Department of Precision and Microsystems Engineering

Over-Actuation and Over-Sensing for Active Damping of Higher-Order Modes in a Flexure-Based Motion Stage

Kiavash Khalili

Report no	:	2026.054
Coach	:	A.M. (Aditya) Natu
Professor	:	Dr. S.H. (Hassan) Hossein Nia Kani
Specialisation	:	Mechatronic System Design
Type of report	:	Master Thesis
Date	:	11 June 2026

TU Delft Delft University of Technology

Challenge the future

Contents

1 Introduction

1.1. Background and Relevance

High-precision motion systems are used in applications where small positioning errors can directly affect performance, such as semiconductor manufacturing, microscopy, optical inspection, and precision metrology. In these systems, both accuracy and speed are important. Higher motion speed usually requires higher accelerations and higher control bandwidth. However, this also makes the system more sensitive to structural vibrations. If flexible modes are excited during operation, they can increase tracking error, cause residual oscillations, and increase settling time[1], [2], [3].

Traditionally, the dynamic performance of precision mechanisms has been improved by increasing the structural stiffness and pushing the first natural frequency to a higher value. This is effective to a certain extent, because a higher first resonance allows a higher usable bandwidth. However, increasing stiffness alone is not always sufficient. It can increase mass, require larger actuation forces, and shift the flexible-mode problem to higher frequencies instead of removing it[3], [4].

The use of active vibration control as a way to enhance the modal behavior of precision systems represents a major improvement in dynamic performance. The primary difference between traditional approaches (i.e., passive damping) and this approach are the use of sensors to measure structural motion; and the application of actuator forces to control structures. Flexure-based nanopositioning systems represent one area where lightly-damped structural resonance modes significantly degrade positioning accuracy and active vibration control could play a role in tackling this issue[4], [5].

Over-actuation and over-sensing extend this idea by adding more actuators and sensors than the minimum required for the main motion degree of freedom. The additional sensing channels can provide information about flexible modes that are weakly visible in the main performance output. At the same time, the additional actuators can be used to suppress these modes more directly. Therefore, over-actuation and over-sensing provide a possible way to improve disturbance rejection and residual vibration behaviour without relying only on increased mechanical stiffness[5], [6], [7].

1.2. Problem Definition

The motion stage that was analyzed within the scope of this thesis is a flexure-based stage. This stage has as its primary function to move along a single axis but it also contains lightly damped higher order modes that could be problematic. Vibrations can occur as a result of rapid motions, environmental perturbations, or through cross-coupling of motion at various parts of the mechanism. If vibrations occurs when sufficient damping is not present for those flexible modes, then both the positioning accuracy and the residual vibrational amplitude will be adversely affected.

The primary problem lies with the fact that most of the important modes do not show up directly in the normal control channel. This means that a controller using just the main displacement signal will likely have trouble seeing those other high order modes sufficiently to dampen them. A second problem is that a given physical sensor and actuator combination might not be able to isolate a particular mode. Therefore, multiple modes could easily appear in the same frequency response channel; this complicates simple active vibration control and leads to problems associated with spillover, phase problmes, or non-minimum phase behavior.

Therefore, the problem addressed in this thesis is not only how to design a controller. The problem is also how the mechanism, actuator layout, sensor layout, and measured outputs can be designed so that the important modes become easier to observe, isolate, and control.

1.3. LITERATURE GAP

Previous work has shown that over-actuation and over-sensing can improve the dynamic behaviour of precision systems by improving modal observability and controllability. This can then enable better active disturbance rejection [5], [6], [7]. However, the effectiveness of active vibration control depends strongly on actuator and sensor placement, because the targeted modes must be both well-excited and well-measured [8], [9].

Centralized Multiple-input and multiple-output (MIMO) control enables multiple modes to be controlled simultaneously; however, it is generally necessary for an accurate high order model of the system, and results in controller configurations which are often more complex to design and apply than decentralized approaches. Mechanical methods (i.e., stiffening structures or eliminating excitation of undesirable mode shapes by smart actuator placement) may also help to mitigate vibration problems but will typically not provide active suppression of unwanted vibrations if they are still excited by external disturbances.

A gap exists in practice that separates purely mechanical design from complex centralized MIMO (Multiple-Input Multiple Output) control. Mechanism design with a focus on control, the placement of actuators and sensors, and modal filtering will address the gap identified above. The goal is to establish mode-dominant channels where a single target mode is predominant in all pairs of inputs and outputs. With these channel established, modular resonance controllers like PPF and band-pass PPF may then be applied to suppress the undesirable resonances in a relatively simple control system.

1.4. RESEARCH OBJECTIVE

The objective of this thesis is to investigate if the vibration behaviour of a flexure-based motion stage can be improved by using over-actuation and over-sensing together with modal filtering and resonant feedback control.

More specifically, a control-oriented stage architecture is designed and evaluated. In this architecture, the mechanism geometry, actuator and sensor placement is optimized, and virtual modal outputs are chosen such that the targeted flexible modes can be observed, isolated, and damped more effectively.

The focus of this thesis is not to design a complex tracking controller. Instead, this thesis focuses on improving the damping of selected flexible modes and reducing the vibration contribution caused by disturbances around these modes.

1.5. RESEARCH QUESTIONS

The main research question of this thesis is:

How can over-actuation and over-sensing be combined with modal filtering and resonant feedback control to reduce vibration caused by targeted flexible modes in a flexure-based motion stage?

This question is answered through the following subquestions:

1. **How can the mechanism geometry be designed and optimized so that the selected flexible modes become easier to separate and control?**
   This is addressed in Chapter 2 and Chapter 3.

2. **How can actuator and sensor placement be used to obtain useful physical channels, and where does single-point placement become insufficient?**
   This is addressed in Chapter 4 and motivates the modal-filtering approach in Chapter 5.

3. **To what extent can modal filtering of redundant sensor signals be used to construct mode-dominant virtual outputs for the targeted modes?**
   This is addressed in Chapter 5.

4. **How effective are PPF and band-pass PPF controllers in attenuating the targeted modal channels, and how much spillover is introduced?**
   This is addressed in Chapter 6 using the 7×7 MIMO frequency-response model. The selected band-pass PPF controllers are then verified on the reduced 5×5 state-space model in Chapter 7.

5. **Under which disturbance conditions and sensor-noise levels does the over-actuated and over-sensed architecture provide a measurable reduction in vibration?**
   This is addressed in Chapter 8 and Chapter 9.

1.6. METHODOLOGY OVERVIEW

The methodology starts with the design of the flexure-based motion stage. First, the design requirements are defined based on the desired degree of freedom, available piezo patches, required stiffness, first natural frequency, and sensing and actuation authority.

Based on these requirements, the leaf-spring dimensions and mechanism layout are selected.

After the initial mechanism design, the geometry is optimized in COMSOL. The optimization is not only used to increase the first natural frequency, but also to improve the modal behaviour of the stage. For this reason, parasitic pitch and roll behaviour are pushed higher in terms of frequency, and the selected rigid-body mode is separated as much as possible. In this way, the mechanism is designed with the later vibration-control problem in mind.

Next, the actuator and sensor placement problem is investigated. The placement method is first tested on the main degree of freedom, where the expected sensor location can be estimated more easily. After this, the same approach is applied to the selected rigid-body mode. The placement results are checked using COMSOL-based criteria and MATLAB-based controllability, observability, and Hankel singular value (HSV) analysis.

After the physical placement study, it is shown that single-point sensing and actuation are not sufficient for all modes. Some higher-order modes have similar spatial deformation patterns, which means that one physical sensor can measure a combination of several modes. This results in mixed channels and makes simple active vibration control more difficult. Therefore, modal filtering is introduced.

In the modal filtering step, multiple sensor signals are combined into virtual outputs. These virtual outputs are constructed so that one targeted mode becomes dominant, while selected nuisance modes are suppressed. In this way, the control problem is moved from physical channels with mixed modal content to more mode-dominant channels.

The mode-dominant channels are then used for the frequency-domain control design. First, the full 7×7 MIMO frequency-response plant is generated from the FEM model. PPF and band-pass PPF controllers are then applied to the selected channels. The controllers are evaluated based on resonance attenuation, sensitivity, and spillover.

To perform time-domain simulations, a state-space model is required, because the simulations cannot be performed directly on the FRD data. A reduced 5×5 state-space model is identified from the frequency-response data. The estimated model is compared with the original FRD data, and the band-pass PPF controllers are implemented on this model. The attenuation, spillover, and sensitivity are checked again before the model is used for the time-domain benchmark.

Finally, the controlled and uncontrolled cases are compared in the time domain. The benchmark includes tracking motion, floor-vibration disturbance, and sensor noise in the over-sensing channels. Both an ideal over-sensing case and a realistic over-sensing case are considered. This allows the theoretical benefit of the proposed architecture to be separated from the practical limitation caused by sensor-noise propagation.

1.7. THESIS CONTRIBUTIONS

The main contributions of this thesis are:

1. A control-oriented design workflow was made for a flexure-based motion stage. The mechanism was not only designed for high stiffness and a higher first natural frequency, but also with the later vibration-control problem in mind. For this reason, parasitic pitch and roll behaviour were reduced, and the selected yaw-mode was separated as much as possible.

2. An actuator and sensor placement approach was used based on modal strain criteria in COMSOL. The placement results were also checked in MATLAB using controllability, observability gramians, and Hankel singular value (HSV) analysis.

3. It was shown that single-point sensing and actuation are not always sufficient when several modes have similar spatial deformation patterns. Even when the actuator and sensor positions were optimized, some physical channels still contained more than one dominant mode.

4. A modal filtering approach was used to solve this problem. Multiple sensor signals were combined into virtual outputs, so that the targeted modes could be isolated more clearly.

5. PPF and band-pass PPF controllers were applied to the mode-dominant channels and evaluated in the frequency domain. The controllers were compared based on resonance attenuation, sensitivity, and spillover.

6. A reduced 5×5 state-space model was made from the frequency-response data so that the controller performance could also be tested in the time domain.

7. A time-domain benchmark was performed to compare the uncontrolled case, the ideal over-sensing case, and the realistic over-sensing case. The comparison was made under tracking motion and modal disturbance excitation.

8. The practical limitations of the proposed architecture were discussed. In particular, it was shown that the benefit of over-actuation and over-sensing depends on how well the targeted modes can be isolated and how clean the additional sensing channels are in terms of sensor noise.

1.8. THESIS OUTLINE

The remainder of this thesis is structured as follows.

Chapter 2 presents the mechanism design. The design requirements, piezo patch constraints, leaf-spring design, and selected mechanism layout are discussed.

Chapter 3 presents the control-oriented structural optimization. In this chapter, the geometry of the mechanism is optimized not only to increase the first natural frequency, but also to improve the modal behaviour of the stage. The focus is placed on reducing parasitic pitch and roll behaviour and separating the selected rigid-body mode as much as possible.

Chapter 4 discusses the actuator and sensor placement problem. The placement method is first tested on the main degree of freedom and is then applied to the selected rigid-body mode. The results are checked using both COMSOL-based criteria and MATLAB-based controllability, observability gramians, and Hankel singular value (HSV) analysis.

Chapter 5 explains the transition from physical channels to modal channels. It is shown that single-point sensing and actuation are not sufficient for all modes, because several modes have similar spatial deformation patterns. For this reason, modal filtering is introduced, where multiple sensor signals are combined to construct mode-dominant virtual outputs.

Chapter 6 presents the frequency-domain control design. The 7×7 MIMO plant is generated, and PPF and band-pass PPF controllers are applied to the selected modal channels. The controllers are evaluated based on resonance attenuation, sensitivity, and spillover.

Chapter 7 presents the reduced 5×5 state-space benchmark model. This model is identified from the frequency-response data because time-domain simulations cannot be performed directly on the FRD data. The estimated state-space model is first compared with the original FRD data. After this, the band-pass PPF controllers are implemented on the state-space model, and the attenuation, spillover, and sensitivity are checked before the model is used for the time-domain benchmark.

Chapter 8 presents the time-domain benchmark. The simulation architecture, tracking trajectory, floor-vibration disturbance, sensor-noise model, and simulation cases are explained. The uncontrolled case, ideal over-sensing case, and realistic over-sensing case are compared to evaluate the practical benefit of the proposed architecture.

Chapter 9 discusses the main findings of the thesis. The benefit of over-actuation and over-sensing is discussed together with the importance of mode isolation, the influence of sensor noise, and the practical limitations of the method.

Finally, Chapter 10 presents the conclusions and future work. The main findings are summarized in relation to the research objective, and recommendations are given for future development and experimental validation.

2 Mechanism Design

2.1. Design Requirements

To start with, a couple of design requirements were set, namely a high first natural frequency in order to make the mechanism more relevant. Also, the lateral stiffness should be high enough to have compliance mainly in the direction of the degree of freedom. The mechanism's stiffness should be no more than 80 percent of what the actuator can push. Materials were also limited, so the design had to be made according to those. To get a higher natural frequency, the mass of the stage and the leaf springs should be as low as possible. Thus, for the stage, Aluminium 6061-T6 was selected, and for the leaf spring, AISI Type 304 Stainless Steel was selected. These were the criteria on which the materials were selected:

- Material selection
  - Better Young's modulus (E) per density
  - Better fatigue characteristics
  - AISI Type 304 Stainless Steel for the leaf springs
  - Aluminium 6061-T6; 6061-T651 for the stage and intermediate pieces

Separated mode frequencies were another requirement that was set. The purpose behind this requirement was to be able to design for control so that, by oversensing, the vibration control channel would isolate the modes in order to simplify the control problem. It was also important to have a symmetrical design and a centered guidance so that the main degree of freedom is isolated, allowing for a clean performance channel and a simple control problem. A large stroke would have been ideal to be able to see the mechanism in action as well; however, the range of motion was limited by the available actuator (PK2FVF1). The mechanism should also be easy to manufacture and to assemble, meaning that the number of parts that need to be connected needs to be as low as possible to avoid misalignments and connection stiffness differences that can impact the system's dynamics. Finally, a safety factor of two was applied to the leaf spring desing to make sure the material won't plastically deform.

2.2. Leaf Spring Design

2.2.1. Piezo patch constraints

Given that the piezo patch available was P-876.SP1[10], the leaf springs had to be designed according to its dimensions. This led to constraints regarding the length as well as the width of the leaf springs. The width of the leaf spring should be about the same as the piezo patch width in order to have optimum actuation authority. As the width of the leaf spring increases, the actuation authority decreases. Thus, to get the best performance, it was decided to set the leaf spring width to be the same as the piezo patch plus

2

1 mm of extra space at the edges to allow for the tolerances of the piezo patch itself. As for the length of the leaf spring, according to the literature, the optimum percentage of the leaf spring that should be covered is about half, and any coverage further than that reaches the point of diminishing returns[11]. Thus, it was decided to set the length of the leaf spring to be twice as long as the piezo patch. The only other dimension left to explore was the thickness. This is investigated in the following section.

2.2.2. SENSING AND ACTUATION AUTHORITY

It is important to make sure that the strain can be observed by the piezo patch and that enough actuation authority is available on the actuation side. This is related to the thickness of the leaf spring. A relationship between the strain and the thickness of the leaf spring was derived and can be found in the Appendix A. The resulting relation between the leaf spring thickness and the strain is illustrated in Figure 2.1.

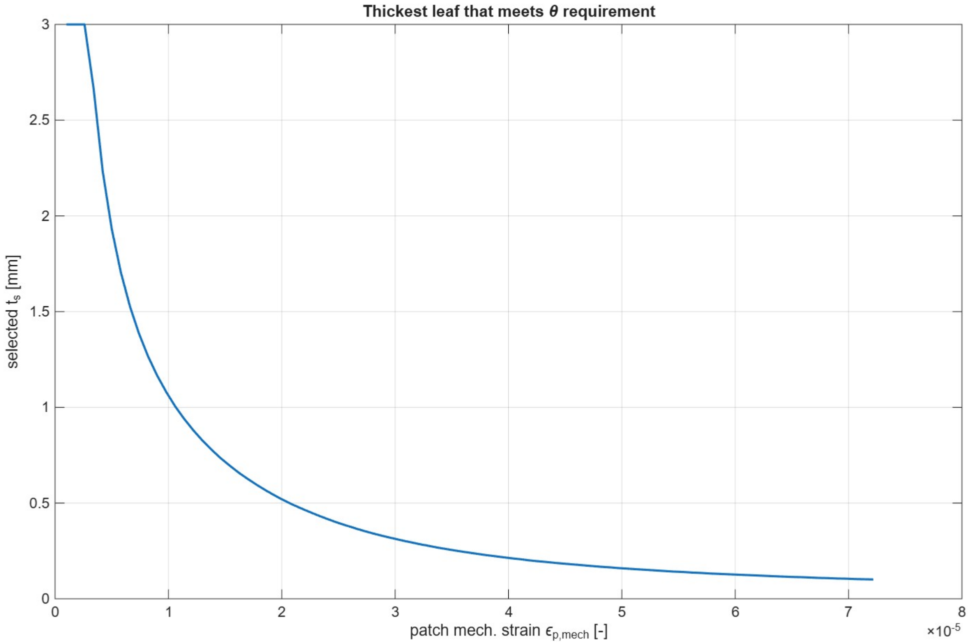

Figure 2.1: Leaf-spring thickness vs. patch mechanical strain

With the relation between the thickness and the patch strain now known, an informed decision can be made about the maximum thickness that can be chosen to allow for a higher natural frequency. The remaining question is how thick the leaf spring needs to be for sufficient actuation authority to be achieved.

For an actuated piezo patch, the free strain is just the piezoelectric strain law:

$$\varepsilon_{\text{free}} = d_{31} E_3 \quad [12] \tag{2.1}$$

The electric field through the patch thickness is:

$$E_3 = \frac{V}{t_p}$$

So:

$$\boxed{\varepsilon_{\text{free}} = d_{31}\frac{V}{t_p}}$$

Using magnitude:

$$\boxed{|\varepsilon_{\text{free}}| = |d_{31}|\frac{V_{\max}}{t_p}}$$

$$d_{31} = -180\times10^{-12}\ \text{m/V} \quad [13] \tag{2.2}$$
$$V_{\max} = 150\ \text{V} \tag{2.3}$$
$$t_p = 0.5\times10^{-3}\ \text{m} \quad [10] \tag{2.4}$$
$$\varepsilon_{\text{free}} = 54\ \mu\varepsilon \tag{2.5}$$

If the strain transfer is about 60%, about $32\mu\varepsilon$ can be expected, which can be applied to a 0.3 mm leaf spring, as is evident from Figure 2.1, and the maximum authority will theoretically still be retained by the actuator. The actuator authority will be degraded by any leaf spring thicker than $0.3mm$, and if maintaining that authority is the goal, the maximum leaf spring thickness should be set to no more than $0.3mm$. For piezo sensing, SNR is important, and as a rule of thumb 20 dB is already good for baseline detection. For the purpose of this design, 30 dB SNR was chosen as the requirement. The relation between SNR and strain is derived and can be found in the Appendix B. The minimum strain to get this SNR was calculated to be about $7\mu\varepsilon$. As can also be seen from the plot, at this strain, the maximum leaf spring thickness is about $1.53mm$. Thus, for a 30 dB SNR, a maximum of $1.53mm$ of thickness is allowed, and a maximum of $0.3mm$ is allowed for the actuation authority. Sensing is not going to be an issue according to these derivations; however, when it comes to actuation authority, the thickness cannot be more than $0.3mm$, which is what was decided. Here, the following trade-off can be established.

Thickness trade-off:

- Stiffness is increased by thicker leaf springs
- Higher bandwidth
- But lower actuation authority

2.3. MECHANISM LAY OUT

For the mechanism layout, different variants were considered, namely a conventional parallelogram, a double parallelogram, and a double four-bar parallelogram(Figure 2.2). The double four-bar parallelogram layout was selected because the requirements mentioned earlier, such as high off-axis stiffness and centered guidance, were met by it. To improve the off-axis stiffness, an extra layer of leaf springs was added. Some assumptions needed to be made; the dimensions of the stage and the intermediate piece, as well as the spacing between the leaf springs, were set to be comparable to other test-beds in the literature. These dimensions were later optimized to fulfill a set of criteria, which will be further outlined in the Section 3.

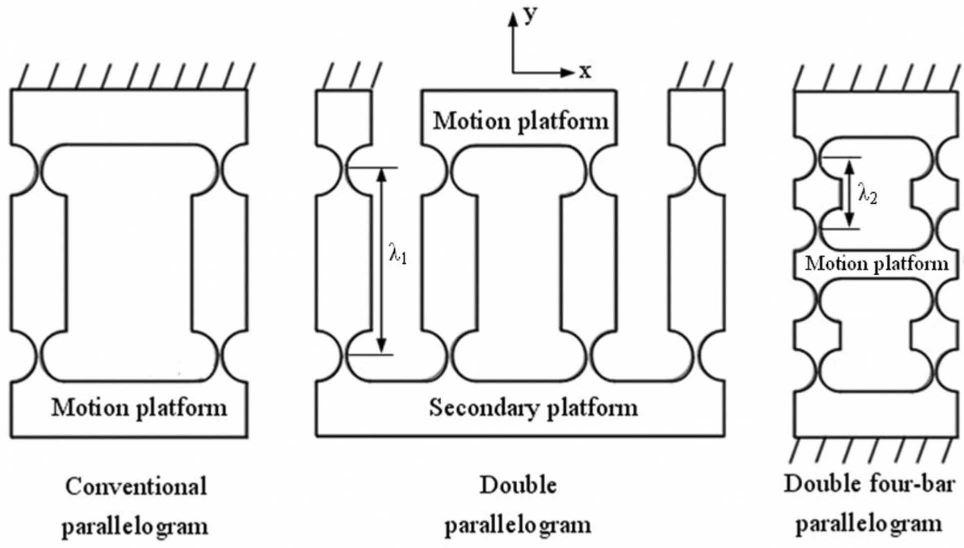

Figure 2.2: Parallelogram-based flexure layout variants considered for the mechanism design[14].

3 Control-Oriented Structural Optimization

3.1. Optimization Objective

The optimization was set up with the following objectives:

1. The first natural frequency was required to be as high as possible.
2. The pitch and roll modes were required to be shifted to higher frequencies.
3. The yaw mode was required to be dominant, instead of being a mixture of motions.
4. The yaw-dominant mode was required to appear at a certain frequency relative to the first natural frequency.

