Department of Precision and Microsystems Engineering

Vibro-Acoustic Reduced Order Modelling: A Modular Approach

Lucas Costa Arslanian

Report no : 2025.022
Supervisors : Bart van der Aa, ASML & Haydar Dirik, ASML
Professor : Peter Steeneken
Specialisation : DMN
Type of report : Master Thesis
Date : 02/07/2025

Vibro-Acoustic Reduced Order Modelling: A Modular Approach

MSc Thesis Report

by

Lucas Costa Arslanian

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on July 2, 2025 at 14:00

Thesis committee:

Chair:	Prof. Dr. Peter G. Steeneken, TU Delft
Assessment committee:	Dr. Jieun Yang, TU Delft
	Dr. Bart van der Aa, ASML
Additional supervisor	Ir. Haydar Dirik, ASML

Place:	Faculty of Mechanical Engineering, Delft
Project Duration:	September, 2024 - July, 2025
Student number:	5905753

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Faculty of Mechanical Engineering · Delft University of Technology

Abstract

Acoustic-structural interactions present significant engineering challenges, particularly in the domains of noise reduction and vibration control. At ASML, measurement-based analyses have revealed that acoustic disturbance paths often dominate the dynamic behavior of atmospheric lithography machines. This project focuses on enhancing ASML's current one-way coupled acoustic-structural modelling approach by developing a two-way coupled, known as vibro-acoustic, modelling framework. However, this advancement introduces substantial computational complexity, necessitating effective model reduction techniques. The primary objective of this work is to reduce vibro-acoustic models in a way that preserves their ability to be modularly coupled with other system components. To this end, three Component Mode Synthesis (CMS)-based reduction methods were evaluated, with only one proving suitable for both academic and industry-scale models. The resulting reduced-order models successfully retained the dynamic fidelity of the full system and enabled efficient coupling with other substructures. When applied to harmonic excitation analyses, the reduced models achieved a dramatic reduction in computational time, from several hours to approximately one minute, while accounting for the cost of model reduction.

Contents

Nomenclature

List of Abbreviations

CB	Craig-Bampton
CMS	Component Mode Synthesis
CTF	Condensed Transfer Function
DoFs	Degrees of Dreedoms
FE	Finite Element
FEM	Finite Element Modelling
FOM	Full Order Model
FRFs	Frequency Response Functions
FSI	Fluid-Structure Interface
HH	Hintz-Herting
MAC	Modal Assurance Criterion
MPC	Multipoint Constraint
NVH	Noise, Vibration and Harshness
ROM	Reduced Order Modelling
RPs	Remote Points
SE	Super Elements

List of Symbols

$\boldsymbol{g}$	Interface forces
$\boldsymbol{x}_a$	Acoustic pressure
$\boldsymbol{x}_b$	Boundary DoFs
$\boldsymbol{x}_i$	Internal DoFs
$\boldsymbol{x}_s$	Structural displacements
$\boldsymbol{\kappa}$	Reduced set of DoFs
$\mathbf{B}$	Mapping Boolean matrix
$\mathbf{C_a}$	Damping matrix of acoustic domain
$\mathbf{C_s}$	Damping matrix of structural domain
$\mathbf{F_a}$	Acoustic excitation
$\mathbf{F_s}$	Structural excitation
$\mathbf{K_a}$	Stiffness matrix of acoustic domain
$\mathbf{K_c}$	Coupled stiffness matrix
$\mathbf{K_s}$	Stiffness matrix of structural domain
$\mathbf{L}$	Localization Boolean matrix
$\mathbf{M_a}$	Mass matrix of acoustic domain
$\mathbf{M_s}$	Mass matrix of structural domain
$\mathbf{T}$	Transformation matrix
$\mathbf{Z}$	Impedance matrix
$\tilde{\mathbf{C}}$	Reduced damping matrix
$\tilde{\mathbf{K}}$	Reduced stiffness matrix
$\tilde{\mathbf{M}}$	Reduced mass matrix
$\boldsymbol{\epsilon}$	Modified relative frequency error
$\boldsymbol{\varepsilon}$	Relative frequency error
ρ_{a}	Fluid density
Z	Acoustic Impedance

List of Figures

List of Tables

Part I

Introduction

1

Literature Review

1.1. Introduction

Context

Acoustic-structural interactions are a major source of ongoing engineering challenges, particularly in areas such as noise reduction and vibration control. These challenges fall under the domain of Noise, Vibration, and Harshness (NVH) engineering, which focuses on optimizing a machine's vibro-acoustic behavior to enhance both comfort and reliability [1]. While NVH techniques are widely applied in the automotive industry, their relevance extends to any field where noise and vibration are critical concerns. Examples include reducing noise from household appliances like blenders and air conditioning units, or isolating vibrations in heavy industrial machinery. In the context of lithography machines, such as those developed by ASML, even the smallest acoustic disturbance can significantly impact system performance. As a result, ASML is actively investing in research and development to better understand and mitigate acoustic-related disturbances, aiming to improve the precision and stability of its machines.

Motivation

The dynamic behavior of ASML machines plays a critical role in their overall performance, and this becomes increasingly significant at higher throughputs. In atmospheric machines, disturbances that affect system dynamics and performance can originate from both structural and acoustic sources. However, measurement-based analyses have shown that acoustic disturbance paths often dominate. To accurately predict the impact of these acoustic disturbances, a multi-physical modelling approach is required, one that integrates both structural dynamics and acoustics. ASML has previously developed a one-way coupled modelling approach to incorporate acoustic effects into its dynamic system modelling framework. This method was validated using measurements from a case study involving the NXT1980Fi machine. While effective in certain scenarios, ASML recognizes that future improvements must focus on developing a two-way coupled modelling approach. This would allow for more accurate performance predictions by capturing the interplay between structural and acoustic domains, especially in cases where the one-way approximation is no longer sufficient.

One of the main challenges in implementing two-way coupling is the computational cost, which is largely determined by the size of the system matrices. As the number of nodes increases, the computational effort typically grows quadratically or cubically, depending on the element types, problem complexity, and the solver used. Two-way coupled models often involve between 10^5 and 10^7 degrees of freedom, making them computationally intensive. To address this, Reduced Order Modelling (ROM) techniques are essential. ASML has already applied these techniques successfully in purely structural simulations, reducing models of similar scale to just a few hundred degrees of freedom. This has transformed simulations that once took hours or days into ones that

complete in seconds or minutes, significantly improving efficiency.

Structure

This literature survey begins with a brief contextualization, providing a theoretical overview of vibro-acoustic analysis and reduced order modelling. Following this, the survey explores the current state of the art in reduced order modelling as applied to vibro-acoustic systems, highlighting key advancements and discussing their limitations. Based on this review and the specific needs of ASML, the report then introduces a project proposal aimed at addressing a novel modelling challenge. This is followed by a detailed plan of approach outlining the steps to be taken in the project.

1.2. Acoustic-Structure Interaction

When an elastic structure is in contact with a fluid, their dynamic interaction is known as an acoustic-structure interaction. In this process, acoustic pressure at the fluid-structure interface (FSI) exerts forces on the structure, altering its vibrations. At the same time, the structure's vibrations influence the pressure field within the fluid. The strength of this coupling depends on several factors, including the geometry of both the structure and the fluid domain, the material properties, and the frequency of the dynamic excitation.

Acoustic-structure interactions can be broadly classified into two categories based on the strength of the coupling between the structural and acoustic domains: uncoupled (one-way) and coupled (two-way) interactions. In one-way coupling, the interaction is relatively weak, meaning that the influence of one domain on the other is minimal. This allows for a simplified modelling approach, which is computationally less demanding. In contrast, two-way coupling involves a stronger interaction where the structure and the fluid continuously influence each other. The acoustic pressure at the fluid-structure interface affects the structural vibrations, and those vibrations, in turn, alter the acoustic field. This mutual dependency requires more complex and computationally intensive modelling techniques to accurately capture the dynamics of the system. The term "vibro-acoustic" is commonly used to describe this type of two-way coupled interaction. In the following sections, the report will explore the defining characteristics, modelling approaches, and typical applications for both one-way and two-way coupling scenarios.

