett, D. F., and Puschell, J. J. (Eds.). (2011). *Space Mission Analysis and Design* (4th ed.). Microcosm Press and Springer.

[5] Dian, W., Yunhua, W., Hongyi, X., Franco, B.-Z., and Chen, X., “Scheduled input shaping based attitude agile control for flexible spacecraft with vibration suppression”, Aerospace Science and Technology, vol. 164, Art. no. 110355, Elsevier, 2025. doi:10.1016/j.ast.2025.110355.

[6] Shi, Jian-Feng & Damaren, Christopher. (2005). Control Law for Active Structural Damping Using a Control Moment Gyro. Journal of Guidance Control and Dynamics - J GUID CONTROL DYNAM. 28. 550-553. 10.2514/1.11269

[7] Kimand, Jae-Jun & Agrawal, Brij. (2006). Experiments on Jerk-Limited Slew Maneuvers of a Flexible Spacecraft. Collection of Technical Papers - AIAA Guidance, Navigation, and Control Conference 2006. 2. 10.2514/6.2006-6187.

[8] Coco, A. and Debois,S., (2025), “Next-Generation Earth Observation: Agile Satellites With CMG Cluster for Enhanced Coverage”