POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Model Order Reduction for Structural Microvibration Analysis: Methods and Comparative Assessment

Laurea Magistrale in Space Engineering - Ingegneria Spaziale

Author: Francesco Monai

Advisor: Prof. Lorenzo Dozio

Co-advisors: Leonardo Peri , Tommaso Robbiani

Academic year: 2024-2025

1. Introduction

The accurate quantification of spacecraft pointing errors requires mapping internal disturbance forces to payload line-of-sight motion through the structural transfer function of an n-DoF second-order system. In high-fidelity finite element (FE) models, the large number of DoF makes frequency response evaluation computationally demanding, as it requires solving large linear systems at multiple frequency points across the target bandwidth. This motivates the adoption of efficient model order reduction (MOR) techniques to ensure practical applicability in microvibration analysis. The conventional Craig-Bampton (CB) methodology is widely used for reducing large-scale structural models but shows limitations when high spectral accuracy is required, since accurate frequency response function (FRF) reconstruction over the target bandwidth requires retaining a large set of fixed-interface modes, reducing compression efficiency. Subsequent developments have introduced two-stage frameworks that combine Component Mode Synthesis (CMS) methods with state-space MOR algorithms, aiming to achieve further dimensionality compression [3]. The primary limitation

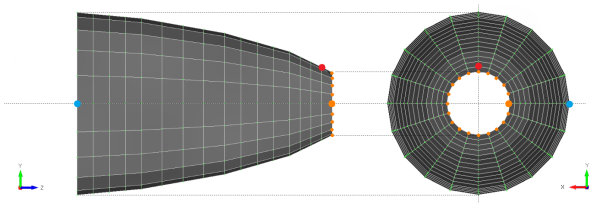

Figure 1: Nozzle FEM model

of state-space methods lies in the first-order linearization, which inherently doubles the system dimensionality ($2n$) and destroys the original second-order physical structure ($\mathbf{M}, \mathbf{C}, \mathbf{K}$). The present work addresses this limitation by developing structure-preserving MOR strategies capable of achieving high spectral fidelity while maintaining computational efficiency within the frequency band relevant to Line-of-Sight jitter analysis. Three MOR paradigms are considered: Balanced Truncation (BT), Moment Matching (MM), and a hybrid procedure that combines the classical CB method with modern second-order reduction techniques. All the investigated methodologies are formulated as projective methods, where the ROMs are generated by constructing appropriate left and right projection matrices to map the full-scale dynamics onto a lower-dimensional subspace. The assessment follows a two-stage strategy where the 960-DoF nozzle model (Fig. 1) serves as a benchmark to evaluate second-order MM and BT against their first-order counterparts. The best-performing configurations identified in the benchmark phase are subsequently applied to the 12U CubeSat system (Fig. 2) to rigorously validate scalability and accuracy in real-world microvibration scenarios through the proposed hybrid reduction framework.

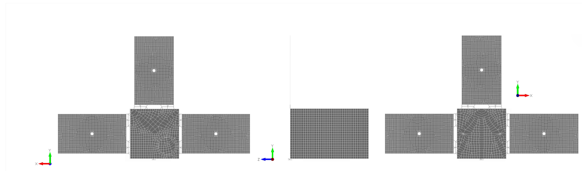

Figure 2: 12U Cubesat FEM model

2. BT Methods

BT is an HSV-based MOR technique that preserves dominant input–output dynamics by ranking system states according to their Hankel Singular Values (HSVs), which quantify joint controllability and observability. The computation of the controllability and observability Gramians $(\mathbf{P}, \mathbf{Q})$ through Lyapunov equations represents the main computational bottleneck but it enables the identification of dynamically significant states that contribute most to the global input-output energy.

First-Order Reduction Conventional first-order reduction utilizes the Cholesky factors of the system Gramians ($\mathbf{P} = \mathbf{S}\mathbf{S}^T, \mathbf{Q} = \mathbf{R}\mathbf{R}^T$) and Singular Value Decomposition (SVD) to compute the projection matrices required from the system's balanced realization. This procedure preserves the states with the highest input-output contribution, discarding those with negligible HSVs as shown in Eq. (1):

$$\mathbf{S}\mathbf{R}^T = [\mathbf{U}_1 \quad \mathbf{U}_2] \begin{bmatrix} \mathbf{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_2 \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^T \\ \mathbf{V}_2^T \end{bmatrix} \tag{1}$$

Although this framework ensures stability and provides an *a priori* $\mathcal{H}_\infty$ error bound, its inherent transition to a state-space representation destroys the second-order physical structure and the physical interpretability of the coordinates.