As a result, a mechanism with a higher potential bandwidth would be obtained, and the yaw mode would be well separated from the first and higher-order modes. The yaw mode was selected as the mode for which the mechanism was designed because in-plane modes would be easier to sense and actuate, while out-of-plane modes would be more challenging to handle.

The reason for aiming at mode separation and a dominant yaw mode was that the oversensing channel could be used to isolate the mode properly without spillover. This would allow the control problem in the oversensing channel to be simplified, since simple mass-spring-damper dynamics would be obtained in that channel.

3.2. Design Variables

The following parameters were optimized. Here, LiveLink was used to update the parameters in SolidWorks and to run the optimization in COMSOL.

- Leaf spring spacing
- Stage
  - Width
  - Height
  - Length
  - Thickness
- Intermediate piece thickness

3.3. OPTIMIZATION COST FUNCTIONS

The optimization cost functions that were used to achieve the aforementioned criteria are shown below. Each term was included to represent one of the optimization objectives.

$$J_{\text{tot}} = w_{\text{rot}} J_{\text{rot}} + w_{\text{freq}} J_{\text{freq}} + w_{\text{shift}} J_{\text{shift}} + w_{\text{yaw}} J_{\text{yaw}} + w_{\text{yawfreq}} J_{\text{yawfreq}}$$

Here, w_{rot}, w_{freq}, w_{shift}, w_{yaw}, and w_{yawfreq} are all the weighting factors of their corresponding objective function. The optimization node made in COMSOL minimizes the total objective function J_{tot}.

3

The term J_{freq} was used to push the first natural to higher frequencies as high as possible.

$$J_{\text{freq}} = \left(\frac{f_{1,\text{ref}}}{f_1}\right)^2$$

The terms J_{rot} and J_{shift} were used to penalize parasitic pitch and roll motion and to push these modes to higher frequencies.

$$J_{\text{rot}} = \sum_{i=1}^{N} \left(|\text{pitch}_i|^2 + |\text{roll}_i|^2\right)$$

$$J_{\text{shift}} = \sum_{i=1}^{N} \frac{|\text{pitch}_i|^2 + |\text{roll}_i|^2}{f_i^2}$$

The term J_{yaw} was used to reward modes that behave dominantly as yaw modes instead of as a mixture of motions.

$$J_{\text{yaw}} = \sum_{i=1}^{N} \left(1 - \frac{\text{yaw}_i^2}{\text{yaw}_i^2 + \text{pitch}_i^2 + \text{roll}_i^2 + \epsilon}\right)$$

The term J_{yawfreq} was used to bias the yaw-dominant response so that it appeared at a certain frequency ratio relative to the first natural frequency. For this, the yaw-weighted average frequency was used $f_{\text{yaw,bar}}$.

$$f_{\text{yaw,bar}} = \frac{\sum_{i=1}^{N} \text{freq}(\lambda_i) \cdot |\text{yaw}_i|^2}{\left(\sum_{i=1}^{N} |\text{yaw}_i|^2\right) + \epsilon_{\text{yawfreq}}}$$

$$r_{\text{yaw}} = \frac{f_{\text{yaw,bar}}}{f_1}$$

$$J_{\text{yawfreq}} = \left(\frac{r_{\text{yaw,target}}}{r_{\text{yaw}}}\right)^2$$

3.4. OPTIMIZATION RESULTS

The final optimized mechanism is shown in Figure 3.1, and the corresponding rigid body modes are shown in Figure 3.2. It can be seen that the optimization was effective, since the first natural frequency which was 57 Hz in the original design was increased to 141 Hz. The yaw mode was placed at approximately ten times the first natural frequency,

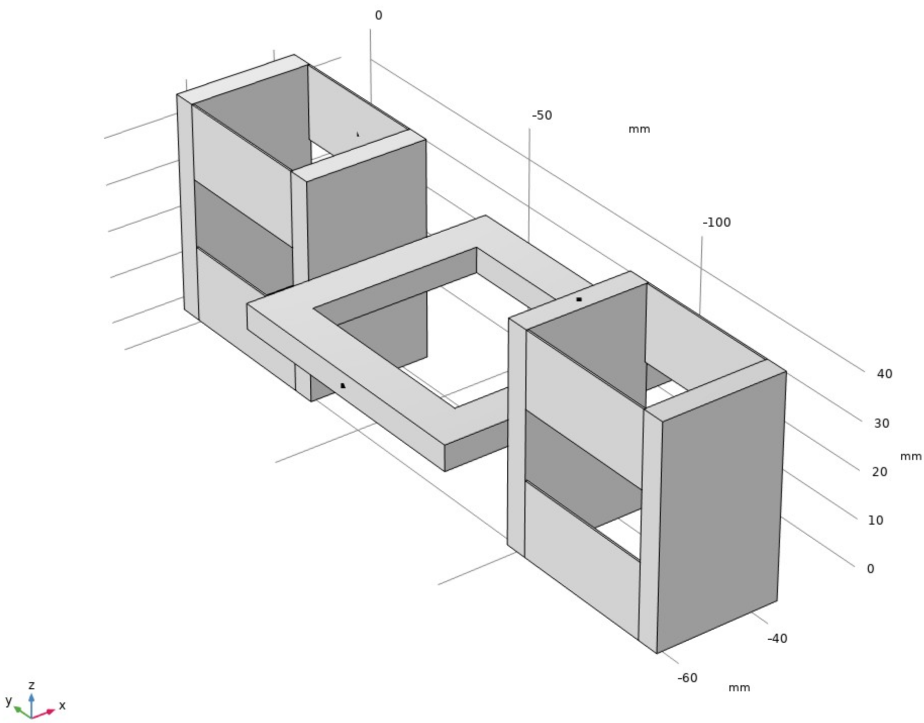

Figure 3.1: Optimized motion stage

and the pitch and roll modes were shifted to higher frequencies.

The goal was to isolate a well-separated yaw mode, and this was achieved. The first rigid-body modes were found at 141.05 Hz, 1401.4 Hz, and 2447.8 Hz. The yaw mode was located between the first and higher rigid-body modes, so it was well separated.

However, the leaf-spring modes were not included in the optimization code. This is a limitation, because these modes can have shapes that result in similar features to the rigid-body modes that were intended to be isolated. As a result, the modal isolation can become less clean, and spillover to neighbouring modes can become more likely. For example, leaf-spring modes were found between 1937.6 Hz and 2064.5 Hz, which could interfere with the isolation of the selected modes. These leaf-spring modes were not treated further in this research, since the focus was placed on the rigid-body modes of the stage.

For future work, these leaf-spring modes should also be included in the optimization code to investigate whether better mode isolation can be achieved. Another drawback of the optimization was that only a single mode was isolated. Therefore, it could also be investigated whether other rigid-body modes can be isolated from each other, so that a simpler controller problem could be obtained for all rigid-body modes.

The yaw mode was separated relatively well; however, the same was not achieved for the other rigid-body modes. For that reason, other approaches had to be explored, namely modal filtering and the use of more sensors. However, the mechanism was potentially made more complicated by this. Therefore, it is worth investigating whether the other rigid-body modes could be isolated in the same way as the yaw mode, so that they could also be isolated by a single actuator and sensor combination.

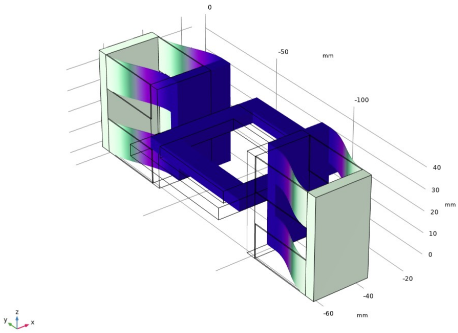

(a) Mode 1 at f_1 = 141.05 Hz.

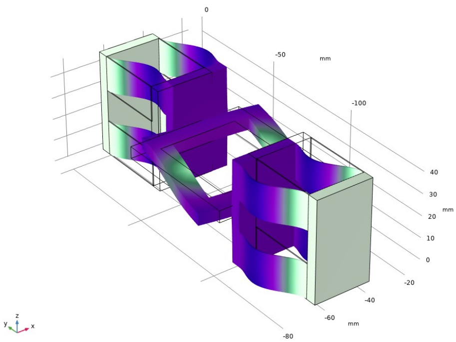

(b) Mode 2 at f_2 = 1401.1 Hz.

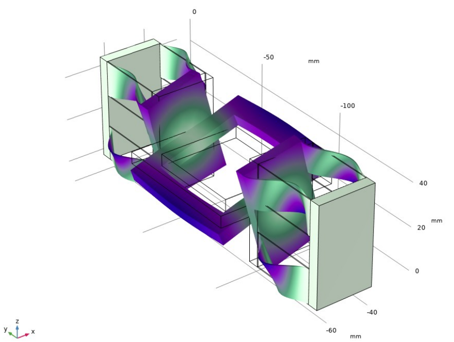

(c) Mode 3 at f_3 = 2447.8 Hz.

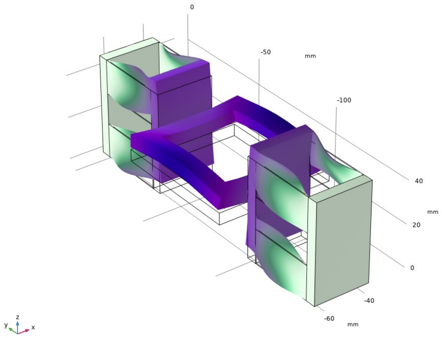

(d) Mode 4 at f_4 = 3015.1 Hz.

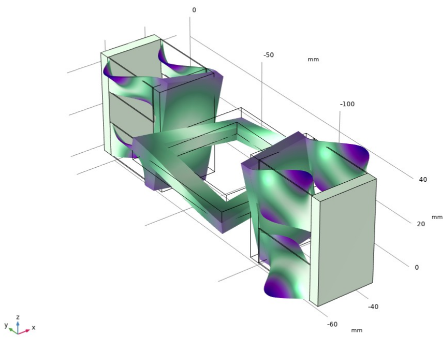

(e) Mode 5 at f_5 = 3535.3 Hz.

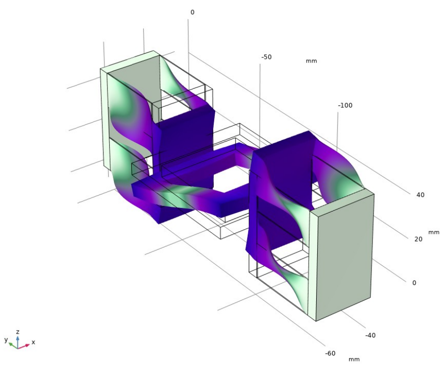

(f) Mode 6 at f_6 = 3570.1 Hz.

Figure 3.2: Selected eigenmodes of the flexure mechanism obtained from modal analysis.

4 Sensor and Actuator Placement Optimization

4.1. Placement Problem Formulation

4.1.1. Placement Criterion

It was found from the literature review that optimum observability and controllability are obtained where the strain is the highest. This means that, when a mode is to be isolated, the sensor and actuator should be placed at the anti-node of the mode of interest and at the node of the other modes. Based on this, an optimization node was made in COMSOL. The optimization node that was used for this optimization is shown in the Appendix C.

4.1.2. Main Degree-of-Freedom Placement

First, the optimization was tested for the main degree of freedom. This was done to confirm whether the optimization pipeline worked, since the position for the main degree of freedom could be estimated more easily. The result confirmed this, because the main DoF sensor was placed at the center of the opposite edge of the stage, where the optimization was run. To make the placement optimization more robust, a separate MATLAB code was made to calculate the observability, controllability, and joint HSV. The code that was used for the calculation of these matrices can be found in the Appendix D. The actuation point, shown by the red arrow, and the sensing point, shown by the green arrow, are depicted in Figure 4.1.

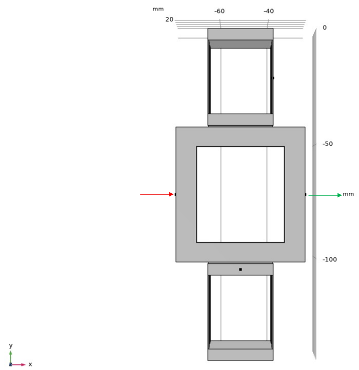

Figure 4.1: Optimum sensor position for main DoF

4.1.3. PLACEMENT FOR THE SECOND RIGID-BODY MODE

After the pipeline had been confirmed, the same approach was used to find the optimum actuator and sensor placement for the second rigid-body mode. For this optimization, different actuator coordinates were swept on one edge of the stage, while the sensor coordinates were swept on the opposite edge. The Hankel singular value (HSV) results are shown in Figure 4.2. According to these results, the optimum actuator and sensor

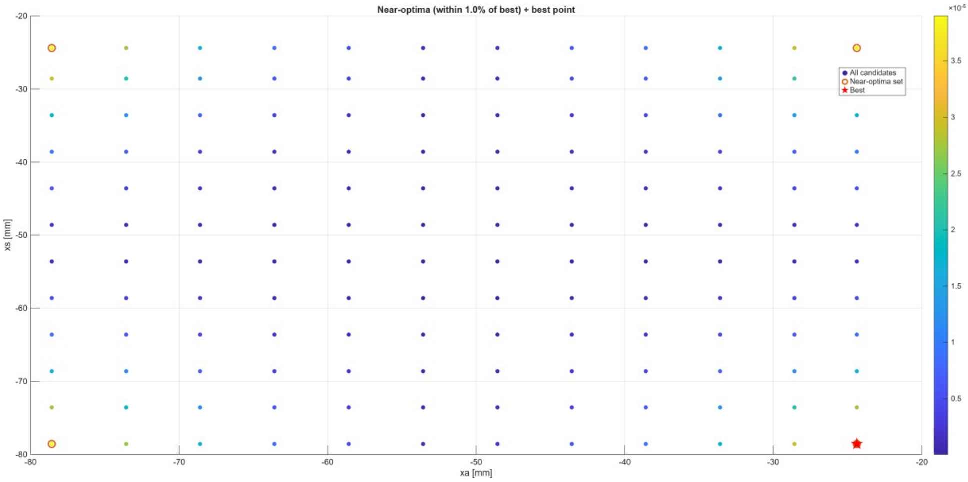

Figure 4.2: HSV of 144 actuator and sensor combinations

placement was found at the edge of the stage on opposite sides. This placement is also depicted in the CAD model in Figure 4.3.

4.1.4. MODE ISOLATION RESULT

As can be seen from the Bode plot in Figure 4.4, the mode at 1401 Hz was isolated well. However, another mode was also observed directly after this mode. This was caused by a leaf-spring mode that produced a similar deformation in the stage. It would be possible for the same actuator and sensor placement to be used for the second rigid-body mode, instead of creating a separate channel for this mode. However, this mode could also be tackled in a separate channel by using piezo patches on the leaf springs. Dealing with the second mode in the already existing channel would be more ideal, since no extra hardware would be required and the system would be less complex.

4.1.5. COMPARISON BETWEEN COMSOL AND HSV RESULTS

The COMSOL results were comparable to the HSV results. However, the COMSOL optimization took about 45 minutes, while the HSV investigation took about 16 hours for a total combination of 144 actuator and sensor positions. It should be mentioned that, even for the second rigid-body mode, an educated guess would have been possible. However, the HSV optimization method could be valuable in more complex designs, where the best actuator and sensor combination cannot be guessed easily.

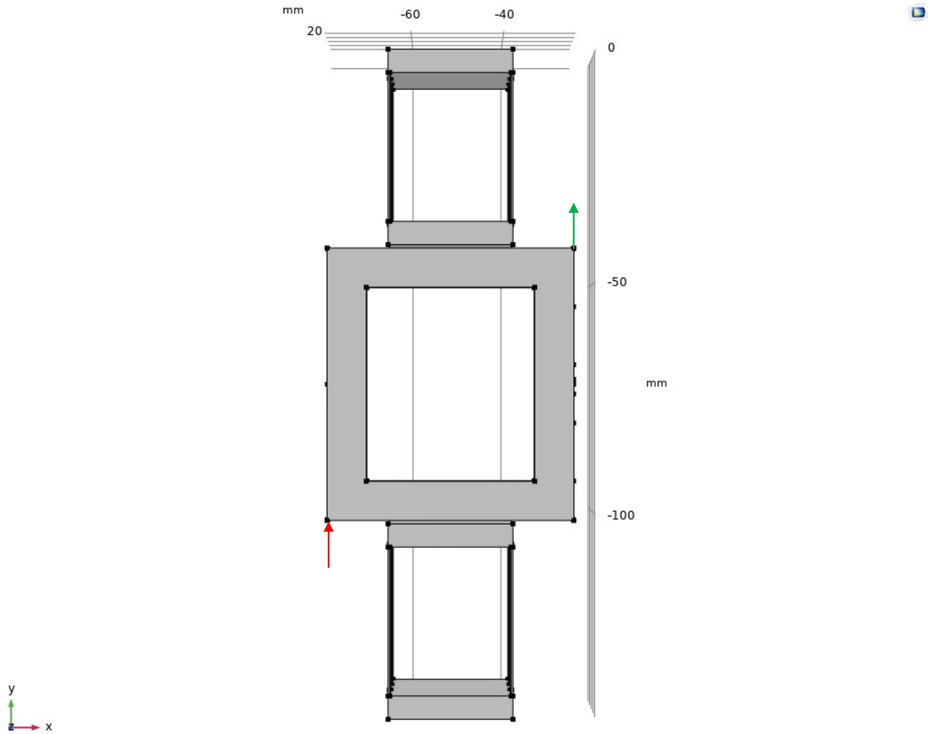

Figure 4.3: Optimum actuator(red arrow) and sensor(green arrow) position for yaw-mode

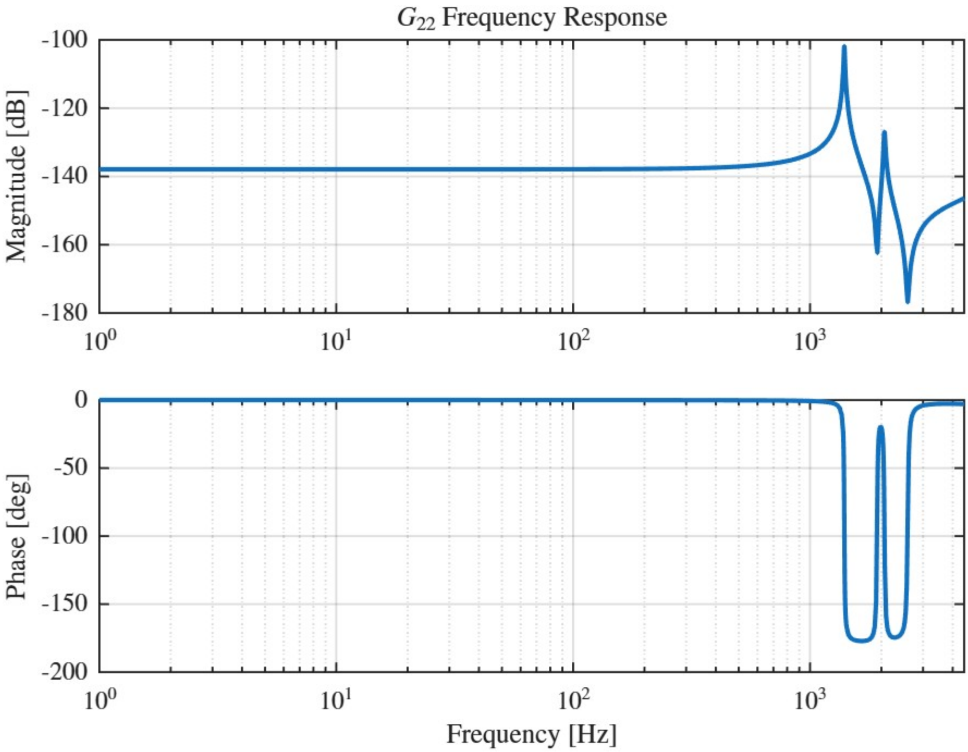

Figure 4.4: Frequency response from the optimized actuator input to sensor output

Even though the COMSOL results were comparable to the HSV results, the HSV optimization should not be considered irrelevant. The COMSOL optimization can be used to quickly find the optimum sensor and actuator position for prototyping and testing, while the HSV optimization can be used to robustly confirm the result.

5 From Physical Channels to Modal Channels

5.1. Limitation of Single-Point Sensing and Actuation

The approach from the previous chapter was shown to be effective in isolating the mode of interest that had to be tackled in the second channel. Therefore, the same approach was applied to extend the active vibration control approach, so that other rigid-body modes could also be isolated and attenuated in their own separate channels.

However, after the optimum actuator and sensor optimization had been run, the results were less promising. The other modes were not isolated clearly, and some of the channels resulted in non-minimum phase behaviour. This behaviour can be seen in the frequency response plot in Figure 5.1.

For a mode to be isolated properly, the actuator and sensor should be placed at the anti-node of the mode of interest and at the node of the other modes. This was what the optimization in COMSOL was looking for. However, this was not successful because the mode shapes shared similar features. Therefore, single-point placement was shown not to be effective in this case, and other alternatives had to be explored.

5.2. Modal Filtering as an Alternative

Since single-point actuation and sensing did not lead to sufficient mode isolation, a more robust alternative had to be investigated. It was decided that modal filtering would be used to isolate the modes more clearly.
For this purpose, five sensors were placed at one edge of the stage and five sensors were placed on the opposite edge of the stage. In addition, one sensor was placed on the intermediate piece, and another sensor was placed on the leaf spring. These locations were selected to explore which combinations of sensor data could result in better mode isolation. The sensing points were chosen because almost all of the modes had some type of deformation in these areas. The sensing locations are shown in Figure 5.2 and are encircled by green markers. The actuation points are shown in Figures 5.3 and 5.4.