There are several solution methods to solve acoustic-structure problems, each suited for a particular case. In order to determine the most appropriate approach, the type of problem should be considered. According to Davis [2], these problems are analysed by either energy, discretization, analytical, or empirical approaches. Energy methods are effective for high frequencies, as they divide the problem into energy storage subsystems and enforce power balance, relying on statistical modal properties. Discretization is the most commonly used technique, and it involves coupling finite elements of the structure and fluid (FE-FE) or finite elements of the structure and boundary elements of the acoustic domain (FE-BE), for exterior problems. However, discretisation becomes computationally challenging for high frequencies, as multiple elements per wavelength are required for accuracy, which results in a very fine mesh of elements. Analytical and empirical models are more limited options, as the analytical approach is only convenient for simple geometries and loading conditions, and empirical methods are just approximations from experimental data and do not provide any physical insight into the problem. For applications relevant to ASML, interior systems are usually considered. Therefore, since both the acoustic and structural domains are bounded in interior cases, they are suitable for finite element discretization (FE-FE) [2, 3].

As stated previously, acoustic-structure interactions problems deal with two domains, structural and acoustic. The structural domain is usually described with respect to displacements ($\boldsymbol{u}$), while the acoustic domain can be described by, most commonly, pressure ($\boldsymbol{p}$) or particle

velocity ($\boldsymbol{\varphi}$). This gives way to the displacement-pressure ($\boldsymbol{u} - \boldsymbol{p}$) and displacement-particle velocity ($\boldsymbol{u} - \boldsymbol{\varphi}$) formulations. These formulations will be presented in the following sections.

1.2.1. One-way Coupled

In uncoupled acoustic-structure systems, the interaction between the structural and fluid components is minimal and can be disregarded. Consequently, this allows the interaction to be simplified to only consider a one-way interaction, resulting in two distinct dynamic problems:

- **Structural Excitation of the Fluid**: In this scenario, the structural components generate vibrations that act as external excitations for the fluid. The vibrations of the structure create an acoustic pressure field in the fluid.
- **Fluid Excitation of the Structure**: In this context, the fluid applies independent force excitations to the structure. The acoustic waves propagating through the fluid generate a pressure distribution along the fluid-structure interface, which induces vibrations in the structure.

Modelling with a one-way interaction is particularly valid in many acoustic-structure systems, especially when a highly stiff elastic structure is surrounded by a low-density fluid. In these cases, the influence of the structure on the fluid is pronounced, while the pressure field radiated by the fluid exerts little to no effect on the vibrations of the structure [3, 4]. This modelling approach has proven effective at ASML for analysing low- and high-pressure zones created by a moving chuck in the NXT machines (atmospheric machines). Here, the chuck (the source) radiates sound which interacts with a receiver, but the receiver in turn is not radiating back into the surrounding. This one-way approximation was adequate in scenarios where the structural component was significantly stiffer than the surrounding acoustic medium. However, in cases involving more flexible structures with multiple connection interfaces, this approach becomes insufficient, as the mutual interaction between the acoustic and structural domains cannot be neglected.

1.2.2. Two-way Coupled

When considering fully coupled vibro-acoustic systems, the interaction between structural and fluid components becomes significant, requiring a unified treatment of the system rather than considering one domain as an independent excitation of the other. This approach offers a more accurate representation of the underlying physics, but also introduces increased complexity and computational demands. Accounting for this interaction becomes especially crucial in situations where an elastic structure is immersed in a dense fluid or a light structure interacts with a low-density fluid.

In this coupling, pressure in the fluid domain induces a force on the interface with the structural domain, and the acceleration of the structural domain is transmitted to the fluid domain. These systems were originally investigated in the popular displacement-pressure formulation by Craggs and Zienkiewicz [5, 6] where the coupled system results in nonsymmetric mass and stiffness matrices, as seen in the equations of motion in Equation (1.1). The complete demonstration of this displacement-pressure formulation can be referenced in Tabak [7].

$$\begin{bmatrix} \mathbf{M_s} & \mathbf{0} \\ \rho_a \mathbf{K_c}^\mathsf{T} & \mathbf{M_a} \end{bmatrix} \begin{Bmatrix} \ddot{\boldsymbol{x}}_\mathbf{s} \\ \ddot{\boldsymbol{x}}_\mathbf{a} \end{Bmatrix} + \begin{bmatrix} \mathbf{C_s} & \mathbf{0} \\ \mathbf{0} & \mathbf{C_a} \end{bmatrix} \begin{Bmatrix} \dot{\boldsymbol{x}}_\mathbf{s} \\ \dot{\boldsymbol{x}}_\mathbf{a} \end{Bmatrix} + \begin{bmatrix} \mathbf{K_s} & -\mathbf{K_c} \\ \mathbf{0} & \mathbf{K_a} \end{bmatrix} \begin{Bmatrix} \boldsymbol{x}_\mathbf{s} \\ \boldsymbol{x}_\mathbf{a} \end{Bmatrix} = \begin{Bmatrix} \mathbf{F_s} \\ \mathbf{F_a} \end{Bmatrix} \tag{1.1}$$

The matrices $\mathbf{M_s}$, $\mathbf{C_s}$, and $\mathbf{K_s}$ correspond to the mass, damping and stiffness of the structural domain, while the corresponding matrices, with subscript a, are the matrices of the

acoustic domain. Furthermore, $\boldsymbol{x_s}$ is the structural displacement of the nodes, $\boldsymbol{x_s}$ is the acoustic pressure of the nodes, and ρ_a is the mass density of the acoustic domain. The right-side terms indicate the excitation in the structural domain ($\mathbf{F_s}$) and acoustic domain ($\mathbf{F_a}$).

The domains are coupled through the off-diagonal elements. The entry $\rho_a \mathbf{K_c}^\intercal$ in the coupled mass matrix describes the imposed acoustic excitation on the fluid due to the displacement of structure, while the entry $-\mathbf{K_c}$ in the coupled stiffness matrix describes the force imposed on the structure due to fluid pressure, which is the area matrix that converts fluid pressures at interface points to structural loads. Since this coupling combines two different physics that have can have different orders, scaling of the system matrices is an alternative for more efficient computation [7].

As stated, vibro-acoustics is a multi-physics phenomena involving the interaction between structural and acoustic components. These two domains are coupled through the FSI, which serves as the boundary where the structural and fluid (acoustic) domains are physically and mathematically coupled. From a numerical modelling perspective, the FSI is characterized by a set of shared DoFs that represent both structural motion and acoustic pressure at the interface nodes, as illustrated in Figure 1.1.

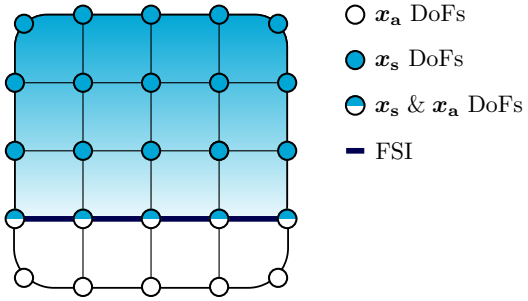

Figure 1.1: Representation of the DoFs and FSI in a vibro-acoustic model

Figure 1.1 presents a schematic of a vibro-acoustic model, where a structural domain (shown in white) is coupled to an acoustic domain (shown in cyan) via the FSI (in navy). The structural domain contains only structural DoFs ($\boldsymbol{x_s}$), and the acoustic domain contains only acoustic DoFs ($\boldsymbol{x_a}$). The exception is at the interface, where both types of DoFs coexist, which enables the necessary coupling between the two domains. The colour scheme used in Figure 1.1 will be consistently applied throughout this thesis for clarity and reference.

1.3. Reduced Order Modular Modelling

Finite element modelling techniques have become essential for analysing increasingly intricate dynamic structures, such as the lythography machines produced by ASML. However, these structures' growing complexity has also led to structural dynamic models of substantial size. This tendency is particularly pronounced in vibro-acoustic simulations, where element sizes must be at least one-sixth of the wavelength corresponding to the highest frequency of interest when employing quadratic elements, as a rule of thumb. Consequently, models can have millions of degrees of freedom (DoFs), resulting in prohibitively long computation times.

ASML's solution is a combination of reduced order modelling (ROM) and substructuring techniques. The latter ensures that the full system is first analysed by its individual components and then later assembled together. By analysing on a component level, the number of DoFs are naturally lower and the model is less complex. However, the main advantages lies in the modularity aspect, allowing for easier exchange of components or the reuse of validated parts, and enables parallel development of various subsystem models.

The ROM enables these components to be further reduced in complexity, for example, by creating super elements (SE). These super elements are a subset of the DoFs and maintain the dominant dynamic behaviour of the original component until a cutoff frequency. Reduced order models make it possible for simulations that used to take hours or days to be computed in seconds. This allows for a larger number of iterations to be performed.