Second-Order Reduction To preserve the inherent structure of the dynamical system, second-order balancing techniques decompose the system Gramians into four submatrices corresponding to position and velocity Gramians [1], as shown in Eq. (2):

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_p & \mathbf{P}_{pv} \\ \mathbf{P}_{vp} & \mathbf{P}_v \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \mathbf{Q}_p & \mathbf{Q}_{pv} \\ \mathbf{Q}_{vp} & \mathbf{Q}_v \end{bmatrix}. \tag{2}$$

Since the underlying Lyapunov equations remain unchanged, these matrices are identical to the Gramians derived for first-order methods. However, this second-order framework enables eight distinct reduction strategies [1] by selectively substituting the full state-space Gramians with specific partitioned submatrices within the standard balancing workflow. This flexibility introduces a limitation: in non-colocated systems, the resulting ROM may exhibit unstable behavior requiring *a posteriori* verification.
In addition to the previously mentioned strategies, this study evaluates two supplementary approaches: **ps** and **vps**. These are specifically designed to guarantee stability [4], regardless of the approximation order or system configuration, by leveraging specific balancing properties that enforce coincidence between test and trial subspaces.
Since both formulations share the same Lyapunov equation, the second-order framework ensures structural preservation without introducing additional computational overhead.

Application to Nozzle Case study The numerical consistency and approximation accuracy are first assessed using the nozzle benchmark. For the second-order BT, three physically equivalent configurations are analyzed, specifically considering position, velocity, and acceleration as outputs. The FRFs are reconstructed *a posteriori* by exploiting Laplace-domain differentiation properties to ensure dynamic consistency among displacement, velocity, and acceleration outputs, with the evaluation in terms of acceleration, as required by the output specifications. The position-based approach demonstrates superior accuracy across the entire frequency spectrum, establishing it as the most reliable configuration for the final model. While the velocity-based formulation also yields competitive results, the acceleration-based alternative fails to provide a valid approximation at low frequencies, regardless of the balancing method or approximation order employed. Consequently, only the position-based results are presented in Table 1.

Table 1: Balanced Truncation summary table. Asterisks ($*$) denote properties observed numerically during validation but not theoretically guaranteed a priori.

Type	Structure	Stability	Realness	RMS err at 1.25%	RMS err at 5%
First-Order BT	✗	✓	✓	0.12	0.043
Second-Order BT					
p	✓	✓*	✓	0.03	0.01
v	✓	✓*	✓	0.03	0.02
pv	✓	✓*	✓	0.04	0.02
vp	✓	✗	✓	-	-
SP					
ps	✓	✓	✓	0.08	0.03
vps	✓	✓	✓	0.085	0.08

The proposed second-order methodologies generally outperform traditional first-order methods in terms of approximation fidelity. The position-based approach (**p**) yields lower RMS errors across the investigated frequency range, making it the most suitable configuration within the considered benchmark. While stability-preserving variants perform slightly worse than standard methods at low reduction orders, their accuracy improves as the model order increases. The **ps** scheme exhibits monotonic accuracy enhancement, whereas the **vps** variant provides only marginal improvements. Figure 3 shows the energy ratio defined as the ratio between the sum of HSVs of the ROM and the FOM, as a function of the approximation order. This dimensionless metric quantifies the amount of *energy* retained by the ROM and provides a direct measure of approximation fidelity.

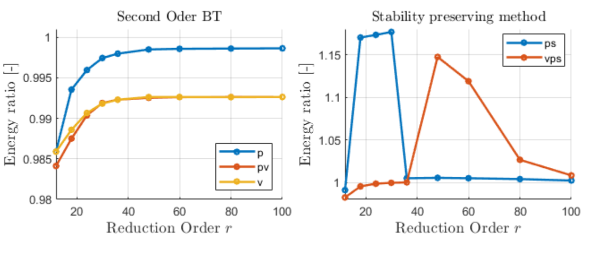

Figure 3: Energy ratio second-order BT methods

The figure on the left illustrates the standard BT methods for second-order systems. The **vp** case is excluded due to its unstable behavior across all approximation orders. Conversely, the **p**, **v**, and **pv** formulations confirm that the system energy is concentrated within a limited number of states, as the energy ratio rapidly approaches unity even at low approximation orders. The stability-preserving methods, reported on the right, offer good energy ratio for very low-order approximations ($r = 12$); however, different trends emerge as the order increases: The **vps** formulation exhibits numerical overshoots at higher reduction orders, whereas the **ps** variant displays two localized peaks at lower approximation orders before asymptotically stabilizing to unity. These results suggest that the one-side projection provides a less accurate FRF reconstruction than the full projection approach leading to energy amplification.