5.3. Actuator and Sensor Layout

In Figures 5.3 and 5.4, each arrow represents an actuator, and its colour represents the corresponding channel.
Channel 1: red, 141 Hz
Channel 2: yellow, 1401 Hz and 2064 Hz
Channel 3: green, 1937 Hz

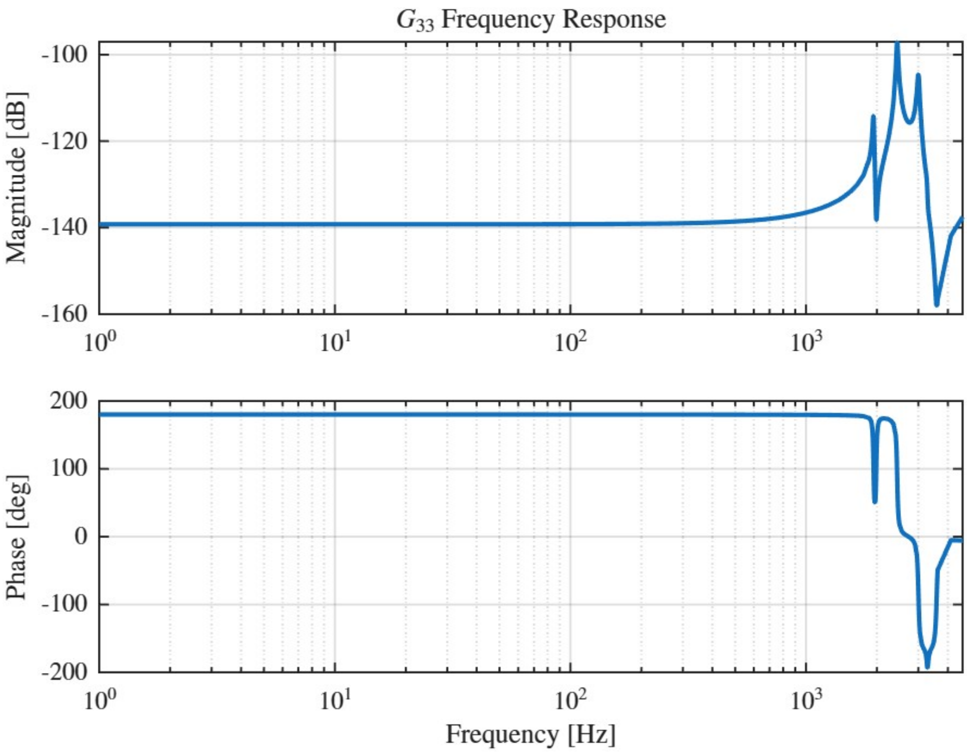

Figure 5.1: Frequency response of the optimized actuator input and sensor output for $f_3 = 2447.8$ Hz.

Channel 4: light blue, 2447 Hz
Channel 5: grey, 3015 Hz
Channel 6: pink, 3535 Hz
Channel 7: black, 3570 Hz

The ten sensors on the stage measured the z-component, while the sensor on the intermediate piece and the sensor on the leaf spring measured the x-component.
As far as the type of sensor is concerned, capacitive sensing could be used. However, finding a reference and a ground might be challenging, because the goal of the ten sensors on the stage was to measure z-displacement. For the sensor on the intermediate piece, a capacitive sensor could also be used. For the leaf spring, a piezo patch could potentially be used.
Having said that, the choice of sensor was outside the scope of this project. When this theory is tested on an actual design, the chosen sensors must be able to measure the aforementioned displacements so that the simulations can be replicated.

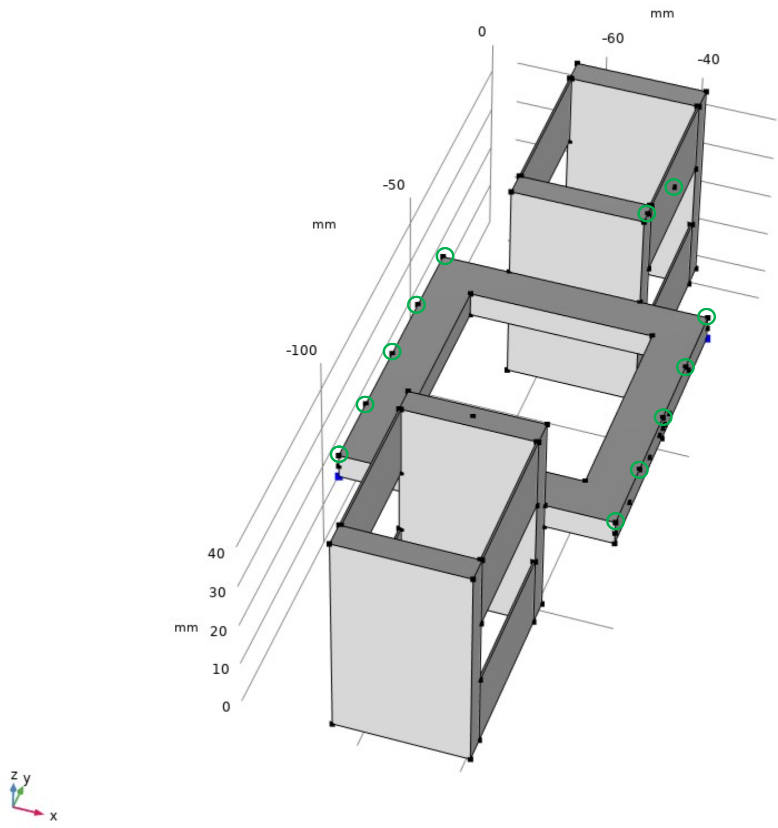

Figure 5.2: Selected sensor locations on the flexure mechanism, indicated by green circles.

5.4. VIRTUAL CHANNEL CONSTRUCTION

Each virtual channel q_m was built by applying a weighting vector over the sensor array. In this way, the weighted combination of the sensor data responded to the mode (m) that was being isolated, while the other selected modes were rejected.
For this approach, a projector was built onto the nuisance subspace and was subtracted from the target mode shape. This resulted in a weighting vector that, when multiplied by the raw sensor data, only contained the mode of interest.
The limitation of this approach was that, when there was a lot of overlap between two mode shapes, little remaining sensor information could be left for constructing the virtual channel. If this happened, other sensor locations or combinations had to be explored to ensure that enough sensor information remained to distinguish the mode of interest from the nuisance modes.
The MATLAB implementation of this approach can be found in the Appendix E.

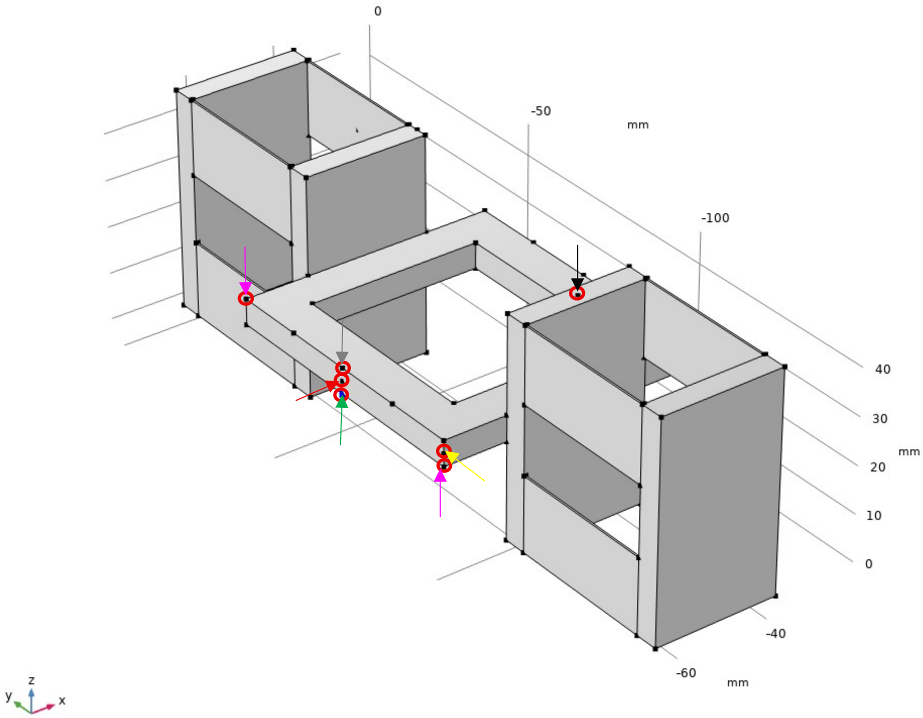

Figure 5.3: Selected actuator positions used for the active vibration control of the selected modes.

5.5. MODAL FILTERING RESULTS

This approach was first applied to channel (G_{33}), which had shown the aforementioned issues with single-point sensing and actuation. As can be seen in Figure 5.5, the troubled channel from the single-point approach was decomposed into three channels, where the modes were isolated more clearly. As a result, the control problem was made much simpler. This shows that the proposed modal filtering method achieved the intended mode isolation. After this, two other rigid-body modes were isolated using the same modal filtering approach. These results are shown in Figure 5.6. The full 7x7 magnitude and phase plot can be found in the Appendix F.

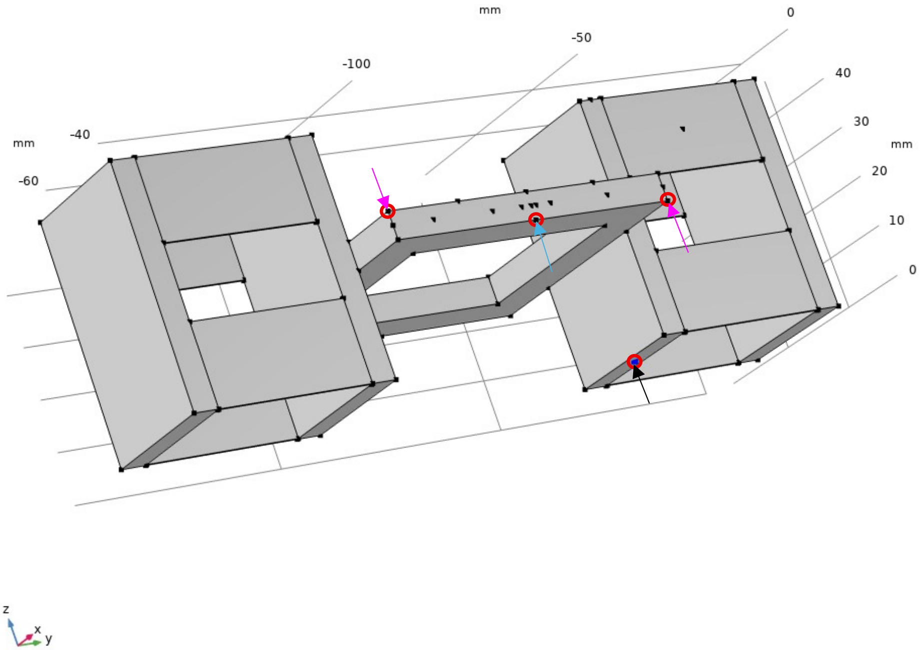

Figure 5.4: Alternative view of the selected actuator positions.

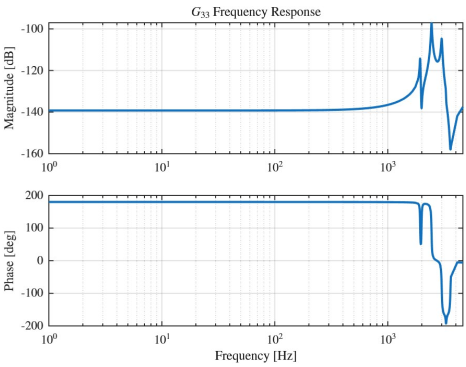

(a) Troubled channel.

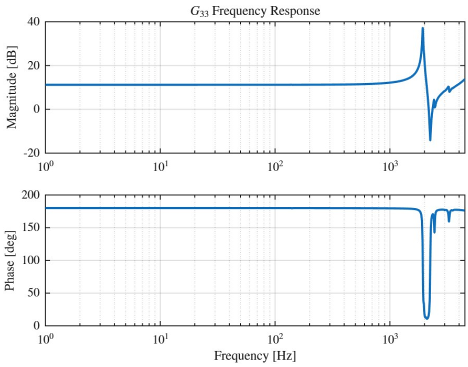

(b) Isolated mode at f = 1937.6 Hz.

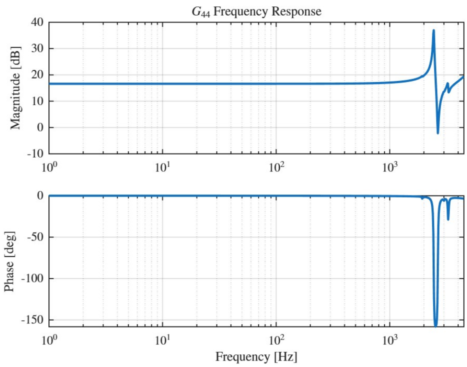

(c) Isolated mode at f = 2447.8 Hz.

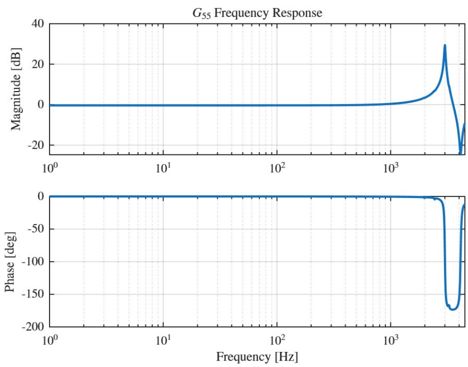

(d) Isolated mode at f = 3015.1 Hz.

Figure 5.5: Decomposed the troubled channel into three well isolated channels

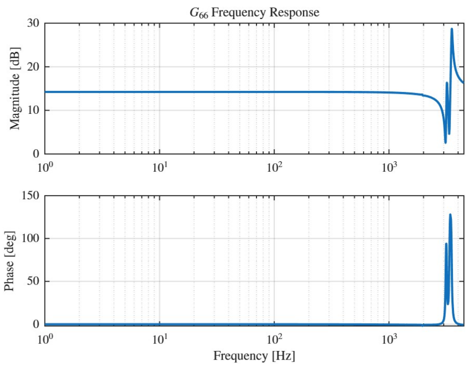

(a) Isolated mode at f = 3535.3 Hz.

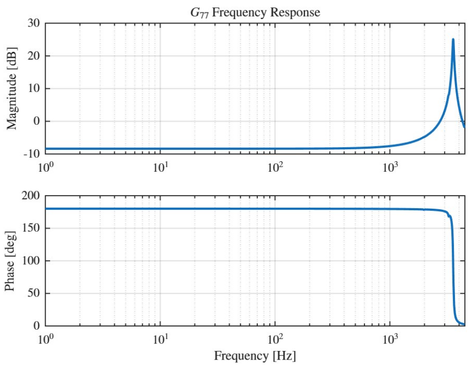

(b) Isolated mode at f = 3570.1 Hz.

Figure 5.6: Remaining two higher-order modes isolated

6 Frequency-Domain Control Design

6.1. Generation of the 7x7 MIMO Plant

The frequency response of the 7x7 MIMO plant was generated by running frequency-domain studies for every input, so that the corresponding outputs could be obtained. In this way, the full 7x7 MIMO plant was generated. The implementation can be found in the Appendix F.
The outputs were composed of the virtual channels as well as the physical channels.

6.2. Default Positive Position Feedback (PPF) Controller Design

For active vibration control, PPF controllers were used in parallel with every channel. Therefore, each channel was given its own PPF controller. Second-order filters were tuned to the target modal frequencies to attenuate the mode in the corresponding loop without spillover into the other channels[15], [16].

The standard PPF controller used in this thesis was written as

$$C_{\mathrm{PPF},k}(s) = \alpha_k \frac{g_k \omega_k^2}{s^2 + 2\zeta_k \omega_k s + \omega_k^2}, \tag{6.1}$$

where g_k is the controller gain, ζ_k is the controller damping ratio, ω_k is the tuning frequency, and α_k is the feedback sign.

The open-loop and closed-loop responses of the diagonal channels are shown in Figure 6.1. As can be seen, attenuation was achieved in most channels without significant spillover or waterbed effect. However, spillover from channel 4 into channel 5 was observed. This was caused by the strong coupling between the actuator placements of these two channels. Even though this spillover was present, an overall improvement in the frequency response was still obtained. If this behaviour is not acceptable, other actuator positions could be investigated so that the coupling and resulting spillover can be reduced.

6.3. Comparison Between Default PPF and Band-Pass PPF

Although promising results were obtained with the default PPF controller, a band-pass PPF controller was also investigated. This was done to determine whether smaller spillover could be obtained when the peak attenuation of both controllers was kept the same for a fair comparison.

The band-pass PPF controller was written as[17]:

$$C_{\mathrm{BP},k}(s) = \alpha_k g_{n,k} \left(\frac{s}{s^2 + 2\zeta_{c,k}\omega_k s + \omega_k^2} \right)^{n_k} h_{m,k}(s+\omega_k)^{m_k}. \tag{6.2}$$

The scaling terms were defined as

$$g_{n,k} = M_k (2\zeta_{c,k}\omega_k)^{n_k}, \tag{6.3}$$

$$h_{m,k} = 2^{-m_k/2}\omega_k^{-m_k}. \tag{6.4}$$

In this thesis, the band-pass magnitude level was matched to the normal PPF magnitude at the tuning frequency and then adjusted using a scaling factor:

$$M_k = \frac{g_k}{2\zeta_k}\beta_k, \tag{6.5}$$

where β_k is the band-pass scaling factor.

The comparison between the two controllers is shown in Figure 6.2. It is evident from the figures that an improvement was obtained when the band-pass PPF controller was used. The band-pass PPF controller performed better in terms of spillover across all channels, except for channel 6, where the DC gain was slightly lower with the default PPF controller.

6.4. SENSITIVITY-FUNCTION ANALYSIS

Furthermore, sensitivity-function analysis was performed to verify the spillover behaviour and to check how much each mode was attenuated. The sensitivity-function plots are shown in Figure 6.3.
The green dashed line was used to indicate the -6 dB line, which was used as the criterion for sufficient attenuation. The red dashed line was used to indicate the +6 dB line, where the spillover and waterbed effect were larger than expected.

It can be seen that the -6 dB criterion was reached by every channel. However, the band-pass PPF controller outperformed the default PPF controller by a large margin in the second channel. The +6 dB limit was exceeded by the default PPF controller, which means that a significant waterbed effect was present. In contrast, the response obtained with the band-pass PPF controller remained relatively flat. It should also be noted that the first channel was completely flat, because no controller was active in that channel.

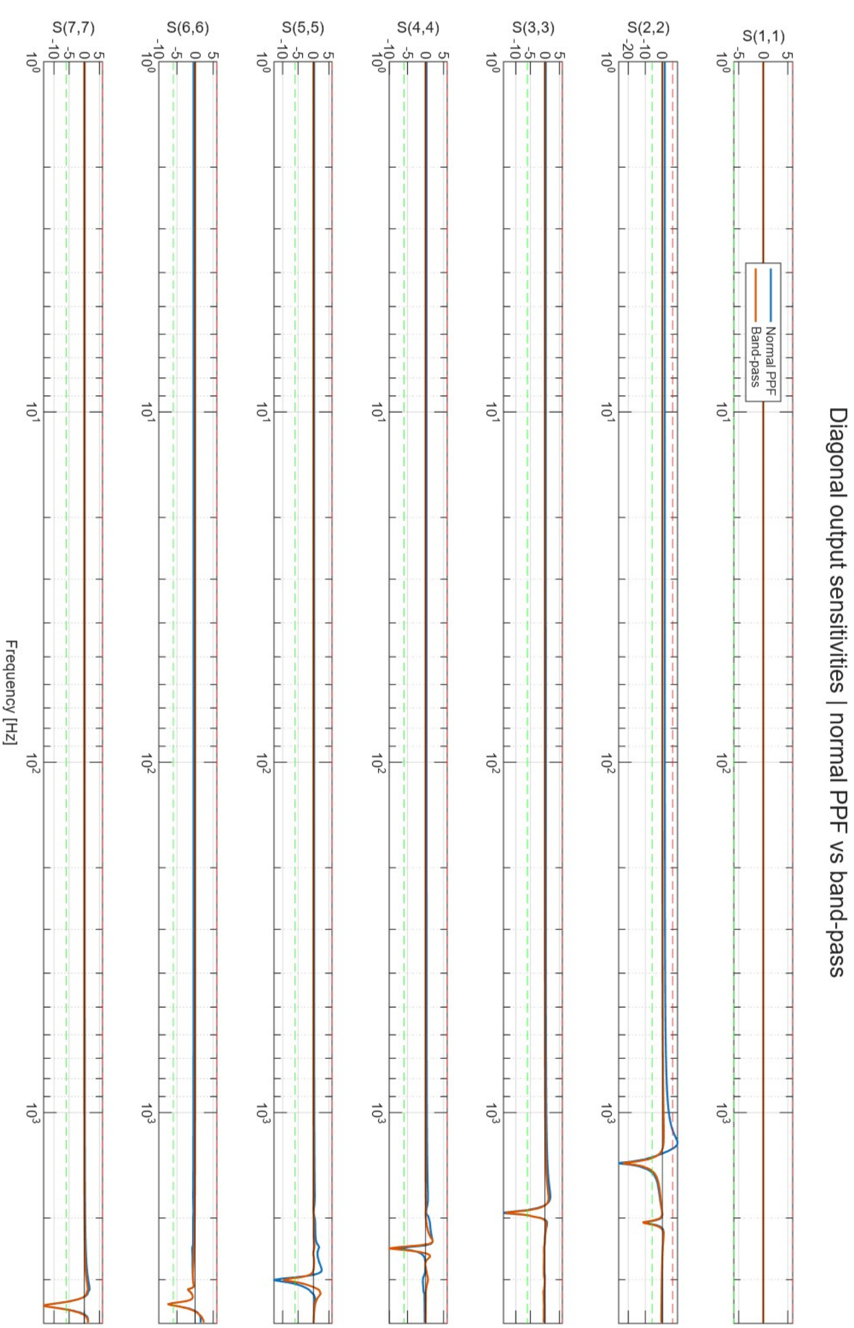

Figure 6.3: Diagonal output sensitivities of the 7×7 MIMO plant for the normal PPF and band-pass PPF controllers.

6

6.5. FULL MIMO SPILLOVER CHECK

From the frequency-response plots and the sensitivity-function plots, it was shown that the band-pass PPF controller outperformed the default PPF controller. The targeted modes were attenuated sufficiently, while the spillover and waterbed effect were kept limited.