Since computing dynamic problems (e.g. vibration modes) requires solving many static-like problems, which requires much longer computational times [8], it is also recommended to simplify complex geometries before generating the reduced order models. This entails maintaining most essential parts, which influence the overall stiffness of the module or local stiffness of the master nodes. For the geometric simplifications, it is advised to remove small fillets, valves, chamfers, pins, covers or holes if they are not in the region of master nodes. Note that this can become a complex procedure, and is usually performed by a specialised design team.

1.3.1. Component Mode Synthesis

The super elements used in ASML models are typically created using Component Mode Synthesis (CMS). CMS is a reduction technique that is often applied to condense large finite element (FE) models with many DoFs into significantly smaller systems with far fewer DoFs. The remaining DoFs, referred to as boundary or master DoFs, usually correspond to the interfaces with the system environment and to actuator or sensor locations. The primary goal of CMS is to reduce the size of the full model while preserving the essential dynamic behavior within the frequency range of interest.

In order to apply the CMS, the equation of motion, typically represented as seen in Equation (1.2), needs to be rewritten by separating the boundary ($\boldsymbol{x}_b$) and internal ($\boldsymbol{x}_i$) DoFs, as depicted in Equation (1.3).

$$\mathbf{M}\ddot{\boldsymbol{x}} + \mathbf{C}\dot{\boldsymbol{x}} + \mathbf{K}\boldsymbol{x} = \boldsymbol{f} \tag{1.2}$$

$$\begin{bmatrix} \mathbf{M}_{bb} & \mathbf{M}_{bi} \\ \mathbf{M}_{ib} & \mathbf{M}_{ii} \end{bmatrix} \begin{Bmatrix} \ddot{\boldsymbol{x}}_b \\ \ddot{\boldsymbol{x}}_i \end{Bmatrix} + \begin{bmatrix} \mathbf{C}_{bb} & \mathbf{C}_{bi} \\ \mathbf{C}_{ib} & \mathbf{C}_{ii} \end{bmatrix} \begin{Bmatrix} \dot{\boldsymbol{x}}_b \\ \dot{\boldsymbol{x}}_i \end{Bmatrix} + \begin{bmatrix} \mathbf{K}_{bb} & \mathbf{K}_{bi} \\ \mathbf{K}_{ib} & \mathbf{K}_{ii} \end{bmatrix} \begin{Bmatrix} \boldsymbol{x}_b \\ \boldsymbol{x}_i \end{Bmatrix} = \begin{Bmatrix} \boldsymbol{f}_b \\ \mathbf{0} \end{Bmatrix} \tag{1.3}$$

It is important to note that Equations 1.2 and 1.3 are formulated on a per-component basis. To obtain the reduced-order model (ROM), a transformation matrix $\mathbf{T}$ is applied to the full system, resulting in a reduced set of degrees of freedom, denoted by $\boldsymbol{\eta}$, as shown in Equation (1.4). The specific form of the transformation matrix depends on the chosen CMS method.

$$\boldsymbol{x} = \mathbf{T}\boldsymbol{\eta} \tag{1.4}$$

Substituting Equation (1.4) into (1.2) and pre-multiplying by $\mathbf{T}^\intercal$:

$$\underbrace{\mathbf{T}^\intercal\mathbf{M}\mathbf{T}}_{\tilde{\mathbf{M}}}\ddot{\boldsymbol{\eta}} + \underbrace{\mathbf{T}^\intercal\mathbf{C}\mathbf{T}}_{\tilde{\mathbf{C}}}\dot{\boldsymbol{\eta}} + \underbrace{\mathbf{T}^\intercal\mathbf{K}\mathbf{T}}_{\tilde{\mathbf{K}}}\boldsymbol{\eta} = \mathbf{T}^\intercal\boldsymbol{f} \tag{1.5}$$

where $\tilde{\mathbf{M}}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{K}}$ represent the reduced mass, damping, and stiffness matrices, respectively. The transformation matrix $\mathbf{T}$ is constructed from a set of deformation modes that define which dynamic behaviors are retained in the reduced model. The reduced coordinates $\boldsymbol{\eta}$ indicate the contribution of each deformation mode to the system's response. The choice of deformation modes used to construct $\mathbf{T}$ directly influences the accuracy and efficiency of the reduced model.

It is important to note that the boundary conditions used to compute these deformation modes are not related to the actual physical boundary conditions of the system. Instead, they are artificial constraints applied solely for the purpose of generating the reduction basis. The methods currently employed by ASML, along with the construction of their respective transformation matrices **T**, are thoroughly discussed by van Steen [9] and are summarized below:

- **Guyan** [10]: This method reduces the number of degrees of freedom in the substructure while preserving its static stiffness characteristics. It does not account for dynamic behavior. The reduction is achieved using static constraint modes, which are the deformations resulting from applying unit displacements to each boundary DoF while keeping the others fixed. This method retains DoFs that significantly influence the static response.
- **Craig-Bampton** [11]: An extension of the Guyan method, it incorporates fixed-interface normal modes to capture the dynamic behavior of the internal structure. It is particularly effective for components with limited motion at the interfaces. However, the transformation matrix must be recomputed if the set of boundary DoFs changes.
- **Hintz-Herting** [12]: Similar to the Craig-Bampton approach, this approach incorporates free-interface normal modes, but it also includes inertia relief and residual flexibility modes in the reduction basis. Unlike Craig-Bampton, it does not require the boundary degrees of freedom to be fixed. As a result, the computed eigenfrequencies and mode shapes remain unaffected by changes in the boundary DoFs.

Computing the mode shape vectors used in mode superposition methods can be computationally demanding. Typically, a finite set of modes is extracted by solving an eigenvalue problem. The expansion process retains the initial modes corresponding to the lowest eigenfrequencies, resulting in a truncated basis that captures the most significant dynamic behavior [13].

Regardless of the specific CMS-based method chosen, the fundamental concept remains the same, as illustrated in Figure 1.2. The process begins with the creation of discretized FE models for each component to be reduced (Figure 1.2(a)). These FE models are then reduced by retaining all boundary degrees of freedom (DoFs) and representing the internal DoFs using a set of mode shapes, which vary depending on the selected CMS method (Figure 1.2(b)). Finally, because all boundary DoFs are preserved during the reduction, the reduced components can be assembled into a complete system model (Figure 1.2(c)), which is a form of dynamic substructuring.

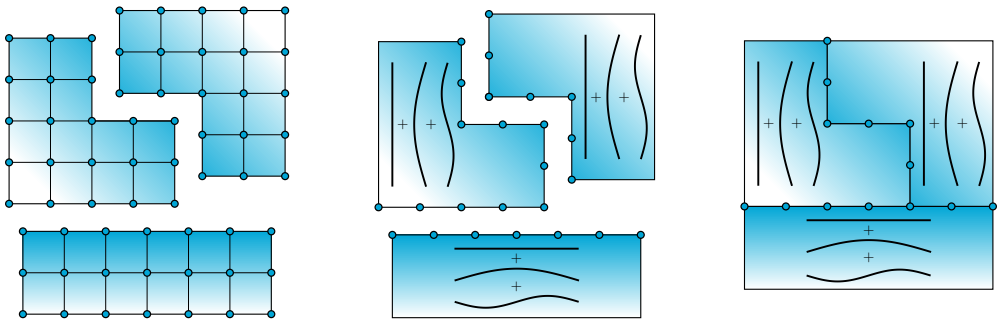

(a) FE discretisation

(b) Reduction of internal DoFs

(c) Assembly of the reduced components

Figure 1.2: Illustration of the CMS steps

1.3.2. Dynamic Substructuring

Dynamics substructuring refers to methods for synthesizing complex dynamic systems through their substructure models. This technique allows the analysis of very complex and large structures that cannot be analysed in a single model, while also reducing computational demand. It assumes substructures have internal and interface DoFs, with dynamic interaction occurring via interface DoFs. Substructures can be considered as superelements, where only the interface dynamics is relevant for substructure assembly [14].