3. MM Methods

The methodology employs a moment-matching framework for FRF approximation, centered on interpolating the transfer function's derivatives (moments) at predefined expansion points [5]. The subspace dimension is thus defined by the aggregate number of moments matched across all selected expansion points. For a ROM of dimension r, the method involves a fundamental trade-off: introducing additional frequency shifts improves global fidelity across the spectral domain, whereas increasing the number of matched moments enhances local approximation accuracy around each shift. For a multi-point expansion with m shifts and q moments per shift, the reduced dimension is $r = m \times q$.

First-Order Reduction A fundamental relationship exists between Krylov subspaces and the moment-matching property of the FRF [2]. For a given system, the q-th order Krylov subspace is defined as:

$$\mathcal{K}_q(\mathbf{S}(s_0), \mathbf{v}(s_0)) = \text{span}\{\mathbf{v}, \mathbf{S}\mathbf{v}, \mathbf{S}^2\mathbf{v}, \dots, \mathbf{S}^{q-1}\mathbf{v}\} \tag{3}$$

where $\mathbf{S}(s_0)$ and $\mathbf{v}(s_0)$ are matrices derived from the first-order realization of the system, and $s_0 = j2\pi f_0$, with f_0 denoting the interpolation center frequency. Krylov-based projection ensures that a ROM spanning the subspace in Eq. (3) matches q moments per expansion point under a Galerkin framework. The $2q$ moment-matching capability of two-sided projection frameworks is not exploited to avoid potential stability issues. At this point, multi-point matching is achieved by constructing the projection matrix through the concatenation of basis vectors spanning distinct Krylov subspaces. However, this first-order transition inherently requires complex-valued matrices and doubles the computational DoFs ($2n$), motivating the move toward second-order structure-preserving implementations.

Second-Order Reduction This methodology employs second-order Krylov techniques to generate projection bases directly from the system's matrices, bypassing state-space linearization. By adopting a one-sided Galerkin framework, the first q moments are matched while preserving the second-order physical structure and symmetry. However, non-zero frequency shifts s_0 inherently yield complex-valued bases, introducing the same memory and computational overhead found in first-order formulations.

To eliminate the complex-valued overhead, a novel *realification* procedure is implemented for multi-point expansion. Valid only for the case of single-MM at each expansion point, this method constructs the Krylov subspace from the undamped system $\mathbf{C} = \mathbf{0}$; consequently, the generating operators remain real-valued, ensuring a real projection matrix. This choice avoids the complex-valued basis components typically introduced by the damping term in second-order Arnoldi or Lanczos iterations. The reduced damping matrix is subsequently computed consistently with the original damping model (e.g., proportional or modal damping), preserving physical coherence with the FOM physical formulation.

Another key aspect of MM methods is the selection of expansion points. While standard techniques rely on linear or logarithmic spacing, this study implements a specialized approach inspired by the Iterative Rational Krylov Algorithm (IRKA) [2]. This approach achieves optimal point placement within the target bandwidth by iteratively selecting the shifts as the mirror images ($-\lambda_i$) of the ROM eigenvalues. To ensure convergence within the target bandwidth for LoS jitter analysis, a spectral filtering criterion is applied: updated shifts falling outside the designated frequency range are discarded and redistributed via interpolation of the remaining shifts. This constraint prevents shift drift outside the target bandwidth and promotes high-fidelity spectral reconstruction in the frequency region most relevant to disturbance propagation.

Nozzle Case study To assess the performance of MM techniques, three scenarios are evaluated by matching $q = \{1, 2, 3\}$ moments per expansion point s_0, resulting in MM1 MM2 and MM3 methods. Numerical evidence indicates that increasing the moments q beyond 2 for a fixed ROM dimension r provides diminishing returns in accuracy. Considering the relatively wide bandwidth required for jitter analysis, increasing the number of matched moments per expansion point tends to reduce the effectiveness of local high-order interpolation, making it more difficult to achieve uniformly accurate approximation across the entire frequency range. Table 2 summarizes the results, comparing equally spaced expansion points for both first- and second-order methods against the final two cases obtained via the Pseudo-IRKA approach. The Logarithmically spaced expansion points approach is discarded due to its poor performance compared to the linear spaced method.