The full MIMO comparison of the base plant, default PPF controller, and band-pass PPF controller can be found in the Appendix G. As can be seen from the full MIMO comparison, the off-diagonal channels overlapped relatively well. Therefore, no major spillover issue was observed from the diagonal channels.

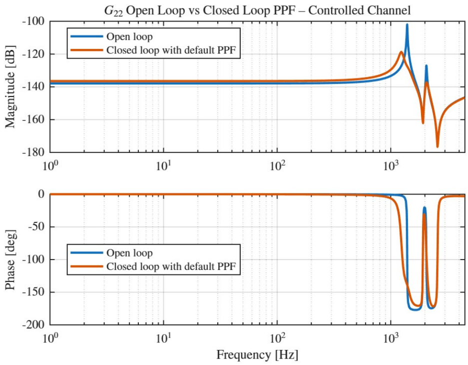

(a) G_{22} response.

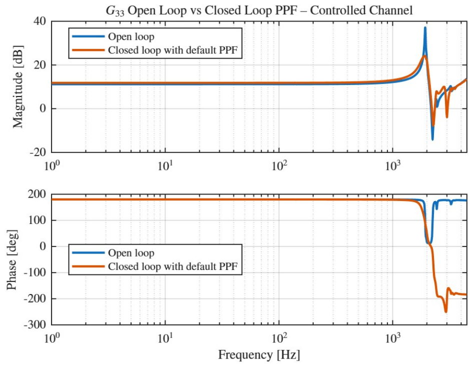

(b) G_{33} response.

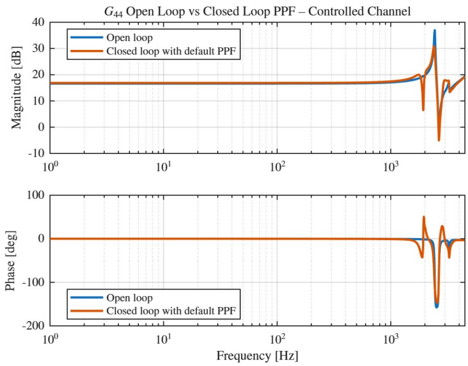

(c) G_{44} response.

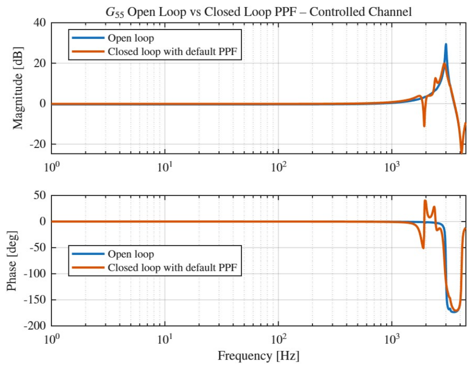

(d) G_{55} response.

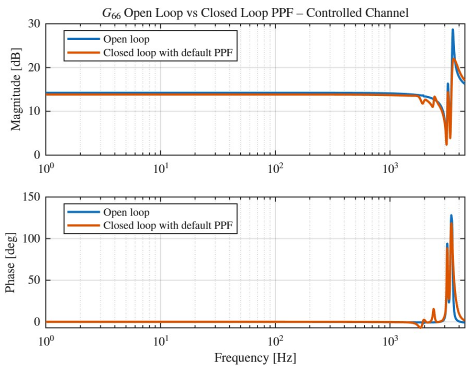

(e) G_{66} response.

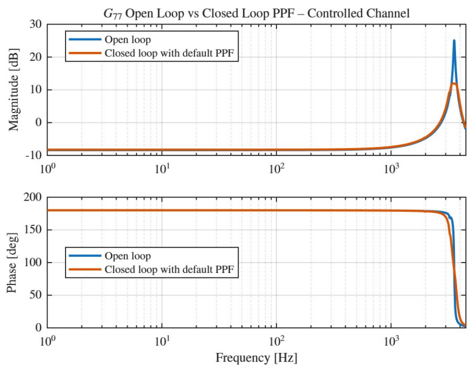

(f) G_{77} response.

Figure 6.1: Open-loop and closed-loop diagonal responses of the controlled channels using the default PPF controller.

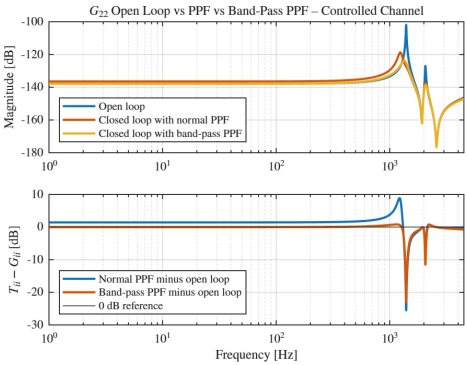

(a) G_{22} response.

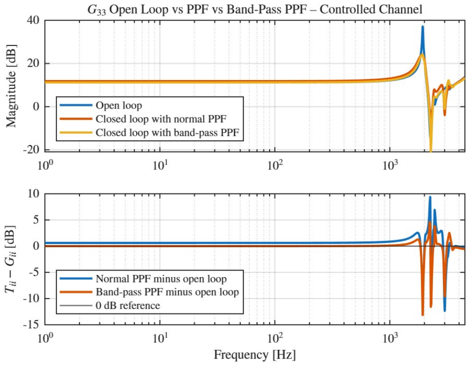

(b) G_{33} response.

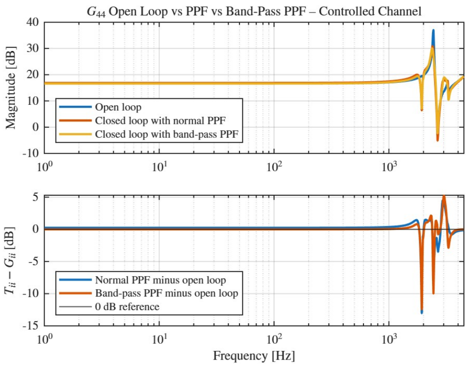

(c) G_{44} response.

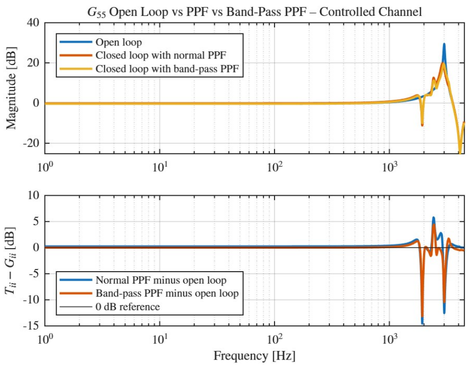

(d) G_{55} response.

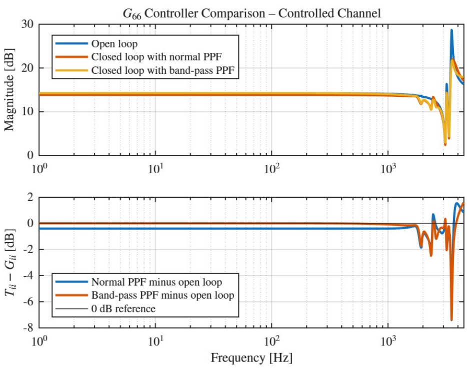

(e) G_{66} response.

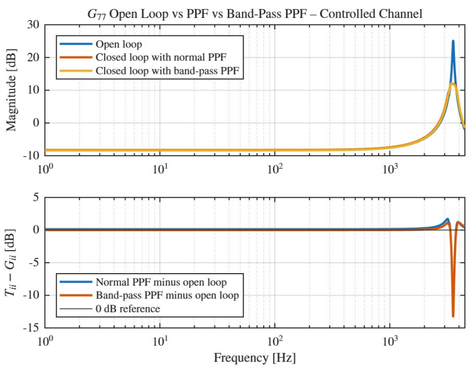

(f) G_{77} response.

Figure 6.2: Comparison of the open-loop, normal PPF, and band-pass PPF responses for the controlled diagonal channels.

7 Reduced-Order State-Space Benchmark Model

7.1. Time-Domain Analysis

After mode isolation using sensor fusion and attenuation using PPF control had been shown to be effective in the frequency domain, a time-domain analysis was performed to further confirm these findings. To perform the simulations in the time domain, a state-space model of the MIMO plant was required, since time-domain simulations could not be performed directly on the generated FRD data.

7.2. State-Space Estimation

To obtain the state-space model from the FRD data, the ssest function in MATLAB was used. More specifically, the AAA algorithm was used to estimate the state-space model for the MIMO system.
Initially, the full 7x7 MIMO plant was considered. However, the MIMO plant had to be reduced to a 5x5 system. This was done because a 5x5 state-space estimate that was close to the FRD data could be obtained, while the higher-dimensional MIMO estimation required too much memory. Around 86 GB of memory was needed due to the high-order estimate, and higher MIMO dimensions caused MATLAB to crash.

The only way to allow the ssest function to complete for the higher-dimensional case was to reduce the state-space order. However, this resulted in an estimated plant in which the MIMO FRD data was not represented properly. Therefore, a reduced 5x5 MIMO model was selected, because it gave a state-space estimate that was very close to the original FRD data. The stability of the estimated state-space plant was checked using the isstable() command in MATLAB. The comparison of the diagonal terms is shown in Figure 7.1. The comparison between the reduced 5x5 estimated state-space MIMO model and the FRD data is shown in the Appendix H.

7.3. Band-Pass PPF Implementation on the State-Space Model

After the 5x5 MIMO state-space model had been obtained and had been shown to be close to the FRD data, the band-pass PPF controllers were added in parallel to the channels.

The tuning values used for the final reduced 5×5 state-space model are listed in Table 7.1. These values were selected based on the frequency-domain tuning and were

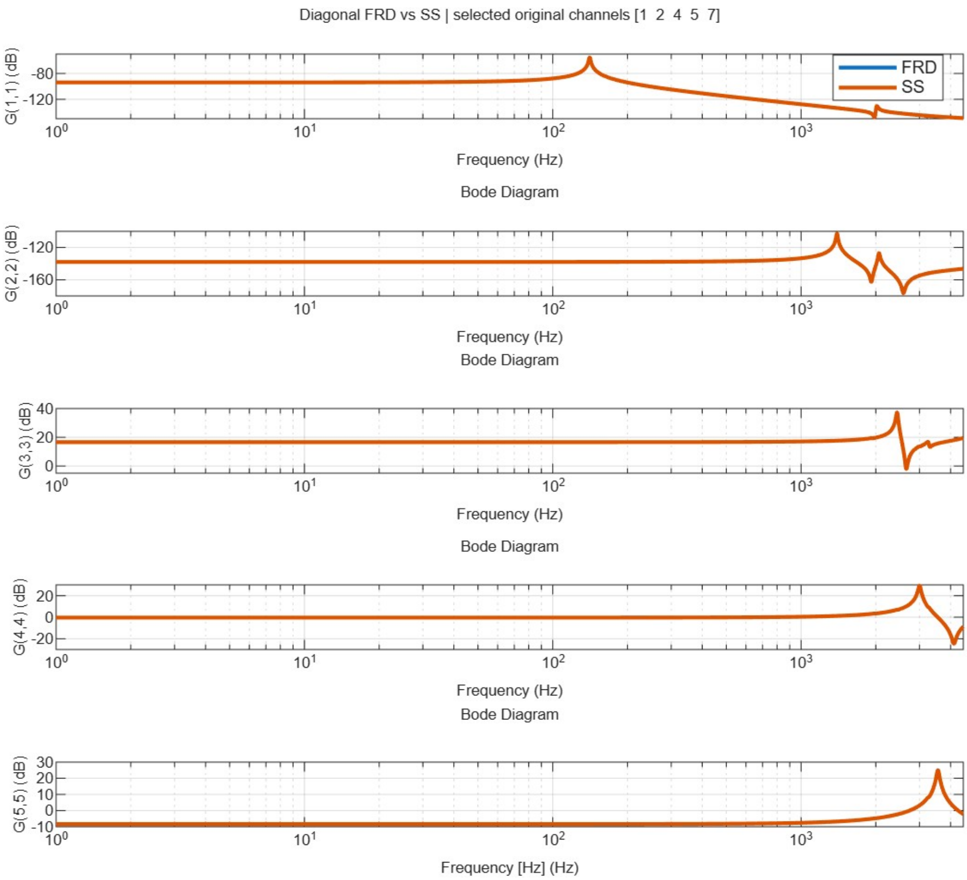

Figure 7.1: Comparison between the original FRD data and the estimated state-space model

then applied to the estimated state-space model.

In total, five PPF controllers were used. Two of these controllers were used to target the two modes in the second channel, while the other three controllers were used to attenuate the other isolated modes in their own channels. The band-pass PPF controllers were applied to the state-space model to verify whether the controllers that had been tuned on the FRD data were still effective on the estimated state-space model. The magnitude responses of the diagonal channels are shown in Figure 7.2. It can be seen that the band-pass controllers attenuated the modes as intended. In addition, the closed-loop system, including all controllers, was stable.
The full 5x5 state-space plant including the band-pass PPF controllers can be found in the Appendix I.

Table 7.1: Band-pass PPF tuning values used for the reduced 5×5 state-space model.

Mode	Channel	f_k [Hz]	α_k	n_k	m_k	$\zeta_{c,k}$
1	2	1401.4	−1	1	2	0.5
2	2	2064.5	−1	1	2	0.1
3	3	2600.0	−1	3	0	0.1
4	4	3200.0	−1	1	1	0.05
5	5	3570.1	+1	1	2	0.1

7.4. Peak Reduction and Spillover Evaluation

The spillover of the diagonal channels is shown in Figure 7.3, and the corresponding results are shown in Table 7.2. The peak reduction was defined as the difference between the open-loop and closed-loop resonance peak magnitudes in the local frequency window of each mode. A mode was considered successfully attenuated when the resonance peak was reduced by at least 6 dB.

The spillover value was defined as the maximum closed-loop amplification outside the targeted modal frequency windows on the corresponding diagonal channel. Therefore, the controller was considered acceptable on the diagonal channels when the off-target amplification remained below +6 dB.

Table 7.2: Summary of the resonance attenuation and spillover check for the targeted modes using band-pass PPF control.

Mode	Mode peak frequency [Hz]	Peak reduction [dB]	Attenuation criterion: ΔG ≤ -6 dB	Max diagonal spillover [dB]	Spillover criterion: ΔG ≤ +6 dB
1	1401.4	21.98	Pass	0.81	Pass
2	2064.5	11.66	Pass	0.81	Pass
3	2447.8	7.62	Pass	3.28	Pass
4	3015.1	9.88	Pass	2.33	Pass
5	3570.1	13.06	Pass	1.19	Pass

7.5. Sensitivity-Function Check

The sensitivity-function plot is shown in Figure 7.4. It can be seen that the response also satisfied the ±6 dB criteria. Therefore, the state-space implementation confirmed the frequency-domain results, since the targeted modes were attenuated while the spillover remained within the acceptable range.

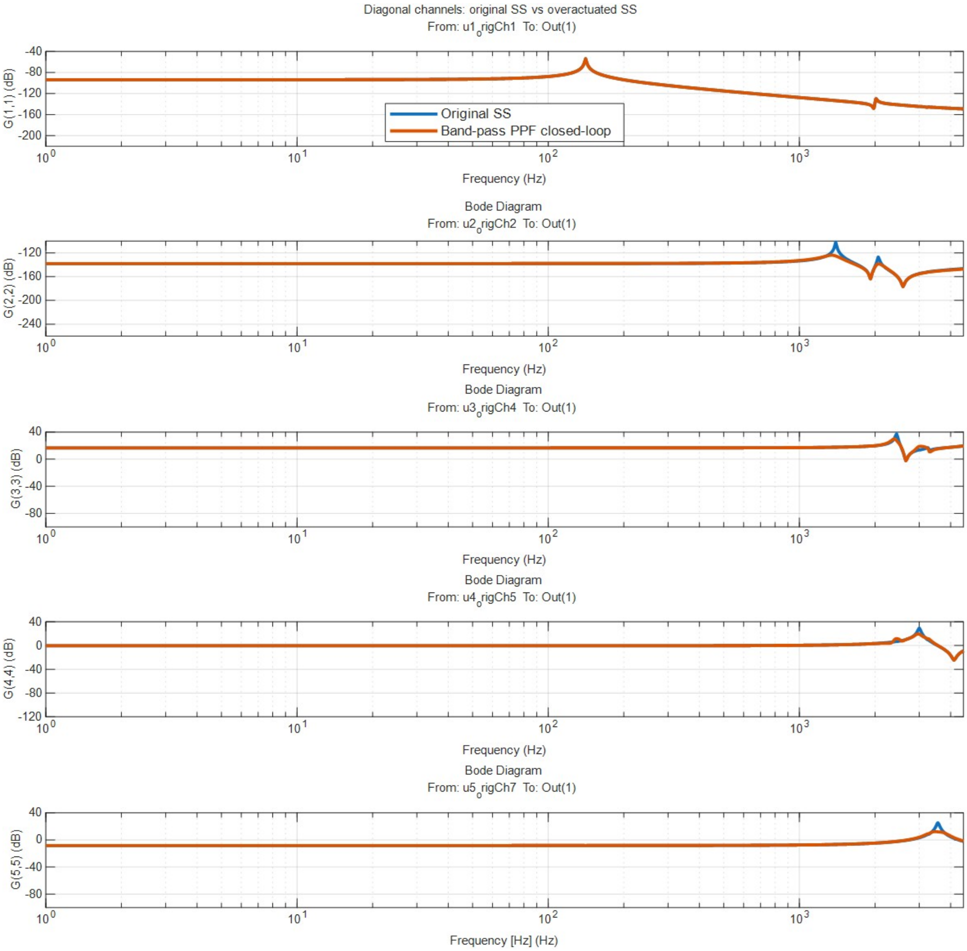

Figure 7.2: Diagonal magnitude responses of the reduced 5×5 state-space model before and after band-pass PPF control.

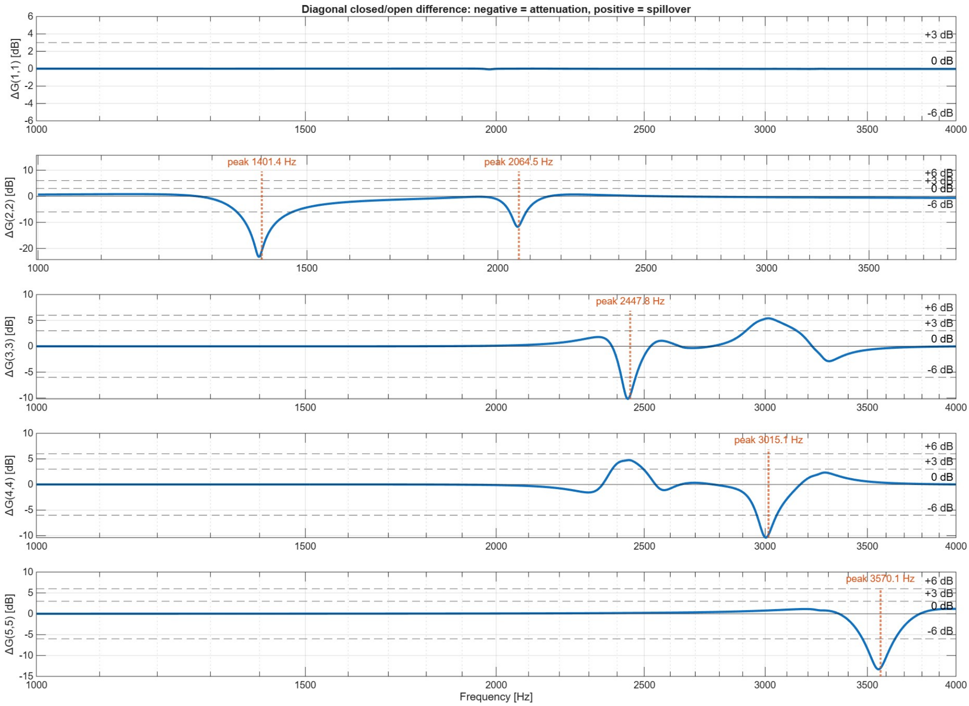

Figure 7.3: Diagonal closed-loop/open-loop magnitude difference for the reduced 5×5 model with band-pass PPF control.

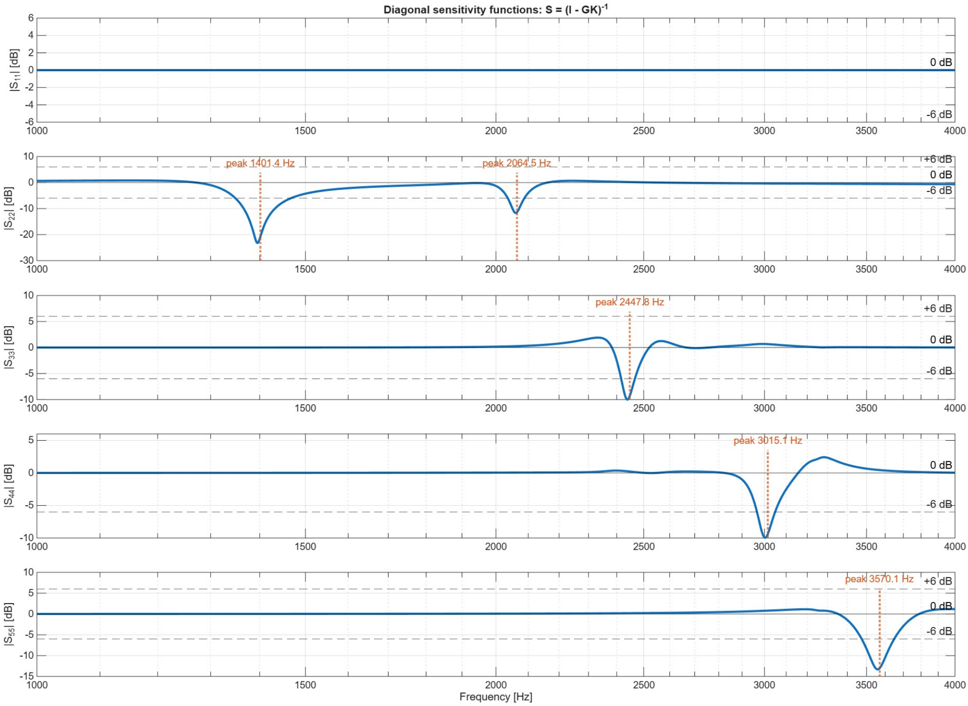

Figure 7.4: Diagonal sensitivity functions of the reduced 5×5 model with band-pass PPF control.