Substructures are subcomponents that interact with their neighbouring subcomponents to form a complete system. When forming a coupled system from initially separate subcomponents, the interface must meet two essential criteria: ensuring the compatibility of primary variables, such as displacement and pressure, and maintaining the equilibrium of flow-related or dual variables, like force and volume flux [15]. Ensuring the equations that describe the system (e.g. equations of motion) satisfy these conditions is enough to grantee that the substructures are assembled. Following the general framework proposed by Klerk et al. [14], these conditions (in the physical domain) can be expressed as:

$$\begin{cases} \mathbf{B}\boldsymbol{x} = \mathbf{0} \\ \mathbf{L}^{\intercal}\boldsymbol{g} = \mathbf{0} \end{cases} \tag{1.6}$$

where $\mathbf{B}$ is a matrix that acts on the interface DoFs and is a known as the Boolean matrix on the matching interface DoFs of the substructures. $\mathbf{L}$ is a matrix that transforms the DoFs of the assembled structure to the DoFs of the substructures to be coupled, localising the corresponding DoFs of the substructures that can be expressed as a single DoF. $\boldsymbol{x}$ and $\boldsymbol{g}$ are the primal and dual variables, respectively. Allen et al. [16] shows an illustrative examples of how these matrices are constructed.

In this way, the dynamics of the substructures are analysed independently, and once the local solutions are defined and represented in a subspace using the DoFs of the physical coordinates of each substructure, the substructures are then coupled. Dynamic substructuring can be applied across various domains suitable for simulating the behavior of multivariable systems, known as multiple-input and multiple-output (MIMO) systems. The choice of domain depends on the components under analysis, the type of analysis performed, the availability of computational and experimental resources, and personal preferences. Five suitable domains are [16, 17]:

- **Physical Domain**: structures are characterized by their mass, stiffness, and damping distributions.
- **Modal Domain**: substructure dynamics are approximated in a reduction basis composed by a combination of modal responses.
- **Frequency Domain**: structures are described by their impedance or admittance Frequency response functions.
- **State-Space Domain**: structures are represented by inputs, outputs, and states related by first-order differential equations.
- **Time Domain**: substructures response are obtained either by time integration or simulated based on impulse response functions [17].

For each domain, many coupling techniques exist and are classified as either a primal or dual assembly, depending if the substructures are assembled by interface displacements (primal) or interface forces (dual) [18]. Dual approaches gained prominence as they can be implemented

with efficient solvers on parallel processing computers [14]. However, they retain all the DoFs of the substructures (generating redundant DoFs on the matching interfaces). FEM packages generally use primal assembly. The main techniques are further discussed in Section 1.4.2.

1.3.3. ASML's Modular Way of Work

Substructuring in combination with reduced order modelling is key to ASML's modular methodology. It boosts computational efficiency, and allows various teams to concurrently work on different components of the machines, to then seamlessly assemble them together. Furthermore, it ensures that components can be easily swapped out with updated versions or reused if they have been validated. Lastly, it enables global sensitivities to be linked with key components.

ASML has a proprietary framework for its modular way of working. Without disclosing any confidential details, the general workflow can be summarized as follows: the process begins with establishing a high-level overview of how individual components interact, ensuring that all relevant interfaces required for system assembly are identified. Once this interface map is defined, each component is modelled using finite element (FE) software.

After the subsystem models are finalized, they must be prepared for reduction into Super Elements (SEs). This preparation involves creating master nodes (the degrees of freedom that will be retained after reduction) and coupling them to a set of nodes in the unreduced component. Subsequently, all interface connections are documented in a spreadsheet, including the specified stiffness and damping values for each connection. It is crucial to ensure that matching interface master nodes align within a defined tolerance. This procedure is further explained in Section 5.2.1.

With the interface definitions in place, each component can be reduced into a super element using a selected Component Mode Synthesis (CMS) method. Finally, using the generated SEs and the interface spreadsheet, the complete system assembly can be constructed. This enables system-level simulations and post-processing analyses. An overview of these steps is illustrated in Figure 1.3, using a simplified test-bench composed of five substructures.

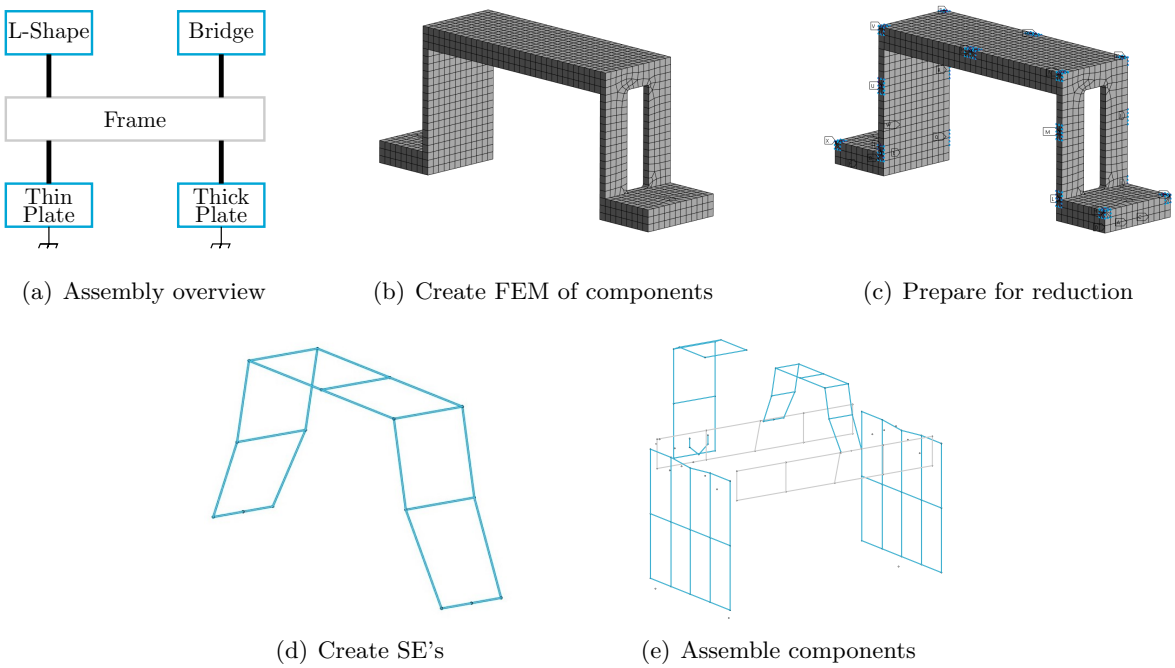

(a) Assembly overview (b) Create FEM of components (c) Prepare for reduction

(d) Create SE's (e) Assemble components

Figure 1.3: Overview of the modular way of work procedure

It is important to note that this modular workflow introduces some specific considerations when applying the CMS methods and coupling the substructures. Unlike traditional CMS applications where all interface DoFs are retained, this approach uses the master nodes as the boundary set. This significantly reduces computational effort, as fewer constraints modes need to be computed, which is especially beneficial for large-scale models such as those used in ASML machines. In this approach, the reduced models are coupled by directly applying specified stiffness and damping relationships between the retained master nodes.

This procedure is described in detail in Section 5.2.1 and illustrated in Figure 1.4. It is worth highlighting that this represents a different type of substructure coupling compared to the primal assembly shown in Figure 1.2. Rather than forming a direct union of substructures, the components are connected through defined mechanical properties (stiffness and damping) between master nodes. This modelling approach is particularly well-suited for representing physical connections commonly found in ASML systems, such as bolted joints.

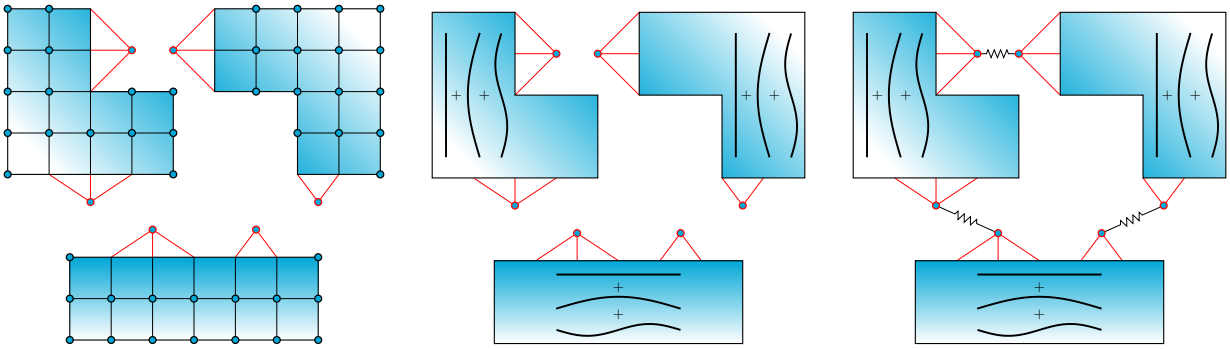

(a) Creating the master nodes

(b) Reducing the modes, retaining only the master nodes

(c) Coupling the master nodes with joint connections

Figure 1.4: CMS method tailored for ASML's modular modelling framework

Figure 1.4 is intended solely to illustrate the concept of master node creation and the way these nodes are used to connect substructures. In practice, these master nodes, also referred to as Remote Points, are typically positioned at the geometric centre of the group of nodes they substitute. Additionally, they must be defined in the same universal coordinate system as the corresponding master node they connect to in the adjacent substructure. While Figure 1.4(b) provides a simplified and intuitive visualization of the concept, it does not reflect the actual implementation. The true result of this coupling approach is more accurately represented in Figure 1.5 below.