Table 2: Summary table for Moment Matching methods. Asterisks (∗) denote results that exhibit sensitivity to the specific first-order state-space realization employed during the reduction process

Type	Structure	Stability	Realness	RMS err at 1.25%	RMS err at 5%
First-Order MM					
MM1	✗	✓*	✗	0.25	$2.0\ 10^{-4}$
MM2	✗	✓*	✗	0.59	$1.9\ 10^{-3}$
MM3	✗	✓*	✗	0.32	$1.6\ 10^{-2}$
Second-Order MM					
MM1	✓	✓	✓	0.32	$3.7\ 10^{-5}$
MM2	✓	✓	✗	0.18	$1.3\ 10^{-5}$
MM3	✓	✓	✗	0.48	$3.7\ 10^{-3}$
Pseudo IRKA MM					
MM1	✓	✓	✓	0.21	$1.3\ 10^{-5}$
MM2	✓	✓	✗	0.28	$2\ 10^{-4}$

Table 2 demonstrates that increasing the approximation ratio r/n from 1.25% to 5% yields significant fidelity improvements for all the considered approaches. However, RMS error primarily captures global approximation performance and may underestimate localized resonance discrepancies. This limitation is critical for lower-order models, where the ROM response diverges significantly from the FOM between expansion points s_0. At higher approximation orders (e.g., 5% of the FOM size), the divergence in fidelity between MM1, MM2, and MM3 becomes pronounced for both first- and second-order methodologies and the optimal variant differs by case. Despite the marginal gains of MM2, MM1 represents a favorable compromise between matrix real-valuedness, numerical robustness, and approximation accuracy. Regarding expansion point selection, the Pseudo-IRKA algorithm yields marginal RMS error improvements while incurring significant iterative overhead, making equally spaced shifts a more practical choice. Overall, second-order MM consistently outperforms first-order methods, providing a superior trade-off between spectral fidelity and computational efficiency.

The *realification* procedure maintains the spectral accuracy of the MM1 method, yielding RMS errors comparable to those of the MM2 formulation. This consistency validates that constructing the projection basis V from the undamped system ($\mathbf{C} = \mathbf{0}_{n\times n}$) does not introduce approximation bias or artificial dissipation.

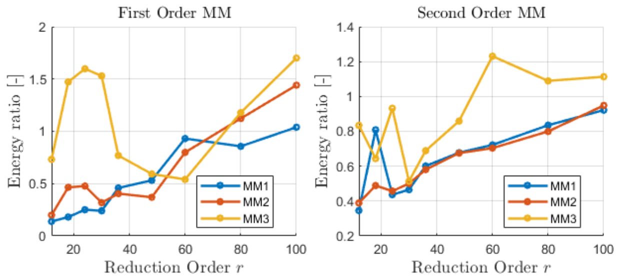

Figure 4: Energy ratio second-order MM methods

Energy ratio trends in Figure 4 confirm that second-order MM methods effectively approximate the system dynamics, with energy ratios asymptotically approaching unity for the primary two formulations. Conversely, first-order approximations exhibit non-physical energy overestimation as the number of matched moments increases; notably, only the MM1 configuration converges to unity. This divergence is attributed to the inherent interpolative nature of the MM paradigm, which lacks an explicit objective function for energy preservation.

4. Cubesat Case Study

Direct BT on large-scale models is computationally demanding due to the cost of solving Lyapunov equations. This necessitates a hybrid MOR framework to enable a baseline comparison with MM. An initial CB step reduces the FOM to a 2200-DOF intermediate model (8% size), preserving 6 critical I/O boundary nodes to maintain interface fidelity. This intermediate system is further reduced to 50 DoFs using **p**-method and stability-preserving second-order BT variants to assess their performances.

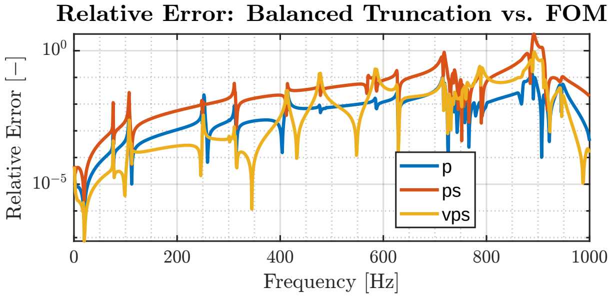

Figure 5: Relative error of CB + second-order BT.

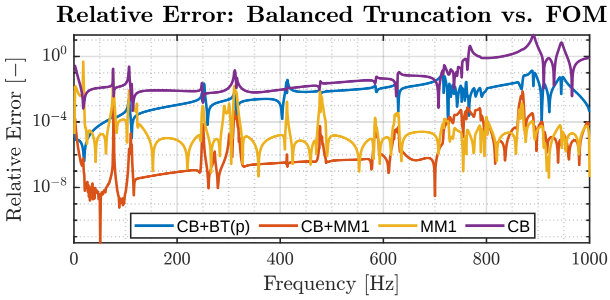

Figure 6: Relative error of different MOR. $r = 50$.