8 Time-Domain Benchmark

8.1. Time-Domain Simulation Objective

After the frequency-domain analysis on the estimated state-space model had been performed and had been shown to be effective, a time-domain analysis was performed. This was done to verify the frequency-domain findings and to gain insight into the challenges that can be introduced when the mechanism is following a trajectory and external disturbances, such as floor vibration and sensor noise, are present during operation.

8.2. Simulation Architecture

The simulation architecture is shown in the block diagram in Figure 8.1. In the inner loop, the band-pass PPF controllers were added in parallel to the plant to improve disturbance rejection and active vibration control.
A reference trajectory was applied to the main plant path, while floor-vibration disturbances were injected in the form of modal forces. TRUE x_1 was defined as the main DoF output of the plant without sensor noise. Sensor noise was added to the oversensing channels in the inner loop.

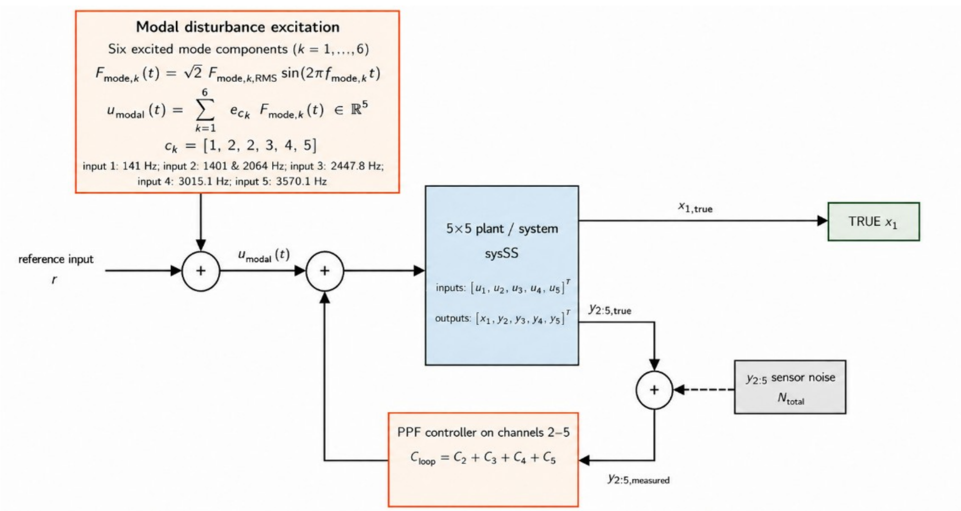

Figure 8.1: Time-domain simulation architecture used to evaluate the reduced 5×5 state-space model with band-pass PPF control.

8.3. TRACKING TRAJECTORY

The reference trajectory was chosen as a triangular displacement signal with a peak amplitude of 10 μm. The trajectory frequency was selected as one tenth of the assumed closed-loop bandwidth, giving

$$f_{\text{track}} = 14.1 \text{ Hz}$$

for a bandwidth of 141 Hz. With this choice, a relatively slow positioning command was represented compared with the first resonance, while repeated dynamic motion was still required from the stage.

In the simulation, the desired x_1 displacement was converted to the main plant input using the magnitude of the identified G_{11} transfer function at the tracking frequency. The tracking error was then evaluated as

$$e_{x_1}(t) = x_{\text{ref}}(t) - x_1(t)$$

For the tracking-only and mixed-disturbance cases, $x_{ref}(t) = x_{track}(t)$ was used. For the disturbance-only case, $x_{ref}(t) = 0$ was used.

8.4. FLOOR VIBRATION DISTURBANCE

Floor vibration was represented by converting an assumed base displacement amplitude into an equivalent inertial force. A displacement amplitude of 1 μm_{RMS} was assigned to each selected modal frequency. The corresponding RMS force was calculated as

$$F_{\text{floor,RMS}} = m_{\text{moving}}(2\pi f_i)^2 x_{\text{floor,RMS}}$$

where m_{moving} is the moving mass of the stage, f_i is the considered modal frequency, and $x_{\text{floor,RMS}}$ is the assumed base displacement amplitude. The resulting force components were then applied as sinusoidal input disturbances at the corresponding plant input channels. Because the force amplitude was scaled with f_i^2, a 1 μm_{RMS} displacement at kHz modal frequencies resulted in large equivalent forces. Therefore, this case should be interpreted as a stress test for modal disturbance rejection rather than as a direct representation of a measured broadband floor-vibration spectrum.

8

8.5. SIMULATION CASES

Three simulation cases were considered:

1. Tracking only
2. Floor vibration only
3. Tracking combined with floor vibration

8.6. SENSOR NOISE MODELS

To evaluate the practical limitation of oversensing, white sensor noise was added to the active BP-PPF sensing channels. The noise was applied to the oversensing channels used by the inner modal damping controller, namely the reduced channels y_2, y_3, y_4, and y_5.

Independent white-noise signals were generated for each active sensing channel and scaled according to selected one-sided amplitude spectral densities in $\mathrm{pm}/\sqrt{\mathrm{Hz}}$. The sample-domain RMS value was obtained from

$$\sigma_{\text{sample}} = S_n \sqrt{\frac{f_s}{2}}$$

where S_n is the sensor-noise amplitude spectral density and f_s is the sampling frequency. The noisy sensor signal was passed through the corresponding diagonal BP-PPF controller. The resulting noisy control action was then propagated through the closed-loop plant. In this way, the realistic over-actuation case captured how oversensor noise could be fed back into the plant as an actuation signal.

The main performance metric was chosen as the amplitude spectral density of the true x_1 error. By using this metric, the mechanical effect of the over-actuated BP-PPF controller could be evaluated directly. The measured x_1 error was not used as the primary metric because measurement noise on the performance sensor could mask the physical vibration reduction. For this reason, sensor noise was not applied to the true x_1.

8.7. TIME-DOMAIN BENCHMARK RESULTS

The time-domain benchmark was evaluated for two different oversensing cases: an ideal oversensing case and a realistic oversensing case. With this comparison, the theoretical benefit of the over-actuated and oversensed architecture was separated from the practical limitations introduced by sensor noise.

The ASD of the true x_1 error was plotted for the three input cases. These plots are shown in Figures 8.2, 8.3, and 8.4.

In the tracking-only case, the true x_1 error ASD was almost unchanged by the band-pass PPF controller. This indicated that the tracking trajectory did not significantly excite the targeted flexible modes. Since the controller was designed to suppress resonant modal motion, little improvement was expected when the error was dominated by low-frequency tracking components. Therefore, it was confirmed that the proposed over-actuated control strategy was not intended to improve low-frequency tracking error directly, but rather to attenuate flexible resonances when they were excited.

8

In the floor-vibration case, the flexible modes of the stage were excited by the disturbance, resulting in clear resonance peaks in the true x_1 error ASD. Compared with the no-controller case, these peaks were significantly attenuated by the ideal band-pass PPF controller. This demonstrated that the over-actuated architecture could reduce modal vibration when the disturbance had sufficient modal content.

The realistic cases showed the influence of oversensor noise. For low noise densities, the realistic response remained close to the ideal response. This meant that the added sensing channels were sufficiently clean to provide useful damping. However, as the oversensor noise density was increased, the broadband ASD level was also increased. This was caused by the sensor noise being fed back through the controller and applied to

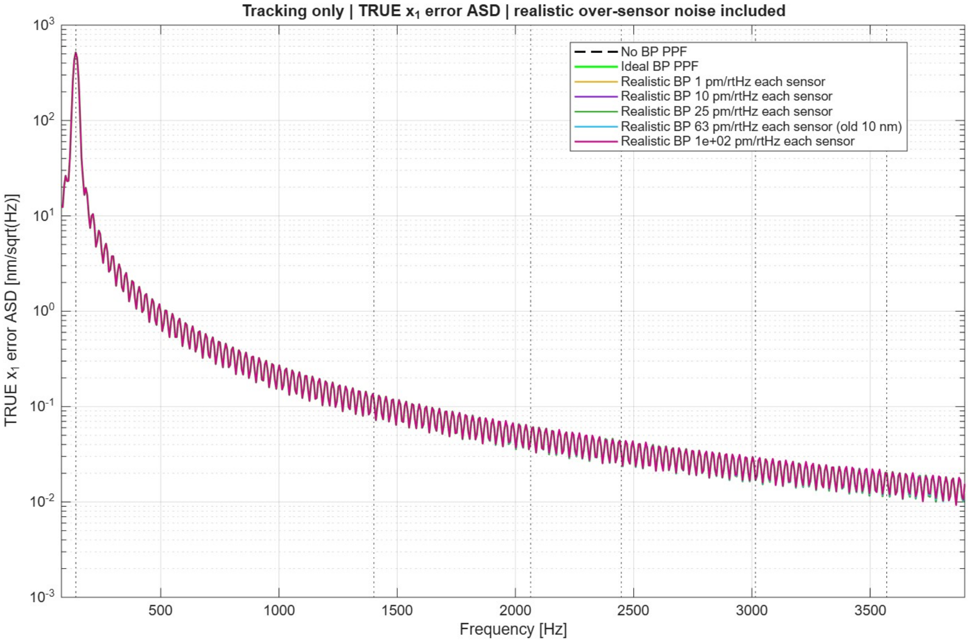

Figure 8.2: ASD of the true x_1 tracking error for the tracking-only case with different over-sensor noise levels.

the plant as actuator force. Therefore, the improvement from over-actuation was limited by the noise level of the additional sensing channels.

In the combined tracking and floor-vibration case, the true x_1 error contained both low-frequency tracking components and high-frequency modal components. The low-frequency part of the ASD was almost unaffected by the controller because the band-pass PPF was not designed to act in this frequency range. However, clear attenuation was obtained at the targeted resonance frequencies compared with the no-controller case. This confirmed that the over-actuated controller suppressed the flexible modal contribution to the error, even though the total error remained strongly influenced by the tracking trajectory.

To summarize, the ideal oversensing case demonstrated the maximum vibration reduction that could be achieved when the modal responses were measured perfectly. The realistic oversensing case showed that this benefit was only retained when the additional sensing channels had sufficiently low noise. If the oversensing noise was too high, this noise was fed back into the plant as actuator force, which reduced or hid the benefit of over-actuation. Therefore, over-actuation and oversensing were most useful for suppressing targeted resonant modes when the added modal sensing channels were sufficiently clean.

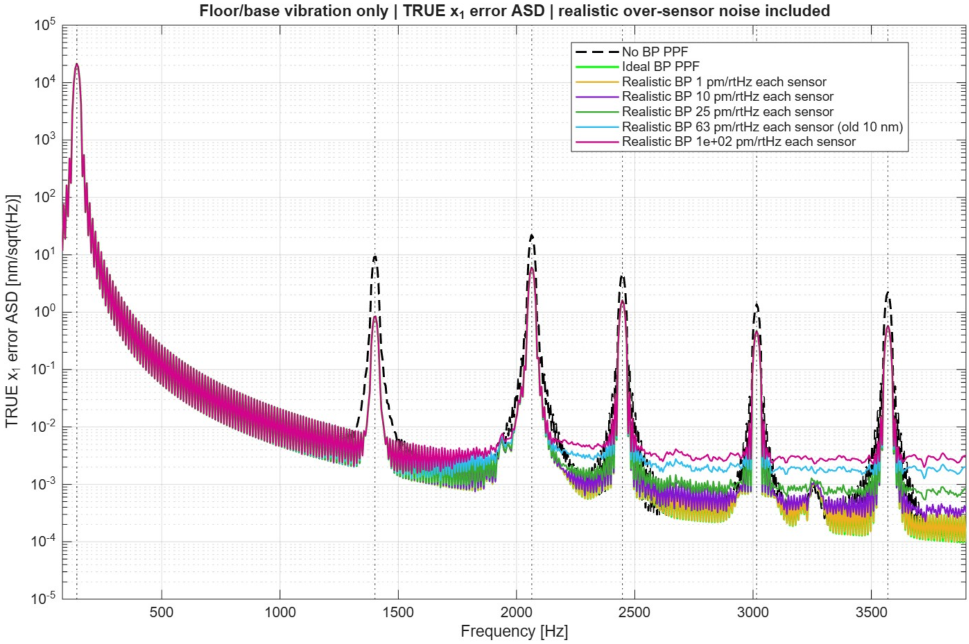

Figure 8.3: ASD of the true x_1 error for the floor-vibration-only case with different over-sensor noise levels.

8.8. SINE-BURST RESPONSE

Furthermore, the response of the controlled modal outputs to targeted sine-burst excitations at the selected resonance frequencies is shown in Figure 8.5. The vertical dotted lines indicate the start and end of the excitation burst.

In all cases, larger residual oscillations were observed after the burst in the uncontrolled system. In contrast, a faster decay and lower vibration amplitude were obtained with the BP-PPF controlled system. This indicated that the BP-PPF controller increased the effective damping of the targeted higher-order modes and reduced residual vibrations.

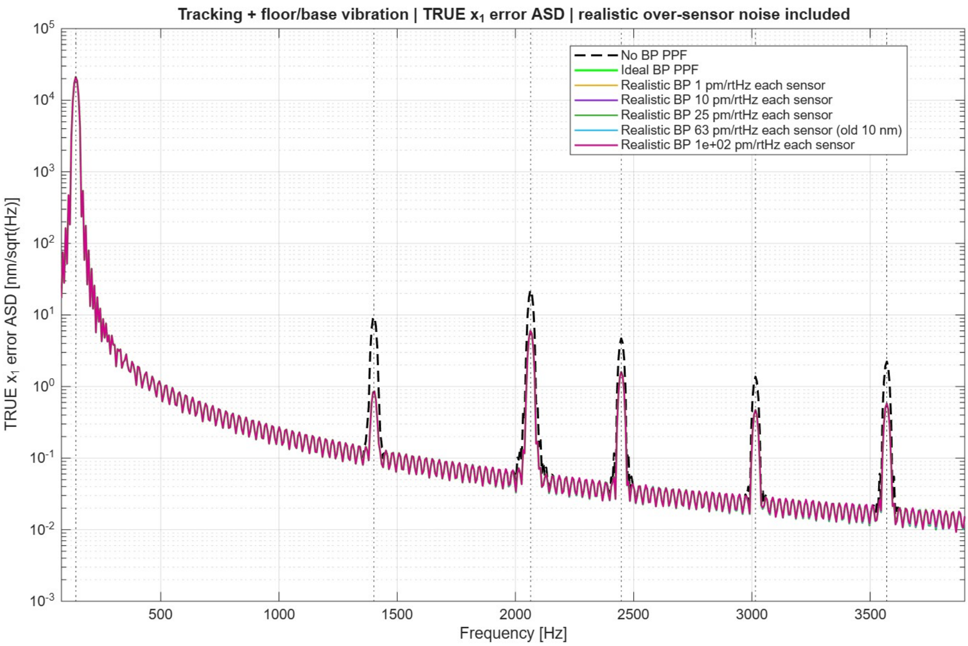

Figure 8.4: ASD of the true x_1 error for the combined tracking and floor-vibration case with different over-sensor noise levels.

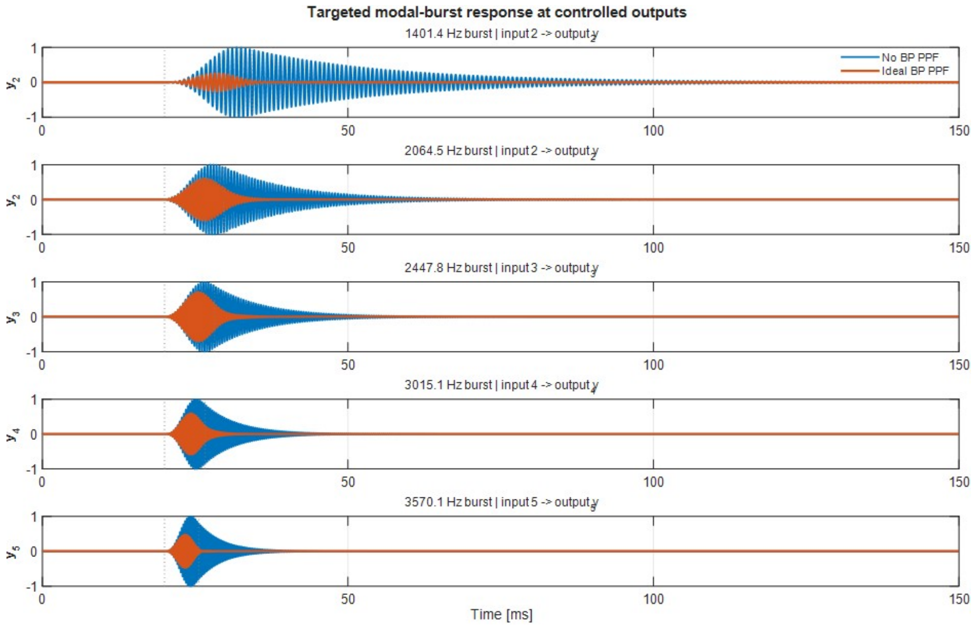

Figure 8.5: Comparison of the targeted modal-burst responses with and without band-pass PPF control.

9 Discussion

9.1. Main Findings

It was shown by the results of this thesis that the dynamic performance of a flexure-based motion stage cannot be improved by controller design alone. The mechanical layout, modal structure, actuator and sensor locations, and the construction of the measured outputs were all found to strongly influence the final control problem.
The initial mechanism design increased stiffness and improved the first natural frequency, while the geometry optimization was used to shape the modal behaviour of the stage in a more control-oriented way. The placement studies showed that actuator and sensor locations could be selected to improve the observability and controllability of specific modes.

However, for some higher-order and out-of-plane modes, single-point sensing and actuation were not sufficient to obtain clean modal separation. This motivated the use of multiple sensors and modal filtering. By combining several physical sensor signals into virtual modal outputs, the originally mixed frequency responses could be transformed into channels in which individual modes were more dominant.
The frequency-domain results showed that PPF and band-pass PPF could effectively reduce the targeted resonance peaks when the modes were sufficiently isolated. The time-domain benchmark confirmed the same conclusion from a disturbance-rejection perspective. Over-actuation and over-sensing were shown to be beneficial when the disturbance excited the targeted flexible modes. However, the benefit was found to depend strongly on the noise level of the additional over-sensing channels.

9.2. Benefit of Over-Actuation and Over-Sensing

The main benefit of over-actuation and over-sensing was that the control system was no longer limited to the information available in the main performance output. In a conventional single-input single-output approach, the controller mainly reacts to the motion measured in the performance direction. If a flexible mode is weakly visible in this signal, the ability to detect and suppress is limited.

By adding additional sensing channels, the flexible modes can be observed more directly. These additional outputs do not necessarily represent the final performance variable, but important modal information is contained in them. When these modal outputs are combined with additional actuation inputs, resonances that would otherwise remain poorly observable and controllable from the main degree of freedom can be targeted.

This interpretation was supported by the time-domain results. When a modal disturbance was injected in a way that excited the controlled modal channel, the modal

response was reduced significantly by the PPF loop.
Therefore, the value of over-actuation was not only that more control inputs were added, but that a path was created to control dynamic behaviour that could not be handled effectively by the main performance channel alone. In other words, the controllability and observability of modes that were poorly sensed and excited by the performance-direction actuator was improved by the extra actuators and sensors.

9.3. IMPORTANCE OF MODE ISOLATION

Mode isolation was found to be a central requirement for successful vibration control. The goal was to make the control problem simple by making the frequency response of the over-sensed channel behave like a simple mass-spring-damper system.
This was important because an active vibration control controller such as PPF is most effective when the controlled channel behaves approximately like a single resonant mode. In that case, the controller can be tuned around one resonance frequency, and damping is mainly added to the intended mode.

However, when several modes are strongly present in the same channel, non-targeted modes can also be affected by the controller. This can lead to spillover, phase problems, and unintended amplification in other channels.
This issue was observed clearly when the same optimization method for the yaw-mode was applied to the higher-order modes. The optimized physical sensor and actuator locations could not improve the response and could not isolate the target mode, because too many similar features were shared between the modes. There was not a single point where the actuator and sensor could be placed at the anti-node of the targeted mode and at the node of the other modes. As a result, the frequency response also contained contributions from other modes. On top of that, non-minimum-phase behaviour appeared. This showed that single-point optimal physical placement alone was not sufficient for higher-order modes.

This limitation was addressed by modal filtering, where more sensors were used to construct virtual outputs. Instead of relying on one measurement point, the mode shape itself was used to distinguish between different modes. This made it possible to transform a mixed physical response into several more mode-dominant channels.
However, it was also noticed that mode isolation through modal filtering could still be challenging when too many mode-shape features were shared. Therefore, other sensor configurations and combinations had to be tested to make the isolation as clean as possible.

9.4. INFLUENCE OF SENSOR NOISE

The time-domain benchmark showed that sensor noise is one of the main practical limitations of the over-actuated and over-sensed control architecture. In the ideal over-sensing case, the true modal output signals were received by the BP-PPF controller. Therefore, only the flexible mode of the plant was acted on. This represented the maximum achievable improvement of the controller.

However, in a realistic implementation, the over-sensing signals are contaminated by measurement noise. In this thesis, white sensor noise was added to the active over-sensing channels y_2, y_3, y_4, and y_5. These channels were used by the inner BP-PPF modal damping loops.

The main performance output x_1 was evaluated as a true mechanical output, so that the physical effect of the controller could be studied without being hidden by measurement noise on the performance sensor ($x_{1,\text{measured}}$).

The results showed that the realistic over-sensing case followed the ideal case when the over-sensor noise density was sufficiently low.
However, as the over-sensor noise density increased, the controller began to react not only to real structural vibration, but also to sensor noise. This noisy control action was then fed back into the plant in the form of actuation force.

Therefore, the benefit of over-sensing was limited by the quality of the additional sensing channels. If the noise was significant, the benefits that the over-sensing channels could potentially offer were reduced or even exacerbated. This means that over-sensing should not only be evaluated based on modal observability, but also based on whether the sensor noise is low enough for feedback control.