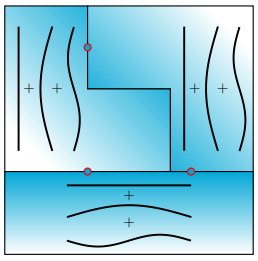

Figure 1.5: Actual representation of the outcome of the coupling procedure

1.4. Coupling Methods

1.4.1. Coupling Approaches for Vibro-Acoustic Systems

In order to incorporate vibro-acoustic simulations within ASML's modular modelling framework, the simplified model shown in Figure 1.6 will be used to showcase the possible modelling approaches that can be implemented. Figure 1.6 shows a vibro-acoustic model whose structural component is coupled to a separate structure through three interface points (represented as springs).

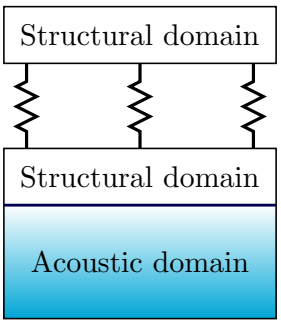

Figure 1.6: Vibro-acoustic system

From Figure 1.6, two possibilities within reduced order modular vibro-acoustic modelling can be explored for ASML. The first corresponds with the vibro-acoustic simulation itself, where the full problem is analysed by its subcomponents and then assembled to obtain the response of the full system. The other, is to model a vibro-acoustic module and then couple it with a separate structural module, with the goal of incorporating the influence of the vibro-acoustic interaction within a larger system. These approaches are seen in the possible cases shown in Figure 1.7, where Cases A and B represent variations of the first alternative, while Cases C and D correspond to the second.

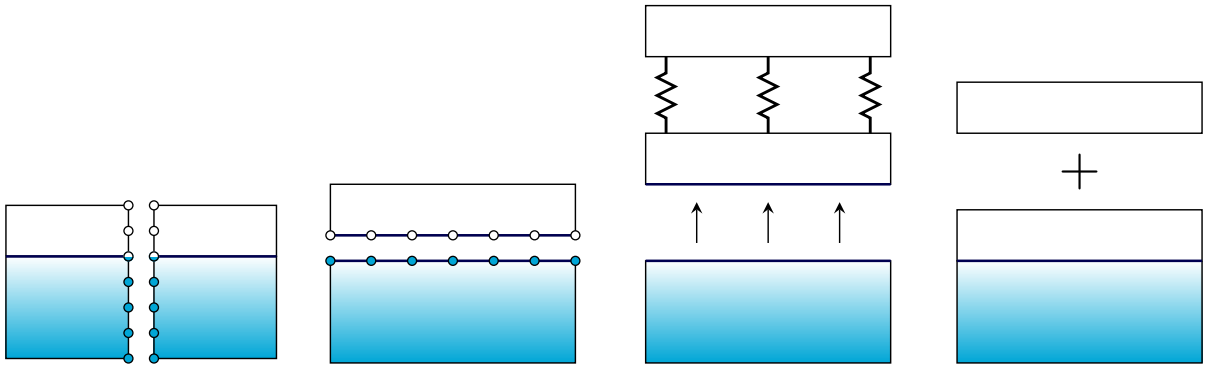

(a) Case A (b) Case B (c) Case C (d) Case D

Figure 1.7: Possible approaches

- **Case A**: This case is used for coupling different vibro-acoustic models together. It is a substructuring technique widely implemented to solve complex systems.
- **Case B**: A versatile modelling technique that has not been solved for. It consists solving an acoustic and a structural domains independently and, determining the coupling conditions for the interface to couple the domains without solving the full system.
- **Case C**: This case is strictly for one-way coupling problems. It consists of solving the

acoustic domain to obtain the pressure filed at the FSI, and then using it as an input for a separate structure simulation.

- **Case D**: An approach consisting of solving a vibro-acoustic problem and then combining it with another structural component.

Cases A and C either have a defined way of work in ASML or are currently in development, while Cases B and D still require the elaboration of a methodology to solve these types of problems. While the four cases are presented as separate, they can be combined into hybrid cases for a multi-step coupling approach if needed.

1.4.2. Coupling Techniques

For the different approaches identified, how the coupling procedure is performed the two-way coupling is an additional decision. Equations (1.7) to (1.10) show the governing equations of the substructuring problem for the for the physical, frequency, modal and state-space domains [14].

$$\begin{cases} \mathbf{M}_d\ddot{\boldsymbol{x}} + \mathbf{C}_d\dot{\boldsymbol{x}} + \mathbf{K}_d\boldsymbol{x} = \boldsymbol{f} + \boldsymbol{g} \\ \mathbf{B}\boldsymbol{x} = \mathbf{0} \\ \mathbf{L}^\intercal\boldsymbol{g} = \mathbf{0} \end{cases} \tag{1.7}$$

$$\begin{cases} \mathbf{Z}(\omega)\boldsymbol{x}(\mathrm{i}\omega) = \boldsymbol{f}(\omega) + \boldsymbol{g}(\omega) \\ \mathbf{B}\boldsymbol{x}(\omega) = \mathbf{0} \\ \mathbf{L}^\intercal\boldsymbol{g}(\omega) = \mathbf{0} \end{cases} \tag{1.8}$$

$$\begin{cases} \tilde{\mathbf{M}}_d\ddot{\boldsymbol{\eta}} + \tilde{\mathbf{C}}_d\dot{\boldsymbol{\eta}} + \tilde{\mathbf{K}}_d\boldsymbol{\eta} = \tilde{\boldsymbol{f}} + \tilde{\boldsymbol{g}} \\ \tilde{\mathbf{B}}\boldsymbol{\eta} = \mathbf{0} \\ \tilde{\mathbf{L}}^\intercal\tilde{\boldsymbol{g}} = \mathbf{0} \end{cases} \tag{1.9}$$

$$\begin{cases} \dot{\boldsymbol{x}} = \mathbf{A}_D\boldsymbol{x} + \mathbf{B}_D(\boldsymbol{u}+\boldsymbol{g}) \\ \boldsymbol{y} = \mathbf{C}_D\boldsymbol{x} + \mathbf{D}_D(\boldsymbol{u}+\boldsymbol{g}) \\ \mathbf{B}\boldsymbol{y} = \mathbf{0} \\ \mathbf{L}^\intercal\boldsymbol{g} = \mathbf{0} \end{cases} \tag{1.10}$$

Note that the system matrices ($\mathbf{M}_d$, $\mathbf{C}_d$ and $\mathbf{K}_d$), admittance matrix ($\mathbf{Z}_d$), reduced matrices ($\tilde{\mathbf{M}}_d$, $\tilde{\mathbf{C}}_d$ and $\tilde{\mathbf{K}}_d$) and state matrices ($\mathbf{A}_D$, $\mathbf{B}_D$, $\mathbf{C}_D$ and $\mathbf{D}_D$) are all block-diagonal, with each block corresponding to a substructure. The different subscripts d and D are only used to differentiate between the damping matrix $\mathbf{C}_d$ and the output matrix $\mathbf{C}_D$.

Coupling in the physical domain, known as direct coupling, demands a complete understanding of the system's mechanical and geometrical properties. These techniques are most effective for numerical models from FEM simulations and analytical models, where such properties are well-defined. Conversely, when dealing with experimental models where these properties are not fully known, frequency techniques like the LM-FBS [19] are more appropriate, as they allow the direct use of experimentally obtained frequency response functions in the assembly process. For large and complex systems, the modal domain is employed, primarily via the established CMS method. This type of coupling is similar to direct coupling, but utilizes reduced matrices obtained through mode truncation [16]. When the internal information of the system is significant or the system has numerous inputs and outputs, another option is to use the state-space representation for coupling. The LM-SSS [20, 21] is a promising recent technique that simplifies the process by requiring only the inversion of a single matrix, although it complicates the physical interpretation of the system. Equations (1.7) to (1.10) each have various solution approaches, which are classified as either a primal or dual.