Figure 5 illustrates the relative errors for the three cases. While **vps** stability-preserving method outperforms the **p**-method at low frequencies, it exhibits accuracy degradation at higher frequencies. **ps**, on the other hand, exhibit poor approximation over the entire band. The performance of MM is evaluated through two distinct strategies: a direct reduction of the FOM to order 50 and a hybrid approach utilizing the 2200-DOF intermediate system. Consistent with the nozzle benchmarking phase, this study exclusively employs the single-point (MM1) method. Figure 6 compares the reductions of the 27000-DoF FOM to a 50 DoF system, via four strategies: (1) hybrid CB-BT, (2) hybrid CB-MM, (3) direct MM (4) direct CB. This benchmark evaluates the impact of the intermediate CMS step on spectral fidelity, compared to single-step methods. All evaluated strategies outperform standalone CB at equivalent orders, with the hybrid CB-MM1 configuration achieving the highest spectral fidelity. The error floor at 100–150 Hz constitutes an intrinsic lower bound inherited from the primary CB stage, confirming that ROM precision is limited by the intermediate CMS fidelity rather than the secondary reduction. While BT improves the CB baseline, it remains less effective than MM1, which ensures superior spectral alignment via targeted interpolation. Conversely, the sharp accuracy degradation of standalone CB beyond 700 Hz reflects modal truncation effects at high frequencies. MM methods represent a favorable compromise between computational efficiency and spectral approximation accuracy.

5. Conclusions

This study evaluates two prominent MOR frameworks applied to structural dynamics problem: energy-based BT and interpolation-based MM. These methodologies are first assessed through a medium-scale nozzle benchmark and subsequently extended to a large-scale 12U CubeSat assembly. The comparative assessment confirms that second-order MOR techniques, particularly when integrated into a hybrid framework, provide the necessary computational efficiency to support high-fidelity spacecraft pointing analysis. By significantly reducing the dimensionality of the structural models while preserving high spectral accuracy in the microvibration frequency range, the proposed methodologies enable the rapid evaluation of pointing budget requirements and jitter performance. The investigation demonstrates that reduction accuracy is highly sensitive to the choice of output variables.

In BT, a substantial performance gap exists between position- and acceleration-based formulations. While BT is highly effective for position-level outputs, it fails to maintain fidelity for direct acceleration outputs. The most robust results are achieved by reducing the system based on position and reconstructing the acceleration in the Laplace domain, ensuring dynamic consistency with output requirements. The superior accuracy of the **p**-method contrasts with its lack of an *a priori* stability guarantee while the **ps** and **vps** formulations ensure stability by construction at the cost of lower spectral fidelity in this specific case study. On the other hand, MM provides superior spectral accuracy and computational efficiency as the approximation order increases. The challenge of complex-valued matrices, typical of multi-point Krylov methods, is successfully mitigated through a novel single-moment *realification* approach. The CubeSat study confirms that MM maintains high-fidelity results across the entire frequency spectrum, whereas BT-based models exhibit accuracy degradation in high-frequency regimes.

Ultimately, the hybrid MOR framework emerges as the optimal solution for high-dimensional spacecraft systems. By integrating CB reduction with second-order algorithms, the framework provides a modular, representative model that significantly reduces computational overhead compared to the FOM while maintaining the structural properties. This approach allows for rapid, targeted secondary reductions to meet shifting mission requirements without necessitating a return to the computationally expensive full-scale FE environment.

References

[1] P. Benner and S. W. Werner. Frequency- and time-limited balanced truncation for large-scale second-order systems. *Linear Algebra and its Applications*, 623:68–103, Aug. 2021.

[2] S. Gugercin, A. C. Antoulas, and C. Beattie. H_2 Model Reduction for Large-Scale Linear Dynamical Systems. *SIAM Journal on Matrix Analysis and Applications*, 30(2):609–638, Jan. 2008.

[3] L. Peri, M. Pagano, L. Dozio, and P. Nali. A Virtual Testing Simulation Environment for the ESA's Micro Vibrations Measurement System. *Aerotecnica Missili & Spazio*, 104(1):3–13, Feb. 2025.

[4] T. Robbiani, F. Sanfedino, and L. Dozio. Second-order stability-preserving balanced reduction methods of structural models. *Proceedings of CEASAIDAA 2025*, 2025.

[5] T.-J. Su and R. R. Craig. Model reduction and control of flexible structures using Krylov vectors. *Journal of Guidance, Control, and Dynamics*, 14(2):260–267, Mar. 1991.