9.5. INTERPRETATION OF THE FREQUENCY-DOMAIN AND TIME-DOMAIN RESULTS

Complementary views of the same control problem were provided by the frequency-domain and time-domain analyses. In the frequency-domain analysis, it was shown that the BP-PPF controller could attenuate the targeted resonance peaks of the estimated state-space model. This analysis was useful for checking resonance reduction, spillover, and the effect of the controller on the selected modal channels.

The time-domain benchmark was then used to verify whether these frequency-domain improvements also appeared under representative operating conditions. Three cases were considered: tracking only, floor vibration only, and tracking combined with floor vibration.
In the tracking-only case, almost no difference was observed between the uncontrolled and controlled responses. This indicated that the selected triangular reference trajectory did not significantly excite the targeted flexible modes. Since the BP-PPF controller was designed to damp resonant modal motion, little improvement was expected when the error was dominated by low-frequency tracking components.

A different result was obtained in the floor vibration case. In this thesis, floor vibration was represented by equivalent modal force excitation at the selected resonance frequencies. Therefore, the flexible modes that the BP-PPF controller was designed to suppress were directly excited. As a result, the resonance peaks in the true x_1 error

ASD were significantly reduced by the ideal over-actuated controller. This confirmed the frequency-domain conclusion that the controller could attenuate the targeted modes when those modes were excited.

In the combined tracking and floor vibration case, both low-frequency tracking components and high-frequency modal components were present in the response. Little effect was obtained on the low-frequency part of the error, but the resonant peaks caused by the modal disturbance were reduced. Therefore, the BP-PPF controller should not be interpreted as a general tracking-error controller. Its main purpose was to reduce vibrations when the disturbance contained energy near the controlled modal frequencies.

Overall, it was shown by the frequency-domain analysis that the controller had the ability to damp the selected modes, while it was shown by the time-domain analysis under which excitation and noise conditions this damping led to practical improvement.

9.6. Practical Limitations

Although the time-domain benchmark confirmed the usefulness of the BP-PPF controller for modal disturbance rejection, several practical limitations should be considered.
First, the floor vibration disturbance used in the benchmark should be interpreted as a modal disturbance stress test rather than as a direct representation of a measured broadband floor-vibration spectrum. The assumed base displacement was converted into equivalent inertial forces at the selected modal frequencies. Since the equivalent force scales with f_i^2, the resulting excitation at kHz modal frequencies can be large. This is useful for testing whether the controller can suppress resonant modal vibration, but it does not represent real floor-vibration conditions.

Second, the results depend on the quality of the estimated state-space model. The time-domain benchmark was performed using the estimated reduced MIMO model rather than the full finite-element model or an experimental plant. Therefore, the conclusions are valid within the accuracy of the identified model.

Finally, implementation complexity is increased by the over-sensed architecture. Each additional modal sensing channel requires suitable sensor fusion for the modal filtering and sufficiently low noise, otherwise noise may be injected by the additional feedback loops, and it could reduce the improvements or even exacerbate it.

9.7. Future Design Recommendations

Future designs should continue to treat the mechanical design, sensing layout, and control architecture as one combined problem. Increasing stiffness and raising the first natural frequency are important, but they are not sufficient by themselves. The structure should also be designed so that the important flexible modes can be observed and controlled through well-separated modal channels.

A key recommendation is that modal isolation should be included as an explicit design objective. It was shown by the results that BP-PPF was most effective when each control channel was dominated by one targeted resonance. If several modes are strongly mixed in the same channel, the controller becomes more difficult to tune and spillover may be introduced.

Therefore, future mechanism and sensor-layout designs should be evaluated not only by their natural frequencies, but also by how clearly the relevant modes appear in the chosen sensing channels. This was already done in this thesis for a single yaw mode and was shown to be effective. Further investigation is required to evaluate whether this approach could be extended to more modes.

Furthermore, sensor noise should be included early in the design process. It was shown by the time-domain benchmark that over-sensing was only beneficial when the additional sensing channels were sufficiently clean. Therefore, the required noise density should be estimated before sensors and electronics are selected.

10 Conclusion and Future Work

10.1. Conclusion

The goal of this thesis was to investigate whether the dynamic performance of a positioning stage could be improved by combining control-oriented mechanical design, over-actuation, over-sensing, modal filtering, and active vibration control. The work started from the mechanical design of the stage and gradually developed toward a MIMO control architecture, in which additional sensing and actuation channels were used to observe and suppress flexible modes more effectively.

It was shown that the performance of a motions stage is not determined only by the controller. The mechanical layout, modal structure, sensor placement, actuator placement, and measurement quality were all found to influence the final control problem. Therefore, good modal observability and controllability should be included throughout the design process.

The structural optimization showed that the geometry of the mechanism could be shaped in a control-oriented way. By optimizing the design variables of the stage and intermediate structure, the first natural frequency was increased, while parasitic pitch and roll behaviour were shifted to higher frequencies. This created a better starting point for later sensor placement and vibration control.

The sensor and actuator placement studies showed that the location of sensors and actuators had a strong influence on the ability to observe and control individual modes. For the well-isolated mode that resulted from the optimization, namely the yaw mode, a physical placement strategy based on strain and modal response was sufficient. However, for the non-optimized higher-order modes, single-point sensing and actuation were not enough to obtain clean modal separation. These modes shared similar deformation patterns, which resulted in mixed frequency responses and non-minimum-phase behaviour.

To address this issue, modal filtering was introduced. By combining multiple physical sensor signals into virtual modal outputs, the mixed responses were transformed into channels in which the targeted modes were more dominant. This made the control problem easier and allowed simple active vibration controllers, such as PPF and band-pass PPF, to be used effectively. The frequency-domain analysis showed that the targeted resonance peaks could be attenuated with limited spillover when the channels were sufficiently isolated.

The time-domain benchmark confirmed the main frequency-domain findings. When the disturbance excited the targeted flexible modes, the BP-PPF controller reduced the resonance peaks and increased the effective damping of the controlled modal outputs. However, the tracking-only case showed little improvement because the reference trajectory did not significantly excite the targeted flexible modes. This confirmed that the proposed approach was mainly useful for suppressing resonant vibration, rather than directly improving low-frequency tracking error.

A final conclusion is that the practical benefit of over-sensing depends strongly on sensor noise. In the realistic case, noise in the over-sensing channels was fed back through the BP-PPF controller and entered the plant as actuator force. If this noise was sufficiently low, the over-actuated system remained beneficial. If the noise was too high, the improvement was reduced. Therefore, over-actuation and over-sensing were only practically useful when the added sensing channels were both modal enough and clean enough.

10.2. CONTRIBUTIONS

The main contributions of this thesis are summarized as follows.

First, a control-oriented design workflow was developed. Instead of optimizing the mechanism only for stiffness or displacement range, the geometry was optimized with the later vibration-control problem in mind. This included increasing the first natural frequency, reducing parasitic pitch and roll behaviour, and shaping the modal structure so that the targeted mode, namely the yaw mode, could be more easily observed and controlled.

Second, it was shown how actuator and sensor placement could be optimized. The COMSOL approach provided a practical and fast placement method, while the Gramian and Hankel singular value analysis provided a more rigorous check of the selected locations.

Third, the limitation of single-point physical sensing was demonstrated for modes that shared similar features. For some modes, even optimized physical sensor and actuator locations could not fully isolate the desired mode. This led to mixed channels and undesirable non-minimum-phase behaviour.

Fourth, a modal filtering approach was used to construct virtual modal outputs from multiple sensor signals. This allowed the control architecture to move from physical channels with mixed modal content toward more mode-dominant virtual channels. This was an important contribution because the MIMO control problem was made easier and simple active vibration controllers were made more effective.

Fifth, PPF and band-pass PPF controllers were applied to the mode-dominant channels and evaluated in the frequency domain. It was shown that the targeted resonances could be attenuated when the corresponding channel had sufficient modal isolation.

Sixth, a time-domain benchmark was used to evaluate the controlled system under tracking, modal disturbance excitation, and realistic over-sensor noise. It was shown by this benchmark that the proposed over-actuated architecture was beneficial when the targeted resonant modes were excited, but that its practical performance depended strongly on the noise level of the over-sensing channels.

10.3. Limitations

The main limitation of this work is that the results are simulation-based. The mechanism, modal filtering, controller design, and time-domain benchmark were evaluated using finite-element and estimated state-space models, but no experimental validation was performed. Therefore, the conclusions are valid within the accuracy of the models that were used.

Actuator dynamics were also not explicitly modelled. The control input was applied directly to the plant, while a real implementation would include actuator behaviour between the controller output and the mechanical input. These dynamics could limit the achievable damping. Therefore, the obtained attenuation should be interpreted as an ideal-actuator result. Another limitation is that the practical sensor implementation was not fully investigated. The sensor locations and virtual outputs were studied from a modelling and control point of view, but the final sensor type, mounting method, calibration, and noise performance were not experimentally validated.

10.4. Future Work

Future work should first focus on experimental validation using a reduced version of the architecture. A practical first step would be to select one or two well-isolated modal channels, measure the open-loop frequency responses, and compare them with the simulated model.

The modal filtering procedure should then be validated experimentally by checking whether the virtual outputs are dominated by the intended modes. After the modal outputs have been verified, BP-PPF loops can be implemented and tested. Additional damping loops should only be added step by step while stability and spillover are monitored. Future work should also include actuator dynamics in the model. This would make the simulation more realistic and would allow the effect of actuator bandwidth, phase lag, and force limits to be evaluated.

Future designs could also benefit from including multi-mode isolation as a design criterion from the beginning. In this way, the mechanical design, sensor placement, and control architecture could be developed together instead of being treated as separate steps.

10.5. Final Remarks

It was shown in this thesis that over-actuation and over-sensing can improve the vibration control of a motion stage when the additional channels provide clean access to the targeted flexible modes. The main benefit was found to be the ability to construct mode-dominant channels, which isolated the relevant modes and simplified the control problem.

Although experimental validation is still required, the simulation results showed that the proposed approach is a promising strategy for suppressing targeted resonant modes in precision positioning systems.

Bibliography

[1] MIKRONIEK, "PROFESSIONAL JOURNAL ON PRECISION ENGINEERING," 2020. [Online]. Available: https://www.dspe.nl/wp-content/uploads/2020/03/Mikroniek-2020-2-low-res.pdf.

[2] *Enhancing Semiconductor Manufacturing with Vibration Control*. [Online]. Available: https://www.azom.com/article.aspx?ArticleID=23783.

[3] "Data-based spatial feedforward for over-actuated motion systems," DOI: 10.1016/j.mechatronics.2014.07.002. [Online]. Available: https://doi.org/10.1016/j.mechatronics.2014.07.002.

[4] C. Collette, S. Janssens, and K. Artoos, "Review of Active Vibration Isolation Strategies," *Recent Patents on Mechanical Engineering*, vol. 4, pp. 0–000, 2011.

[5] T. Van Der Graaf, "Active Vibration Control Using over-sensing and over-actuation," 2020. [Online]. Available: http://repository.tudelft.nl/..

[6] J. Schneiders et al., "Integrated design of a lightweight positioning system," pp. 1–10, 2004. [Online]. Available: www.tue.nl/taverne.

[7] *(PDF) Modal Framework for Closed-Loop Analysis of Over-Actuated Motion Systems*. [Online]. Available: https://www.researchgate.net/publication/228945241_Modal_Framework_for_Closed-Loop_Analysis_of_Over-Actuated_Motion_Systems.

[8] V. Gupta, M. Sharma, and N. Thakur, "Optimization criteria for optimal placement of piezoelectric sensors and actuators on a smart structure: A technical review," *Journal of Intelligent Material Systems and Structures*, vol. 21, no. 12, pp. 1227–1243, Aug. 2010, ISSN: 1045389X. DOI: 10.1177/1045389X10381659. [Online]. Available: http://www.sagepub.co.uk/journalsPermissions.nav.

[9] Z. c. Qiu, X. m. Zhang, H. x. Wu, and H. h. Zhang, "Optimal placement and active vibration control for piezoelectric smart flexible cantilever plate," *Journal of Sound and Vibration*, vol. 301, no. 3-5, pp. 521–543, Apr. 2007, ISSN: 0022-460X. DOI: 10.1016/J.JSV.2006.10.018.

[10] *P-876 DuraAct Patch Transducer*. [Online]. Available: https://www.piceramic.com/en/products/piezoceramic-actuators/patch-transducers/p-876-duraact-patch-transducer-101790.

[11] N. Kumar and S. P. Singh, "Vibration and damping characteristics of beams with active constrained layer treatments under parametric variations," *Materials & Design*, vol. 30, no. 10, pp. 4162–4174, Dec. 2009, ISSN: 0261-3069. DOI: 10.1016/J.MATDES.2009.04.044. [Online]. Available: https://www-sciencedirect-com.tudelft.idm.oclc.org/science/article/pii/S0261306909001836.

[12] E. F. Crawley and E. H. Anderson, "Detailed Models of Piezoceramic Actuation of Beams," *Journal of Intelligent Material Systems and Structures*, vol. 1, no. 1, pp. 4–25, 1990, ISSN: 15308138. DOI: 10.1177/1045389X9000100102. [Online]. Available: /doi/pdf/10.1177/1045389X9000100102?download=true.

[13] E. Balmes and A. Deraemaeker, "Modelling structures with piezoelectric materials : full, reduced order, and state-space models Theory and SDT Tutorial," 2001. [Online]. Available: http://www.sdtools.com.

[14] H. Tang, Y. Li, and J. Huang, "Design and analysis of a dual-mode driven parallel XY micromanipulator for micro/nanomanipulations," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 226, no. 12, pp. 3043–3057, Dec. 2012, ISSN: 09544062. DOI: 10.1177/0954406212442272;JOURNAL:JOURNAL:PICB;WEBSITE:WEBSITE:SAGE;REQUESTEDJOURNAL:JOURNAL:PICB;WGROUP:STRING:PUBLICATION. [Online]. Available: https://journals-sagepub-com.tudelft.idm.oclc.org/doi/10.1177/0954406212442272.

[15] T. K. Caughey, "On the stability problem caused by finite actuator dynamics in the collocated control of large space structures," *International Journal of Control*, vol. 41, no. 3, pp. 787–802, 1985, ISSN: 13665820. DOI: 10.1080/00207178508961163;WEBSITE:WEBSITE:TFOPB;PAGEGROUP:STRING:PUBLICATION. [Online]. Available: https://www-tandfonline-com.tudelft.idm.oclc.org/doi/pdf/10.1080/00207178508961163.

[16] M. I. Friswell and D. J. Inman, "The relationship between positive position feedback and output feedback controllers," *Smart Mater. Struct*, vol. 8, pp. 285–291, 1999.

[17] A. M. Natu and H. HosseinNia, "Narrow-Band Active Damping Using Higher-Order Band-Pass Control: A Loop-Shaping Perspective," 2025. DOI: 10.2139/SSRN.5947393. [Online]. Available: https://papers.ssrn.com/abstract=5947393.

[18] J. Sirohi and I. Chopra, "Fundamental Understanding of Piezoelectric Strain Sensors," 2001. DOI: 10.1106/8BFB-GC8P-XQ47-YCQ0. [Online]. Available: http://jim.sagepub.com.

[19] "Signal conditioning for piezoelectric sensors," [Online]. Available: www.ti.com/aaj.

[20] H. Sun, Y. Wang, X. Qing, and Z. Wu, "Optimal Bonding Adhesive Design," DOI: 10.3390/s18082554. [Online]. Available: www.mdpi.com/journal/sensors.

[21] *Material parameters of PIC 255. | Download Table.* [Online]. Available: https://www.researchgate.net/figure/Material-parameters-of-PIC-255_tbl2_259230226.

[22] *Projection matrix - Wikipedia.* [Online]. Available: https://en.wikipedia.org/wiki/Projection_matrix.

A Leaf-Spring Thickness and Patch Strain Relation

The derivation is based on an Euler–Bernoulli beam assumption, where the patch-covered part of the leaf spring is assumed to bend with approximately constant curvature. For small bending deformation, the normal strain at a distance z from the neutral axis is

$$\varepsilon = -\kappa z$$

where κ is the bending curvature. A relation found in literature gives the strain in the beam as

$$\varepsilon = -\kappa z = -\frac{6\left(1+\frac{1}{T}\right)\left(\frac{2}{t_s}\right)\Lambda}{(6+\Psi_b)+\frac{12}{T}+\frac{8}{T^2}}z \qquad [12]$$

with

$$T = \frac{t_s}{t_a}, \quad z = \frac{t_s}{2}$$

and

$$\Psi_b = \Psi_e = \frac{(EA)_s}{(EA)_a} = \frac{E_s b_s t_s}{E_a b_a t_a}.$$

Here, t_s is the leaf-spring thickness, t_a is the actuator thickness in the literature formulation, E_s and E_a are the Young's moduli of the substrate and actuator, and b_s and b_a are their widths.

Actuator authority:

$$\Lambda = N_{\text{pairs}} \frac{Y_p}{E_s}\frac{B_p}{b}\Lambda_1 \frac{2y_p}{t_s},$$

where

$$\Lambda_1 = d_{31}\frac{V}{t_p}.$$

In this expression, N_{pairs} is the number of active piezo patch pairs, Y_p is the piezo Young's modulus, E_s is the leaf-spring Young's modulus, B_p is the patch width, b is the leaf-spring width, d_{31} is the piezoelectric strain coefficient, V is the applied voltage, and t_p is the piezo patch thickness.

The distance from the leaf-spring neutral axis to the centre of the piezo patch is

$$y_p = \frac{t_s}{2} + t_a + \frac{t_p}{2},$$

where t_a is now the adhesive layer thickness. The factor

$$\frac{2y_p}{t_s}$$

acts as a dimensionless lever arm. A larger distance between the piezo patch and the neutral axis increases the bending authority of the patch.

Using the above relations, the mechanical strain can be evaluated for different values of thickness which was used to obtain the thickness-strain trend shown in Figure 2.1.

B SNR-Based Strain Requirement

The relation between SNR and strain can be written as follows.
Charge generated by piezoelectric patch mechanical strain is:

$$Q = \eta d_{31} Y_p A_p \varepsilon_{p,\mathrm{mech}} \qquad [18]$$

where Q is the generated charge, η is an effectiveness factor, d_{31} is the piezoelectric coefficient, Y_p is the Young's modulus of the piezo material, A_p is the effective electrode area, and $\varepsilon_{p,\mathrm{mech}}$ is the mechanical strain in the piezo patch.

The charge converts to voltage according to

$$V_{\mathrm{sig}} = \frac{Q}{C_f},$$

where C_f is the feedback capacitance. Substituting the generated charge gives

$$V_{\mathrm{sig}} = \frac{\eta d_{31} Y_p A_p \varepsilon_{p,\mathrm{mech}}}{C_f}.$$

The signal-to-noise ratio is defined as:

$$\mathrm{SNR}_{\mathrm{lin}} = \frac{V_{\mathrm{sig}}}{V_n},$$

where V_n is the voltage noise floor. Therefore,

$$\mathrm{SNR}_{\mathrm{lin}} = \frac{\eta d_{31} Y_p A_p \varepsilon_{p,\mathrm{mech}}}{C_f V_n}.$$

And minimum detectable strain:

$$\varepsilon_{\min} = \frac{\mathrm{SNR}_{\mathrm{lin}} V_n C_f}{\eta |d_{31}| Y_p A_p}.$$

Required decibels is converted to linear amplitude ratio:

$$\mathrm{SNR}_{\mathrm{lin}} = 10^{\mathrm{SNR}_{\mathrm{dB}}/20}.$$

Then the minimum strain:

$$\varepsilon_{\min} = \frac{10^{\mathrm{SNR}_{\mathrm{dB}}/20} V_n C_f}{\eta |d_{31}| Y_p A_p}.$$

Parameter	Symbol	Value	Unit / note
Noise voltage density	V_n	434[19]	$\mathrm{nV}/\sqrt{\mathrm{Hz}}$
Feedback capacitance	C_f	680[19]	pF
Feedback resistance	R_f	270[19]	MΩ
Target signal-to-noise ratio	$\mathrm{SNR}_{\mathrm{target}}$	30	dB
Effectiveness factor	η	0.6[20]	-
Piezoelectric coefficient	d_{31}	-180×10^{-12}[13]	m/V
Effective piezo area	A_p	10×10[10]	mm^2
Piezo Young's modulus	Y_p	63[21]	GPa
Minimum detectable strain density	$\varepsilon_{\mathrm{min}}$	$1.37 \times 10^{-5} \times \sqrt{\mathrm{BW}}$	$\mu\varepsilon/\sqrt{\mathrm{Hz}}$
Example bandwidth	BW	5	Hz
Minimum detectable strain at 5 Hz	$\varepsilon_{\mathrm{min}}$	3.064×10^{-5}	$\mu\varepsilon$

Table B.1: Parameters used for the minimum detectable strain calculation.

C Optimization objective

The optimization objective was defined to favor a high sensor response at the desired mode around 1401 Hz, while reducing the response at the undesired mode around 141 Hz. Therefore, the following objective was maximized:

$$J_2 = \frac{G_{1401}}{G_{141} + \epsilon_J} \tag{C.1}$$

where

$$G_{141} = \left|\mathrm{SensorDisp}(141\,\mathrm{Hz})\right|^2 \tag{C.2}$$

and

$$G_{1401} = \left|\mathrm{SensorDisp}(1401\,\mathrm{Hz})\right|^2. \tag{C.3}$$

Here, $\epsilon_J = 10^{-12}$ was used to avoid division by zero.
In COMSOL, the global variable was implemented as:

```
epsJ 1e-12 ""

is141 "if(abs(freq-141[Hz]) < 50[Hz], 1, 0)" ""
is1401 "if(abs(freq-1401[Hz]) < 50[Hz], 1, 0)" ""

Gbad "is141*abs(comp1.SensorDisp)^2" ""
Ggood "is1401*abs(comp1.SensorDisp)^2" ""

J2 "G1401/(G141 + epsJ)" ""

G141 "withsol('sol29', abs(comp1.SensorDisp)^2, setval(xa,xa), setval(xs,xs),
↪ setind(freq,1))" ""
G1401 "withsol('sol29', abs(comp1.SensorDisp)^2, setval(xa,xa), setval(xs,xs),
↪ setind(freq,2))" ""
```

The optimization was performed by maximizing J_2.