1.5. State of the Art

1.5.1. Vibro-acoustic model order reduction

Within vibro-acoustic model order reduction, van de Walle et al. [22] introduces two novel stability-preserving methods for ROM of vibro-acoustic finite element models. Stability-preserving reduction methods are necessary for time-domain active noise control and auralisation (technique of artificially

making an acoustical situation audible). The first method reformulates the displacement-fluid velocity potential (u–ϕ) formulation, adopting stability conditions through one-sided projection, and offers an asymmetric system matrix structure proven to preserve stability. The second method uses a displacement-pressure (u–p) formulation with decoupled projection spaces, ensuring system stability. Two numerical validation cases illustrate the limitations of conventional model order reduction and the effectiveness of these new techniques. Both methods maintain original system stability, though the first requires non-standard form description and the second necessitates increased ROM size for similar accuracy.

Davis [2] proposes a method to construct uncoupled structural and acoustic FE models and then attempts to couple them outside of the FE software. The method begins by constructing finite element models to derive the natural frequencies and modes of both the structure and the rigid wall fluid cavities. The nodal displacements for each acoustic and structural mode are exported to data files, and MATLAB is used to compute their wavenumbers via Fast Fourier Transform (FFT) of the nodal displacement data. The coupling extent of individual acoustic and structural modes is then assessed through numerical integration over the fluid-structure interface. The full problem is formulated as a set of coupled ordinary differential equations, resulting in an nonsymmetric eigenvalue problem. However, by expressing it in modal degrees of freedom, the problem involves fewer equations than typical FE-FE models, which are based on physical node degrees of freedom.

Leopoldo et al. presents a methodology for simulating vibro-acoustic systems with the goal of including this coupled model in a closed loop control simulation framework that also takes into account the interaction between the system and the control sensors and actuators. This requires a reduced modal state-space formulation. In [23], Leopoldo et al. show that by substituting the structural and acoustic component mode expansions, Equation (1.11) into Equation (1.1), then pre-multiplying the the whole equation by the transpose of the structural and acoustic modal vectors can transform, and finally applying the variable substitution $\bar{q}_s = \omega^2 q_s$, the problem can reach a symmetric form. Since this form is based on the uncoupled eigenvalue solution of each component, the assembly of the system of equations is less demanding when compared to the solution of Equation 1.1.

$$\begin{aligned} \boldsymbol{x}_s &= \boldsymbol{\Phi}_s \, \boldsymbol{q}_s \\ \boldsymbol{x}_a &= \boldsymbol{\Phi}_a \, \boldsymbol{q}_a \end{aligned} \tag{1.11}$$

In this way, Leopoldo et al. demonstrates in [24] how this fully coupled vibro-acoustic system can be written in a reduced modal state-space formulation from the left and right eigenvectors ($\boldsymbol{\Phi}_L$ and $\boldsymbol{\Phi}_R$), as seen in Equation (1.12). This formulation features coupled structural and acoustical inputs and outputs. With this methodology, it was shown that the original problem analysed with 24192 DoFs was reduced to the state-space model with 214 states.

$$\begin{aligned} \begin{Bmatrix} \dot{\boldsymbol{q}} \\ \ddot{\boldsymbol{q}} \end{Bmatrix} &= \begin{bmatrix} 0 & \mathbf{I} \\ -\Omega^2 & -\Gamma \end{bmatrix} \begin{Bmatrix} \boldsymbol{q} \\ \dot{\boldsymbol{q}} \end{Bmatrix} + \begin{bmatrix} \mathbf{0} \\ \boldsymbol{\Phi}_L^{\intercal}[b] \end{bmatrix} \begin{Bmatrix} \mathbf{F_s} \\ \mathbf{F_a} \end{Bmatrix} \\ \boldsymbol{y} &= \begin{bmatrix} \begin{bmatrix} c \end{bmatrix} \boldsymbol{\Phi}_R^{\intercal} & \mathbf{0} \end{bmatrix} \begin{Bmatrix} \boldsymbol{q} \\ \dot{\boldsymbol{q}} \end{Bmatrix} \end{aligned} \tag{1.12}$$

Kim et al. [25] introduce a ROM technique for vibro-acoustic simulations. The strategy integrates fluid and structure modes through a two-phase reduction, as seen in Equation (1.13).

Firstly, structure modes are primarily reduced, and fluid modes are modified accordingly, allowing the fluid's modal properties to complement the structure's, clarifying the transference of vibrational energy. This strongly coupled method enhances modelling accuracy by simplifying the reduction process compared to previous techniques, only requiring minor changes to the fluid domain's mass matrix for negligible computational impact, as the system size is unchanged when compared to the weakly coupled. Its effectiveness is validated through four numerical examples. Although the study primarily used eigenvectors as the projection basis, the method is also compatible with the Krylov subspace basis.

$$\mathbf{T}_{sc} = \mathbf{T}^{s}_{wc}\mathbf{T}^{f}_{wc} = \begin{bmatrix} \mathbf{\Phi}_d & \mathbf{\Psi}\tilde{\mathbf{\Xi}}_d \\ \mathbf{0} & \tilde{\mathbf{\Xi}}_d \end{bmatrix} \tag{1.13}$$

where $\mathbf{\Psi}$ is a coupling matrix that defines the relationship between the structural and pressure DoFs at the interface. The matrices $\mathbf{\Phi}_d$ and $\tilde{\mathbf{\Xi}}_d$ represent the structural and fluid dominant eigenvectors, respectively. Transitioning from a weakly coupled to a strongly coupled system solely involves modifying $\mathbf{\Xi}_d$ to $\tilde{\mathbf{\Xi}}_d$, achieved by using the reduced mass matrix of the fluid domain from the initial transformation.

Interpolation and moment-matching techniques seek for low-order approximations that replicate the behavior of the original transfer function at specific expansion points and are another group of ROM techniques. Puri et al. [26, 27, 28] introduces a novel approach to create efficient reduced order models for comprehensive vibro-acoustic analysis. In this method, the basis vectors required for model reduction are generated using the Arnoldi algorithm, which calculates the projection vectors that form the Krylov subspace, ensuring the maximum quantity of moments of the system. This method presents an alternative to many reduction techniques, especially useful for vibro-acoustic optimization, where reducing computational time is crucial. Furthermore, it guarantees complete output approximation by accurately matching both the displacement of the structure and sound pressure levels in the fluid. This method is demonstrated through academic examples, including a plate with a cavity enclosed by rigid walls and a simplified vehicle framework composed of beams and plates, where the reduced models where compared with ANSYS simuloations of the full models.

Unlike the commonly used finite element methods, Desmet [3] presents a new deterministic wave-based prediction technique for coupled vibro-acoustic analysis. In this method, the steady-state dynamic field variables are represented using acoustic and structural wave functions that solve the homogeneous parts of the dynamic equations exactly, as well as particular solution functions for the non-homogeneous equations. This transforms the continuous problem of identifying field variable distributions into a discrete task of evaluating the contributions of these wave functions to the expansions of the field variables. These expansions inherently satisfy the dynamic governing equations, and therefore determining the wave function contributions depends only on the acoustic and structural boundary conditions. Compared to finite element approaches, the wave-based prediction method offers distinct advantages: it is much more compact, and refining the model is straightforward. Nonetheless, for each frequency of interest, the wave model requires re-computation, and natural frequencies and mode shapes cannot be derived through conventional eigenvalue analysis, which requires a full set of wave functions and appropriate particular solution functions.

An alternative wave-based approach is introduced by Donders et al. [29], the Wave-Based Substructuring method. This approach involves portraying interface DoFs as a linear combination of a set of basis functions, referred to as waves. Interface DoFs connections, which are typically defined between substructures, are substituted by wave connections that maintain displacement

and force continuity. Since the number of basis functions is generally significantly smaller than the number of interface DoFs, this significantly simplifies the model reduction process for components of interest, as it reduces the number of enrichment vectors needed. Consequently, any local modifications within the finite element representation of components can be efficiently processed to predict changes in structural performance concerning modes and responses.

1.5.2. Vibro-acoustic substructuring

Tabak [7, 30] suggests an innovative and computationally efficient reduction method for substructuring of vibro-acoustic problems, based in the Craig-Bampton reduction. This approach transforms the original nonsymmetric component matrices, shown in Equation (1.1), into κ-symmetric and τ-symmetric forms. The κ-symmetric form is achieved by multiplying Equation (1.1) with κ as defined in Equation (1.14), while the τ-symmetric form results from a transformation of $\boldsymbol{u} - \boldsymbol{p}$ formulation to the $\tilde{\boldsymbol{u}} - \boldsymbol{p}$ format, as illustrated in Equation (1.15). This allows symmetric solvers to be used for determining the elements of the Craig-Bampton basis at the component level.

$$\kappa = \begin{bmatrix} \mathbf{I} & \mathbf{K_c}\mathbf{K_a}^{-1} \\ \mathbf{0} & \mathbf{M_a}\mathbf{K_a}^{-1} \end{bmatrix} \tag{1.14}$$

$$\begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{p} \end{bmatrix} = \begin{bmatrix} \mathbf{K_a}^{-1}\boldsymbol{M_s} & \mathbf{K_s}^{-1}\mathbf{K_c} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{u}} \\ \boldsymbol{p} \end{bmatrix} \equiv \tau \tilde{\boldsymbol{q}} \tag{1.15}$$

Tabak shows how to apply the Craig-Bamptom reduction to vibro-acoustic components and then assemble them using the framework proposed by Klerk et al. [14], explained in Section 1.3.2. However, the peculiarities for of applying this framework for vibro-acoustic system are highlighted. This reduction and assembly procedures are validated for two versions of the system shown in Figure 1.8, which clearly corresponds to a Case A type of coupling from Figure 1.7. This technique was implemented in Section 4.2, but was not successful in reducing the analysed substructure.

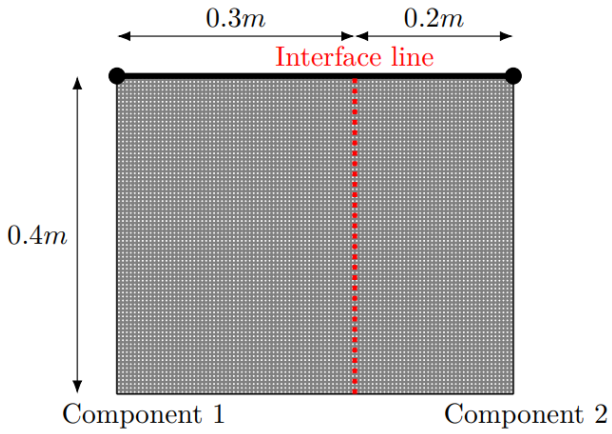

Figure 1.8: 2D Beam coupled to acoustic domain. Reference: Tabak [7]

Herrmann [31] also proposes and extension of the Craig–Bampton method for systems with coupled acoustic and structural elements. Herrmann proposes an interface reduction by Rutz vectors. The method is tested using harmonic analysis on a multifaceted pipe system filled with fluid, consisting of five interconnected substructures, depicted in Figure 1.9. This dynamic substructuring method facilitates the efficient assessment of various 3D pipe setups and improves vibro-acoustic optimization, achieving about half the DoFs relative to the traditional Craig-Bampton method without interface reduction.

In [32], Herrmann applies this approach to two automotive vibro-acoustic scenarios to enhance model order reduction techniques and speed up finite element analysis. One application involves a fluid-filled brake-pipe system constructed from superelements using fully coupled fluid

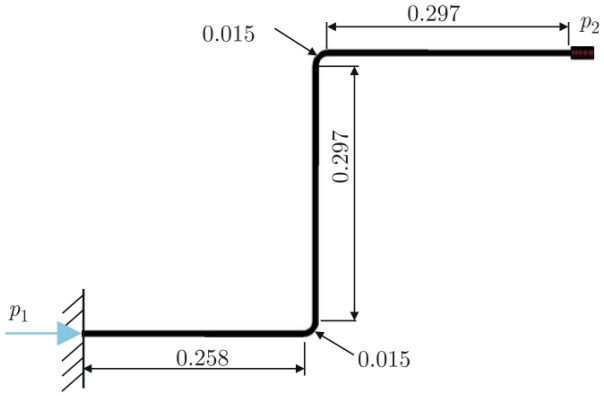

Figure 1.9: Vibro-acoustic pipe assembled system. Reference: Herrmann [31]

and structure components to detail interactions between the pipe shell and the internal fluid. The second application focuses on a rear muffler excited by air-borne sound. Here, the Rubin and Craig-Bampton methods are utilized to compute modes separately for acoustic and structural components, which are then used to develop a reduced model with full acoustic-structure coupling. This technique, which is described in more detail in Section 3.3.1, served as the primary foundation for the methodology developed in this project.

Lindberg et al. [33] also introduces an innovative CMS technique, which utilises undeformed coupling interfaces. This approach offers advantages over traditional CMS methods, which are restricted by the size and compatibility of coupling interfaces. This novel approach allows for a substantial reduction of the original problem, whereas conventional CMS methods are constrained to reducing only the degrees of freedom not linked with coupling interfaces. It also effectively addresses issues related to non-conforming meshes between different components. Nonetheless, its application is limited to structures with interactions between flexible and rigid substructures, where the deformation is predominantly influenced by the rigid substructure and the interface remains mostly undeformed. This formulation is validated on a cross structure depicted in Figure 1.10, which serves as a model for the linking arm in a vehicle suspension system. The proposed reduction technique achieves a relative difference off-resonance of approximately 1% compared to the classical Craig-Bampton CMS method. In [34], Lindberg et al. performs a sensitivity analysis for this methodology on a model of a vehicle suspension system. In this study, it was verified that the power of the structure-borne sound transferred to the car body can be altered by orders of magnitude by rotating the connecting rubber bushings of the suspension system.

The substructuring methodology presented, where an undeformed coupling interface formalism is used to model a vehicle suspension system, gives at hand a powerful tool for investigating the complex mechanisms associated with structure-borne sound in vehicles. By using appropriate models for each substructure, a study of the influence of the rubber bushings used to connect the link arm to the car body is undertaken. It is found that, by rotating the bushings, the power transferred to the car body can be altered by orders of magnitude, on one hand opening up for radically decreasing the disturbances transmitted through the suspension system, but on the other hand as well identifying a source of strong sensitivity to imperfections in the actual physical hardware installation.

Alternative substructuring methods to the CMS methods are also applied to vibro-acoustic systems. Meyer [35] shows the Condensed Transfer Function (CTF) approach as a potential substructuring technique to predict vibro-acoustic behaviour of complex systems. It is a substructuring method based on the concept of subsystem condensed transfer functions that allows

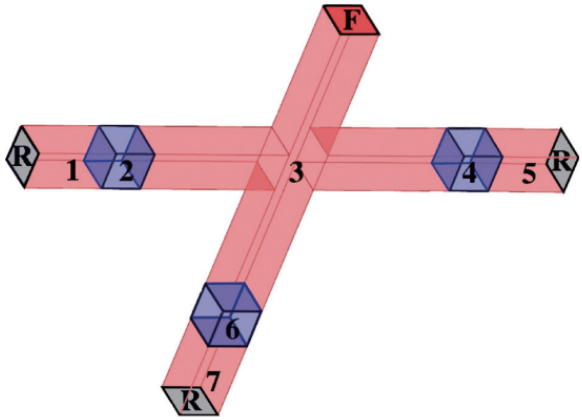

Figure 1.10: Cross 7-substructure model. Reference: Lindberg et al. [33]

assembling acoustical or mechanical subsystems coupled along lines or surfaces. The admittances along the coupling line are condensed to decrease the number of coupling forces that need to be calculated. Meyer validates the method for simple structural models. In [36], Meyer showcases three applications of the CTF approach for substructuring vibro-acoustic systems, going into depth on the selection of the condensation functions and convergence criterion for each case.

The CTF method is also investigated by Dumortier in [37], where he presents a reverse formulation of the CTF method for the decoupling of substructures. This formulation allows the prediction of the behavior of a subsystem that is part of a larger system, from the knowledge of the condensed transfer functions of the global system and of the residual subsystem that must be removed. This approach was applied and validated for the prediction of the purely acoustic behavior of a rigid sphere immersed in water.

Engelmann [15] proposes a novel method where vibro-acoustic modes can be obtained by employing their respective uncoupled modes. Integrate the modal substructuring method with non-conforming grids to facilitate the direct application of experimental and numerical modes in describing weakly coupled vibro-acoustic systems. By utilizing Delaunay triangulation within this framework, mechanical modes—originating from any source and discretized as point clouds—can be effectively linked with computational acoustic modes and the coupling matrix is therefore able to be determined. Validation of this approach is achieved through a two-sided vibro-acoustic box, where it was shown that systems utilizing coarsely discretized experimental modes can be effectively employed alongside computational modes when appropriate interpolation techniques are applied. This technique is limited for cases where the coupled mechanical displacements and pressures can be adequately represented as a superposition of weighted uncoupled modes, which is typical for low-density fluid mediums such as air.