D Sensor Placement Optimization Code

D.1. Main DoF Sensor Placement Code

```
    clc; clear; close all;

%% Sensor locations

% 8 sensor points
xS = -24.37204675 * ones(8,1);

yS = [-93.391;
      -42.000;
      -54.848;
      -74.119;
      -80.543;
      -70.907;
      -67.695;
      -72.513];

zS = 23.2895 * ones(8,1);

Ns = numel(xS);

%% COMSOL frequency-response extraction

model = mphopen('V9.1.13.mph');

stdTag = 'std2';
dsetFR = 'dset8';
F_in   = 100;      % input force [N]

model.study(stdTag).run;

% Frequency vector
f = mphglobal(model, 'freq', 'dataset', dsetFR, 'solnum', 'all');
f = f(:);
Nf = numel(f);

% Pull displacement at all 8 sensor points at once
coordS = [xS.'; yS.'; zS.'];   % 3 x Ns

out = mphinterp(model, 'u', ...
    'coord', coordS, ...
    'dataset', dsetFR, ...
    'solnum', 'all', ...
    'unit', 'm');

out = squeeze(out);

% Ensure output is arranged as Nf x Ns
if size(out,1) == Ns && size(out,2) == Nf
    out = out.';
elseif size(out,1) == Nf && size(out,2) == Ns
    % already correct
else
```

```
    error('Unexpected mphinterp output size.');
end

%% Build FRD model

% Build FRD response array H: Ny x Nu x Nf = 8 x 1 x Nf
H = zeros(Ns, 1, Nf);

for k = 1:Nf
    H(:,1,k) = out(k,:).' / F_in;   % displacement/force [m/N]
end

sysFRD = frd(H, f, 'FrequencyUnit', 'Hz');

%% Plot FRD

opts = bodeoptions;
opts.MagUnits   = 'dB';
opts.PhaseUnits = 'deg';
opts.FreqUnits  = 'Hz';
opts.Grid       = 'on';

figure;
bodeplot(sysFRD, opts);
title('COMSOL FRD: 8 sensors, 1 actuator');

%% State-space estimation

w = 2*pi*f;                    % rad/s
dataID = idfrd(H, w, 0);        % frequency-domain identification data

nx = 17;                       % selected state-space order

optID = ssestOptions;
optID.EnforceStability = true;
optID.Display = 'on';

sysID = ssest(dataID, nx, optID);
sysSS = ss(sysID);

%% Sensor scoring settings

modesHz = 141.05;
bwFrac  = 0.08;

ScoreWo  = zeros(numel(modesHz), Ns);
ScoreHSV = zeros(numel(modesHz), Ns);

%% Frequency-limited Gramian scoring

for k = 1:numel(modesHz)

    f0 = modesHz(k);
    f1 = f0 * (1 - bwFrac);
    f2 = f0 * (1 + bwFrac);

    % Frequency interval is given in rad/s
    optG = gramOptions('FreqIntervals', [2*pi*f1, 2*pi*f2]);

    % Controllability Gramian of the full system
    Wc = gram(sysSS, 'c', optG);

    for i = 1:Ns

        % Single sensor output
        sys_i = sysSS(i,1);

        % Observability Gramian for sensor i
```

```
        Wo = gram(sys_i, 'o', optG);

        ScoreWo(k,i) = trace(Wo);

        % Joint controllability-observability score
        eigProd = eig(Wc * Wo);
        eigProd = max(real(eigProd), 0);

        hsv = sqrt(eigProd);
        ScoreHSV(k,i) = sum(sort(hsv, 'descend'));

    end
end

%% Ranking table

T = table((1:Ns).', xS, yS, zS, ...
    ScoreHSV(1,:).', ScoreWo(1,:).', ...
    'VariableNames', {'Sensor','x','y','z','HSV_141Hz','Wo_141Hz'});

T141 = sortrows(T, 'HSV_141Hz', 'descend');

disp('Top sensors for 141 Hz:');
disp(T141);
```


D.2. Yaw-Mode Sensor Placement Code

```
    clc; clear; close all;

%% ========== USER SETTINGS ==========
csvFile  = "sweep2D_v3.csv";     % <--- your COMSOL exported table
modelFile = "V9.1.13_mimo_distributed_nonrigid.mph";

stdTag = "std4";              % freq-domain study tag
dsetFR = "dset17";             % dataset tag for the freq response
F_in   = -100;               % N


probeExpr = "comp1.SensorDisp";


% Identification settings
nx = 8;
optID = ssestOptions;
optID.EnforceStability = true;
optID.Display = "on";


K = 144;                       % top-K from COMSOL sweep

% Mode bands (Hz)
modesHz = [141.05 1401.4];
bwFrac  = 0.08;               %  8 % band

%% ==========  READ SWEEP CSV ==========
T = readtable(csvFile);
vn = lower(string(T.Properties.VariableNames));

% auto-detect columns
c_xa = find(contains(vn,"xa"),1);
c_xs = find(contains(vn,"xs"),1);
c_J  = find(contains(vn,"j2"),1);
xa_all = T{:,c_xa};  xa_all = xa_all(:);
xs_all = T{:,c_xs};  xs_all = xs_all(:);
J_all  = T{:,c_J};   J_all  = J_all(:);
```

```matlab



%% ========== UNIQUE (xa,xs) AND KEEP BEST J PER PAIR ==========
tol = 1e-9;   % meters;
xa_q = round(xa_all./tol).*tol;
xs_q = round(xs_all./tol).*tol;

% Build a table and group by (xa,xs)
Tall = table(xa_q, xs_q, J_all, 'VariableNames', {'xa','xs','J2_COMSOL'});

% For duplicates, keep the maximum J2_COMSOL for each (xa,xs)
Tuniq = groupsummary(Tall, {'xa','xs'}, 'max', 'J2_COMSOL');
Tuniq = renamevars(Tuniq, "max_J2_COMSOL", "J2_COMSOL");
Tuniq = removevars(Tuniq, "GroupCount");

%% ========== SORT AND TAKE TOP K ==========
Tuniq = sortrows(Tuniq, "J2_COMSOL", "descend");
Cand  = Tuniq(1:min(K,height(Tuniq)), :);

disp("Top UNIQUE candidates from COMSOL sweep:");
disp(Cand);
%% ==========  OPEN COMSOL MODEL ==========
model = mphopen(modelFile);

%% ==========  LOOP CANDIDATES: FRD -> SS -> band-Gramians -> HSV ==========
out = Cand;
out.HSV_141    = nan(height(out),1);
out.HSV_1401   = nan(height(out),1);
out.trWc_141   = nan(height(out),1);
out.trWo_141   = nan(height(out),1);
out.trWc_1401 = nan(height(out),1);
out.trWo_1401 = nan(height(out),1);


%%
for r = 1:height(out)




    xa = out.xa(r);
    xs = out.xs(r);
    zs = 23.2895;


    ya = model.param.evaluate("ya_p")*1e3;   % mm
    ys = model.param.evaluate("ys_p")*1e3;   % mm
    model.geom("geom1").run;
    model.mesh("mesh1").run;
    model.study(stdTag).run;

    f = mphglobal(model, 'freq', 'dataset', dsetFR, 'solnum', 'all');

    f = f(:);
    w = 2*pi*f;  % rad/s

    % Evaluate rotation proxy at the sensor point across frequency
    coordS = [xs; ys; zs];              % 3x1, mm
    y = mphinterp(model, probeExpr, "coord", coordS, "dataset", dsetFR, "solnum", "all",'unit','m');
    y = y(:);                           % Nf x 1 complex


    % Transfer function H = y / F
    H = y / F_in;

    % Build FRD and identify SS
    sysFRD = frd(H, f, 'FrequencyUnit','Hz');
```

```
    dataID = idfrd(H, w, 0);          % continuous-time FRF
    sysID  = ssest(dataID, nx, optID);
    sysSS  = ss(sysID);

    % Compute band-limited scores for each mode
    for k = 1:2
        f0 = modesHz(k);
        f1 = f0*(1-bwFrac);
        f2 = f0*(1+bwFrac);

        optG = gramOptions("FreqIntervals", [2*pi*f1 2*pi*f2]);

        Wc = gram(sysSS, "c", optG);
        Wo = gram(sysSS, "o", optG);

        hsv = sqrt(max(real(eig(Wc*Wo)),0));
        hsvScore = sum(sort(hsv,"descend"));

        if k==1
            out.HSV_141(r)  = hsvScore;
            out.trWc_141(r) = trace(Wc);
            out.trWo_141(r) = trace(Wo);
        else
            out.HSV_1401(r)  = hsvScore;
            out.trWc_1401(r) = trace(Wc);
            out.trWo_1401(r) = trace(Wo);
        end
    end

    fprintf("Done %d/%d: xa=%g, xs=%g\n", r, height(out), xa, xs);
end

%% ==========  RANKING / COMPARISON ==========
disp("Ranking by HSV around 1401 Hz (best first):");
disp(sortrows(out, "HSV_1401", "descend"));

%%

% Sort
rank1401 = sortrows(out, "HSV_1401", "descend");

% Write to Excel
outXlsx = "ranking_HSV1401.xlsx";
writetable(rank1401, outXlsx, "Sheet", "HSV_1401", "WriteMode", "overwrite");

fprintf("Saved: %s\n", outXlsx);
%%
opts = bodeoptions;
opts.MagUnits='dB';
opts.PhaseUnits='deg';
opts.FreqUnits='Hz';
opts.Grid='on';
opts.PhaseMatching = 'on';
opts.XLim = {[1, 2000]};


%% ==========  VISUALIZATION ==========

useCombined = true;
w141 = 0.5;

if useCombined
    score = out.HSV_1401 + w141*out.HSV_141;
    scoreName = sprintf("HSV_1401 + %.2g*HSV_141", w141);
else
    score = out.HSV_1401;
    scoreName = "HSV_1401";
end
```

```
[bestScore, bestIdx] = max(score);
xaBest = out.xa(bestIdx);
xsBest = out.xs(bestIdx);

% --- Plot: show ALL near-optimal solutions
rtol = 0.01;  % within 1% of best
idxNear = score >= bestScore*(1-rtol);

figure;
scatter(out.xa, out.xs, 25, score, "filled"); colorbar;
hold on;
scatter(out.xa(idxNear), out.xs(idxNear), 70, "o", "LineWidth", 1.5);
plot(xaBest, xsBest, "rp", "MarkerSize", 14, "MarkerFaceColor", "r");
grid on;
xlabel("xa [mm]");
ylabel("xs [mm]");
title(sprintf("Near-optima (within %.1f%% of best) + best point", 100*rtol));
legend("All candidates", "Near-optima set", "Best", "Location", "best");
```

E Modal Filtering

E.1. Mathematical Derivation

The following is the sensor response vector:

$$\mathbf{y} = [y_1 \quad y_2 \quad \cdots \quad y_n]^T,$$

where n is the number of sensing points. For each mode the collected sensor data was collected in a vector $\boldsymbol{\phi}_m$. The aim was to construct a weighting vector $\mathbf{q}_m$, such that the weighted sensor signal

$$q_m = \mathbf{q}_m^T \mathbf{y}$$

mainly responded to the target mode m, while the contribution of selected unwanted modes was rejected.

To achieve this, first the unwanted modes are collected in the nuisance matrix

$$\mathbf{B}_m = [\boldsymbol{\phi}_{i_1} \quad \boldsymbol{\phi}_{i_2} \quad \cdots \quad \boldsymbol{\phi}_{i_r}],$$

where the columns represent the sampled mode shapes that should be rejected. A projection matrix onto this nuisance subspace is done as following:

$$\mathbf{P}_m = \mathbf{B}_m(\mathbf{B}_m^T \mathbf{B}_m)^{-1} \mathbf{B}_m^T. \qquad [22]$$

The part of the target mode shape that lies inside the nuisance subspace was removed by applying

$$\tilde{\mathbf{q}}_m = (\mathbf{I} - \mathbf{P}_m)\boldsymbol{\phi}_m.$$

This resulted in a vector that was orthogonal to the selected nuisance modes, while still retaining the part of the target mode that could be observed by the sensor array. Finally, the vector was normalized such that the virtual channel had unit response to the target mode:

$$\mathbf{q}_m = \frac{\tilde{\mathbf{q}}_m}{\tilde{\mathbf{q}}_m^T \boldsymbol{\phi}_m}.$$

The resulting virtual output was then obtained as

$$q_m = \mathbf{q}_m^T \mathbf{y}.$$

E.2. Implementation

```
    clc; clear; close all;
%%
model  = mphopen("V9.1.13_MIMO_sys_two_sensors_third_mode_with_projection_10_sensors_7x7_6th.mph");
stdTag = "std4";
dsetFR = "dset17";

```

```
F_in = 1;

xs1 = -24.37204675;
ys1 = -72.513;
zs1 = 23.2895;

xs2 = -78.57704675;
ys2 = -102;
zs2 =  23.2895;

xs3 = -24.37204675;
ys3 = -102;
zs3 =  23.2895;

xs4 = -24.37204675;
ys4 = -72;
zs4 = 26.00345;

xr1 = -24.37204675;
yr1 = -102;
zr1 = 26.00345;

xr2 = -24.37204675;
yr2 = -87;
zr2 = 26.00345;

xr3 = -24.37204675;
yr3 = -72;
zr3 = 26.00345;

xr4 = -24.37204675;
yr4 = -57;
zr4 = 26.00345;

xr5 = -24.37204675;
yr5 = -42;
zr5 = 26.00345;

xl1 = -78.57704675;  yl1 = -102;  zl1 = 26.00345;
xl2 = -78.57704675;  yl2 =  -87;  zl2 = 26.00345;
xl3 = -78.57704675;  yl3 =  -72;  zl3 = 26.00345;
xl4 = -78.57704675;  yl4 =  -57;  zl4 = 26.00345;
xl5 = -78.57704675;  yl5 =  -42;  zl5 = 26.00345;

xleaf1 = -38.09454675;
yleaf1 = -21;
zleaf1 = 38.579;

xleaf2 = -38.09454675;
yleaf2 = -39.5;
zleaf2 = 46.579;
%%
sens1 = struct("x", xs1, "y", ys1, "z", zs1);   % main DoF sensor point
sens2 = struct("x", xs2, "y", ys2, "z", zs2);   % rotation sensor point
sens3 = struct("x", xs3, "y", ys3, "z", zs3);

expr_y1 = "comp1.SensorDisp";
expr_y2 = "comp1.point2";
expr_y3 = "comp1.point3";

% Frequency vector (Hz)
f = mphglobal(model,"freq","dataset",dsetFR,"solnum","all"); f = f(:);
w = 2*pi*f;
Nf = numel(f);

Ny = 7; Nu = 7;
H  = zeros(Ny,Nu,Nf);   % Ny x Nu x Nf
```

```
coordS1 = [sens1.x; sens1.y; sens1.z];
coordS2 = [sens2.x; sens2.y; sens2.z];
coordS3 = [sens3.x; sens3.y; sens3.z];
%%
expr_y4raw = "w";

coordA = [xr1 yr1 zr1;
          xr2 yr2 zr2;
          xr3 yr3 zr3;
          xr4 yr4 zr4;
          xr5 yr5 zr5];

coordB = [xl1 yl1 zl1;
          xl2 yl2 zl2;
          xl3 yl3 zl3;
          xl4 yl4 zl4;
          xl5 yl5 zl5];

coordC1 = [xleaf1; yleaf1; zleaf1];
coordC2 = [xleaf2; yleaf2; zleaf2];

coord10 = [coordA; coordB];

% ===== eigenmode dataset and mode numbers =====
dsetEig    = "dset6";
mode1 = 3;     % first peak
mode2    = 11;    % second peak
mode3 = 12;    % third peak
mode4 = 21;
mode5 = 22;
mode6 = 13;
mode7 = 1;
mode8 = 2;
mode9 = 4;
mode10 = 5;
mode11 = 6;
mode12 = 7;
mode13 = 8;
mode14 = 9 ;
mode15 = 10;
mode16 = 14;
mode17 = 15;
mode18 = 16;
mode19 = 17 ;
mode20 = 18;
mode21 = 19 ;
mode22 = 20 ;
mode23 = 23;

keepModes  = [mode1 mode2 mode3 mode4 mode5];
extraModes = [mode6 mode7 mode8 mode9 mode10 mode11 mode12 mode13 mode14 mode15 mode16 mode17 mode18
     mode19 mode20 mode21 mode22 mode23];
%%
nSens = 10;

allModes = [keepModes extraModes];
nAll     = numel(allModes);

Phi10  = zeros(nSens, nAll);        % stage only
Phi11a = zeros(nSens+1, nAll);      % stage + leaf1
Phi11b = zeros(nSens+1, nAll);      % stage + leaf2
Phi12  = zeros(nSens+2, nAll);      % stage + both leafs

for m = 1:nAll
    thisMode = allModes(m);

    phi10 = zeros(nSens,1);
    for ii = 1:nSens
```

E

```
        phi10(ii) = mphinterp(model, expr_y4raw, ...
            "coord", coord10(ii,:).', ...
            "dataset", dsetEig, ...
            "solnum", thisMode, ...
            "unit", "m");
    end

    phi_leaf1 = mphinterp(model, 'u', ...
        "coord", coordC1, ...
        "dataset", dsetEig, ...
        "solnum", thisMode, ...
        "unit", "m");

    phi_leaf2 = mphinterp(model, 'u', ...
        "coord", coordC2, ...
        "dataset", dsetEig, ...
        "solnum", thisMode, ...
        "unit", "m");

    Phi10(:,m)  = phi10;
    Phi11a(:,m) = [phi10; phi_leaf1];
    Phi11b(:,m) = [phi10; phi_leaf2];
    Phi12(:,m)  = [phi10; phi_leaf1; phi_leaf2];


end
%%

% Output 3
Phi_use = Phi12;
phi_keep = Phi_use(:,1);
B = [Phi_use(:,2),Phi_use(:,3), Phi_use(:,8), Phi_use(:,16), Phi_use(:,19), Phi_use(:,20)];
P = B * ((B' * B)\B');
q3 = (eye(size(Phi_use,1)) - P) * phi_keep;
q3 = q3 / (q3' * phi_keep);

%%

% Output 4
Phi_use = Phi10;
phi_keep = Phi_use(:,2);
B = [Phi_use(:,[1 3])];
P = B * ((B' * B)\B');
q4 = (eye(size(Phi_use,1)) - P) * phi_keep;
q4 = q4 / (q4' * phi_keep);

%%

% Output 5
Phi_use = Phi10;
phi_keep = Phi_use(:,3);
B = [Phi_use(:,[1 2 ])];
P = B * ((B' * B)\B');
q5 = (eye(size(Phi_use,1)) - P) * phi_keep;
q5 = q5 / (q5' * phi_keep);

%% output 6

% Output 6
Phi_use = Phi10;
phi_keep = Phi_use(:,4);
B = [Phi_use(:,[  9  11 12 13])];
P = B * ((B' * B)\B');
q6 = (eye(size(Phi_use,1)) - P) * phi_keep;
q6 = q6 / (q6' * phi_keep);

%%
```

```
% Output 7
Phi_use = Phi10;
phi_keep = Phi_use(:,5);
B = [Phi_use(:,[]) ];
P = B * ((B' * B)\B');
q7 = (eye(size(Phi_use,1)) - P) * phi_keep;
q7 = q7 / (q7' * phi_keep);

%%

for j = 1:Nu

    % excite one actuator at a time

    model.param.set("F1_amp", "0[N]");
    model.param.set("F2_amp", "0[N]");
    model.param.set("F3_amp", "0[N]");
    model.param.set("F4_amp", "0[N]");
    model.param.set("F5_amp", "0[N]");
    model.param.set("F6_amp", "0[N]");
    model.param.set("F7_amp", "0[N]");
    model.param.set("F8_amp", "0[N]");
    model.param.set("F9_amp", "0[N]");

    if j==1
        model.param.set("F1_amp", sprintf("%g[N]", F_in));
    elseif j==2
        model.param.set("F2_amp", sprintf("%g[N]", F_in));
    elseif j==3
        model.param.set("F3_amp", sprintf("%g[N]", F_in));
    elseif j==4
        model.param.set("F4_amp", sprintf("%g[N]", F_in));
    elseif j==5
        model.param.set("F5_amp", sprintf("%g[N]", F_in));
    elseif j==6
        model.param.set("F6_amp", sprintf("%g[N]", -2*F_in));
        model.param.set("F7_amp", sprintf("%g[N]", 2*F_in));
    elseif j==7
        model.param.set("F8_amp", sprintf("%g[N]", -F_in));
        model.param.set("F9_amp", sprintf("%g[N]", F_in));
    end


    fprintf("j=%d -> F1_amp=%s, F2_amp=%s, F3_amp=%s, F4_amp=%s, F5_amp=%s, F6_amp=%s, F7_amp=%s, F8_amp=%s, F9_amp=%s\n", ...
    j, ...
    model.param.get("F1_amp"), ...
    model.param.get("F2_amp"), ...
    model.param.get("F3_amp"), ...
    model.param.get("F4_amp"), ...
    model.param.get("F5_amp"), ...
    model.param.get("F6_amp"), ...
    model.param.get("F7_amp"), ...
    model.param.get("F8_amp"), ...
    model.param.get("F9_amp"));


    model.study(stdTag).run;

    % extract outputs across ALL frequencies
    y1 = mphinterp(model, expr_y1, "coord", coordS1, "dataset", dsetFR, "solnum", "all",'unit','m'); y1 = y1(:);
    y2 = mphinterp(model, expr_y2, "coord", coordS2, "dataset", dsetFR, "solnum", "all", 'unit','m'); y2 = y2(:);
    y3 = mphinterp(model, expr_y3, "coord", coordS3, "dataset", dsetFR, "solnum", "all", 'unit','m'); y3 = y3(:);


```

```
    R10 = zeros(nSens, Nf);

    for ii = 1:nSens
    ri = mphinterp(model, expr_y4raw, ...
        "coord", coord10(ii,:).', ...
        "dataset", dsetFR, ...
        "solnum", "all", ...
        "unit", "m");
    R10(ii,:) = ri(:).';
    end


    u_leaf1 = mphinterp(model, 'u', ...
        "coord", coordC1, ...
        "dataset", dsetFR, ...
        "solnum", "all", ...
        "unit", "m");
    u_leaf1 = u_leaf1(:);

    u_leaf2 = mphinterp(model, 'u', ...
        "coord", coordC2, ...
        "dataset", dsetFR, ...
        "solnum", "all", ...
        "unit", "m");
    u_leaf2 = u_leaf2(:);


    for k = 1:Nf
        H(1,j,k) = y1(k)/F_in;
        H(2,j,k) = -(y2(k) - y3(k))/F_in;
        w10 = R10(:,k);
        w11a  = [w10; u_leaf1(k)];
        w11b  = [w10; u_leaf2(k)];
        w12 = [w10; u_leaf1(k); u_leaf2(k)];
        H(3,j,k) = -(q3' * w12 ) / F_in;
        H(4,j,k) = (q4' * w10 ) / F_in;
        H(5,j,k) = (q5' * w10 ) / F_in;
        H(6,j,k) = (q6' * w10) / F_in;
        H(7,j,k) = (q7' * w10 ) / F_in;

    end

    fprintf("Done input %d/%d\n", j, Nu);
end

%%
% Build MIMO FRD
sysFRD = frd(H, f, "FrequencyUnit","Hz");
opts = bodeoptions;
opts.MagUnits="dB"; opts.PhaseUnits="deg"; opts.FreqUnits="Hz";
opts.Grid="on"; opts.PhaseMatching="on";

figure; bodeplot(sysFRD, opts);
xlim([min(f) max(f)]);
legend("COMSOL FRD");
title("7x7 MIMO: FRD");
```

F 7×7 MIMO Frequency Response

F.1. Magnitude Response

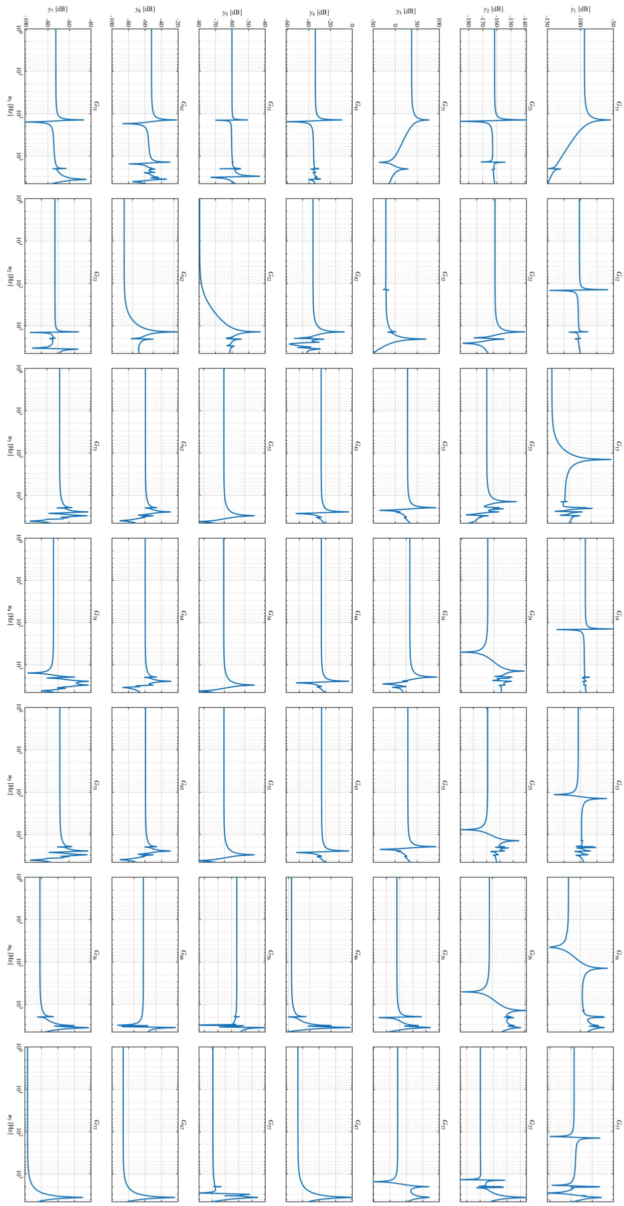

7 × 7 MIMO FRF Magnitude Matrix

F.2. PHASE RESPONSE

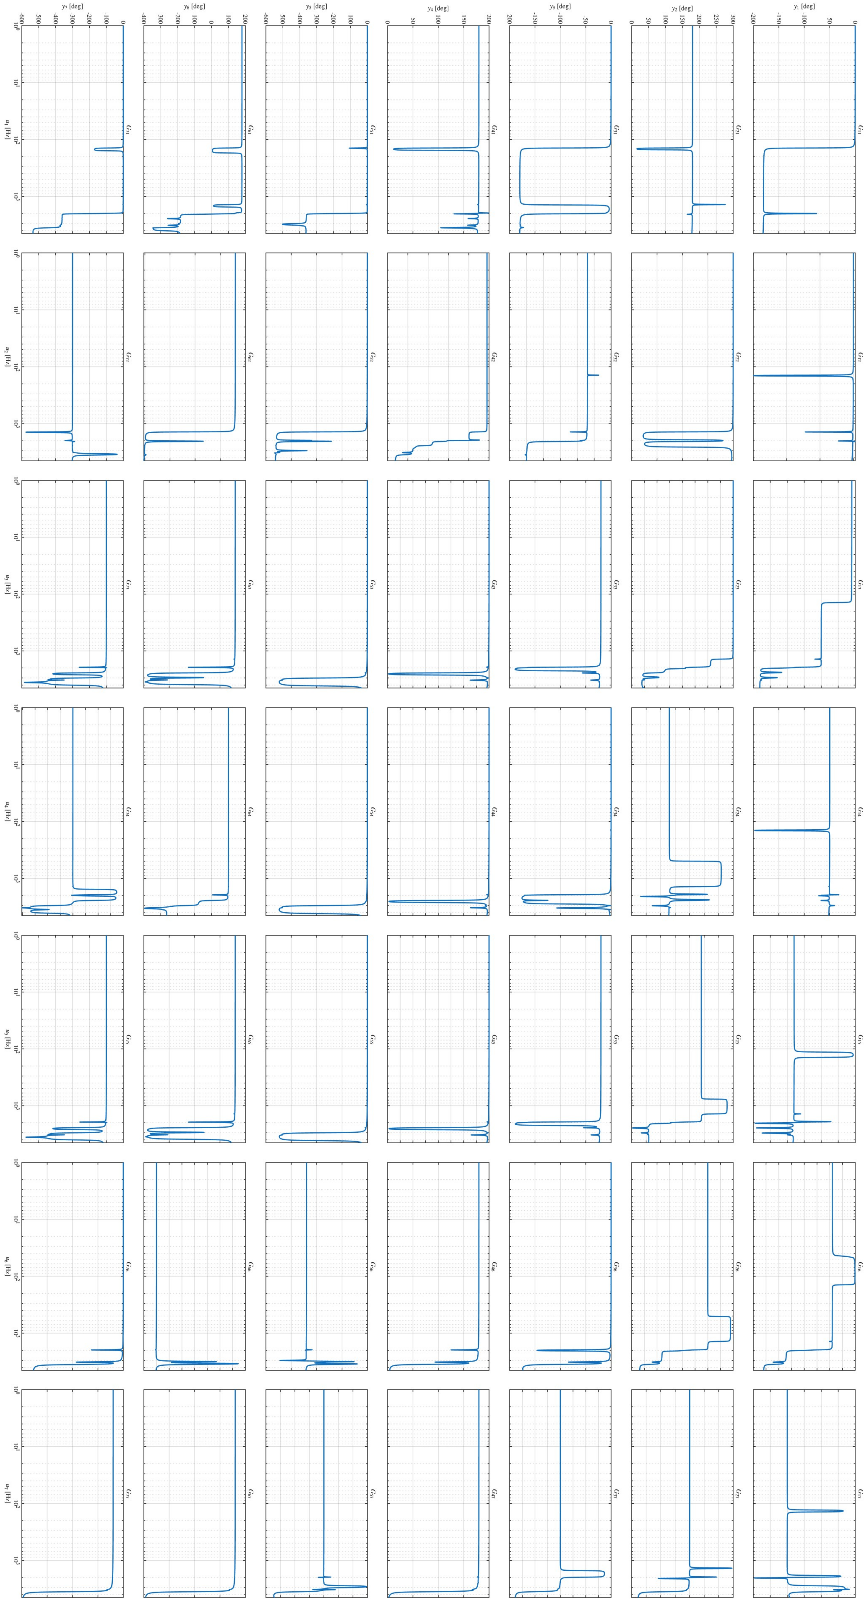

7×7 MIMO FRF Phase Matrix

G 7×7 Controller Comparison

G.1. Magnitude Plot

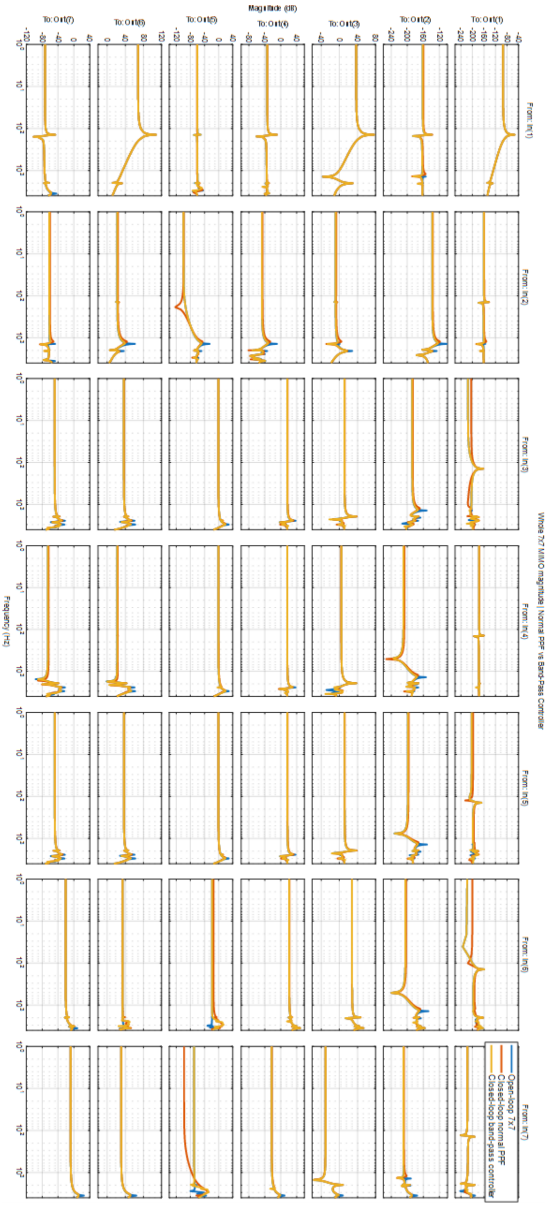

G.2. Phase Plot

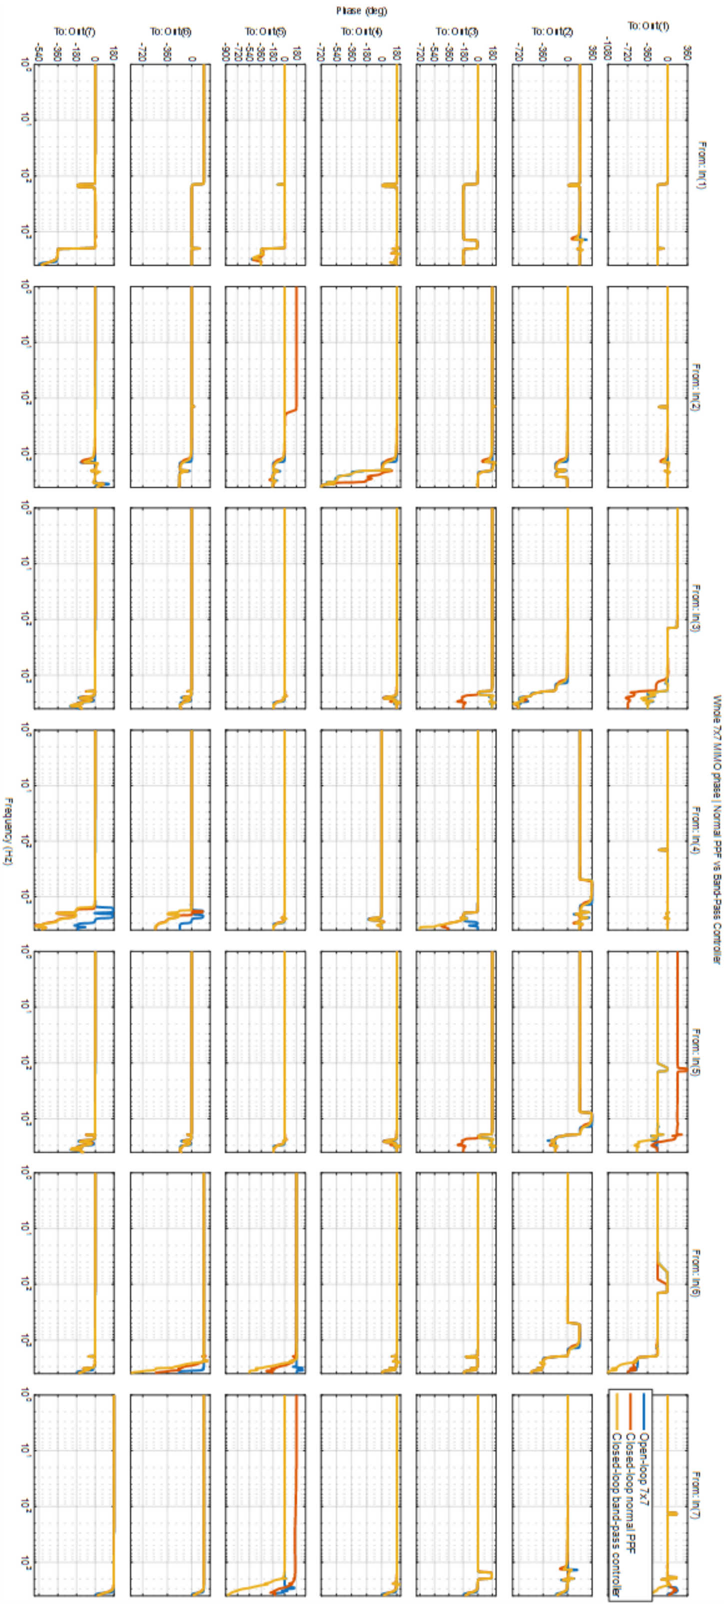

H 5×5 State-Space Model Validation

H.1. Magnitude Response

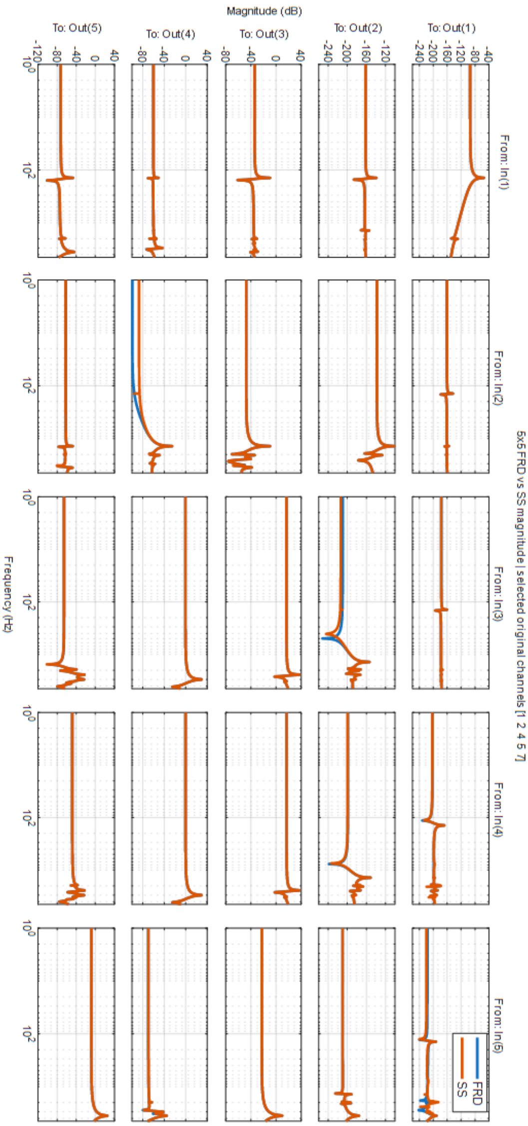

H.2. PHASE RESPONSE

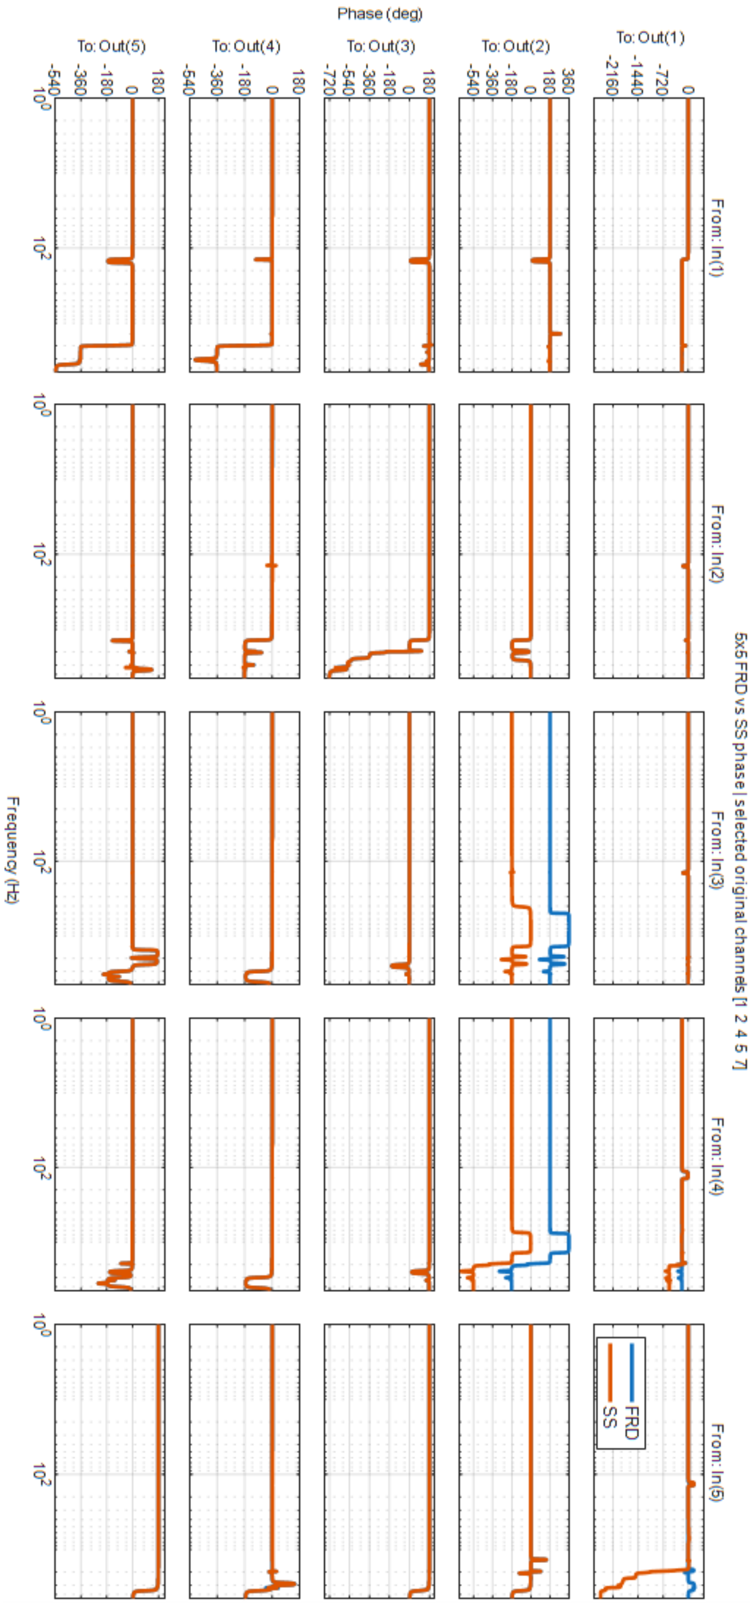

I 5×5 Open-Loop and Band-Pass PPF Closed-Loop Response

I.1. Magnitude Response

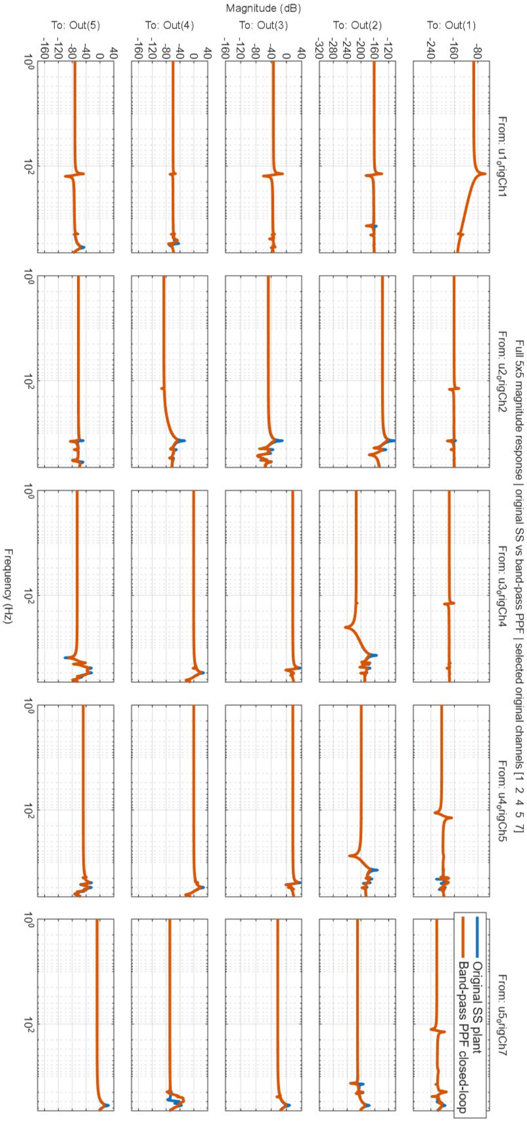

I.2. PHASE RESPONSE