2

Project Proposal

2.1. Research Gap

The previous chapter presented an overview of reduced order modelling techniques for vibro-acoustic systems. These systems are known to be challenging due to their complexity and the significant computational resources required for accurate simulations. The literature review highlights several strategies to address these challenges. These include simplifying the models through reduction techniques or analysing systems by dividing them into sub-components, which are then recombined. Table 2.1 provides a summary of the different coupling approaches introduced in Figure 1.7, along with references that have explored each type. In several of these cases, ASML already has established workflows. However, the literature does not yet offer a clear method for incorporating vibro-acoustic models into a broader multiphysics system, specifically the configuration referred to as Case D in Figure 1.7. This gap is important for ASML, as addressing it would make it possible to understand how specific vibro-acoustic interactions affect other components within a system. Integrating these models into ASML's modular modelling approach would greatly improve the speed and efficiency of performing and applying vibro-acoustic analyses in system-level simulations.

Table 2.1: Reference cases for the identified coupling approaches

Coupling Type	Case A	Case B	Case C	Case D
One-Way	[38, 39][1]	Current ASML	Current ASML	Gap
Two-way	[7, 31, 32]	[2, 23]	N.A.	Gap

2.2. Research Proposal

This research addresses a key gap: a streamlined approach for integrating two-way coupled vibro-acoustic models into ASML's existing structural modelling framework. The primary objective of this project is to develop a standardized workflow for creating such vibro-acoustic modules, enabling the inclusion of acoustic-structure interaction effects in other parts of ASML's machines. A central focus of this work is the development of Reduced Order Models (ROMs) for vibro-acoustic (sub)components. These ROMs must be compatible with ASML's modular modelling methodology used for dynamic system simulations. A key constraint is that the final model must be fully executable in MATLAB, without relying on external finite element method (FEM) software such as COMSOL or ANSYS. To achieve these goals, the following research questions

[1] A variation of Case A, not exactly as it is depicted in Figure 1.7.

will be addressed:

- Which reduction techniques are suitable for vibro-acoustic models while remaining compatible with ASML's modular modelling methodology?
- How to ensure the reduced models maintain compatible interfaces?
- What are the improvements in accuracy of the fully coupled vibro-acoustic system over the uncoupled acoustic-structure system?

Existing methods are first analysed as a starting point, and subsequently extended, as detailed further in Section 3.3.

2.3. Approach

The goal of this project is to develop a comprehensive modelling strategy that integrates three key procedures: model reduction, substructuring, and two-way coupled vibro-acoustic analysis. Each of these procedures presents its own complexities and challenges, which must be carefully addressed during the development process. To manage this, the project will begin with an Initial Investigation, where each modelling procedure will be implemented incrementally using a simplified academic model. This phase allows for focused testing and refinement of each method in isolation. Once the individual methodologies have been validated and optimised, they will be combined and applied in a Main Study. This final phase will involve implementing the complete modelling strategy on a real-world case study relevant to ASML, ensuring the approach is both practical and effective.

An insightful analysis initially planned for this project involved a validation study, comparing equivalent systems modelled using one-way and two-way vibro-acoustic coupling. The goal was to benchmark these models against either established reference results or available experimental data, thereby demonstrating the necessity of a two-way coupling strategy. A potential follow-up to this study could have been an investigation into how structural stiffness and fluid density influence the accuracy of one-way coupling. This would help identify the conditions under which a one-way approach provides a sufficiently accurate approximation. However, due to time constraints, these analyses could not be included within the scope of this project.

2.3.1. Initial Investigation

The objective of the Initial Investigation is to formulate the methodology for assembling a reduced vibro-acoustic module with a structural module. This begins analysing the simplest system satisfying these criteria, depicted in Figure 2.1(a). This system comprises a vibro-acoustic module with one structural and one acoustic domain sharing a common FSI, where the structural domain has a matching interface with an additional structural component. With this model, the following steps will be conducted:

I. Create and solve the unreduced vibro-acoustic and structural modules.

II. Develop a coupling procedure (physical/modal/state-space/...) to obtain an assembly of both modules.

III. Validate the obtained assembly with by comparing with the simulation of the whole system.

IV. Create reduced order models of both modules and repeat steps I - III.

These steps will be validated using a simple academic model. This model will act as a proof of concept, allowing for the exploration and comparison of different modelling approaches. As potential candidates for Case C, Figure 2.1 presents several possible configurations for the

academic model. Among these, the configuration shown in Figure 2.1(b) involves multiple FSIs, which could introduce significant challenges during both the model reduction and assembly phases. Additionally, it would require solving an extra acoustic-structure interaction problem between the two modules. To keep the validation focused and manageable, the academic model will instead follow the simpler configuration shown in Figure 2.1(a), as described in Section 3.1. This setup is better suited for implementing and testing the foundational aspects of the reduction procedure. It is important to note that the springs depicted in Figure 2.1 are symbolic and merely represent the interface connections between different components.

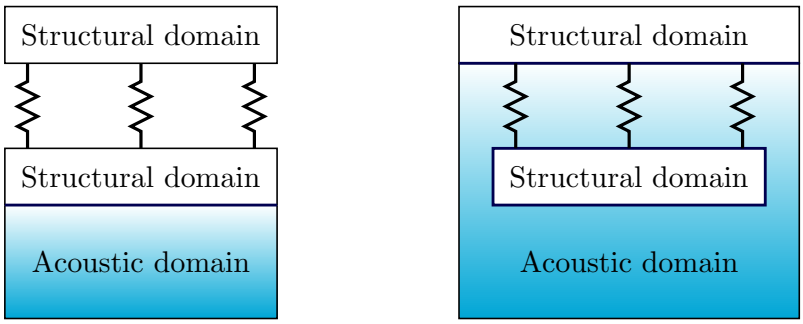

(a) Initial model (b) Interior model

Figure 2.1: Schematics of potential models of the Initial Investigation

2.3.2. Case Study

This phase of the project focuses on applying the validated modelling procedure to a case study relevant to ASML. While the exact application is still to be determined, it is expected to resemble the test-bench model shown in Figure 2.2. This model consists of five structural components surrounded by an acoustic domain and can be logically divided into vibro-acoustic and purely structural modules. It is important to note that the test-bench model does not include any actual ASML machine components. This ensures that the thesis does not disclose any confidential information, while still reflecting the physical conditions and challenges present in real ASML applications.

For the case selected in the Main Study, the same procedure outlined in the Initial Investigation will be followed. The first step involves creating both the fully coupled vibro-acoustic module and the structural module, as illustrated in Figure 2.3(a) and Figure 2.3(b), respectively. Next, the developed assembly procedure will be applied to combine these modules into the complete system, shown in Figure 2.3(c). Due to the increased complexity and size of this simulation, it will be carried out using the reduced-order models of the individual components. To ensure the accuracy of the assembly process, the full system will also be simulated using unreduced models for validation purposes.

As a potential final step for the project, the developed technique will be analysed with regards to its compatibility with ASML's modular modelling way-of-work. Therefore, it will be verified if any adjustments on how the coupling and model reductions are performed are required to ensure that it can be included as a new tool. Lastly, this also entails creating proper documentarian so other can easily use of this tool. If deemed necessary, modifications to the proposed model types or additional studies may be performed to further develop the proposed methodology.

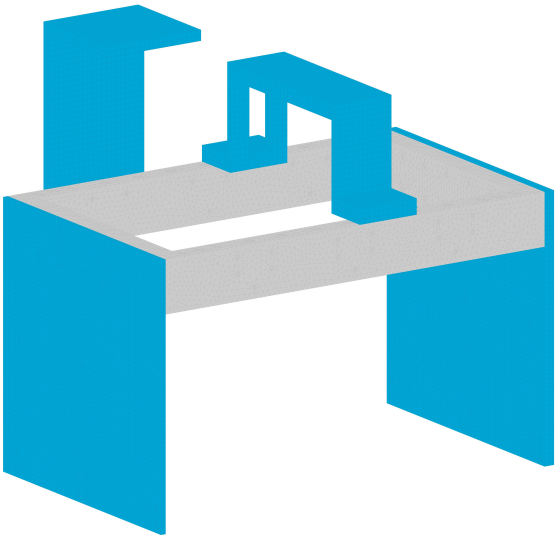

Figure 2.2: Test-bench full FEM model

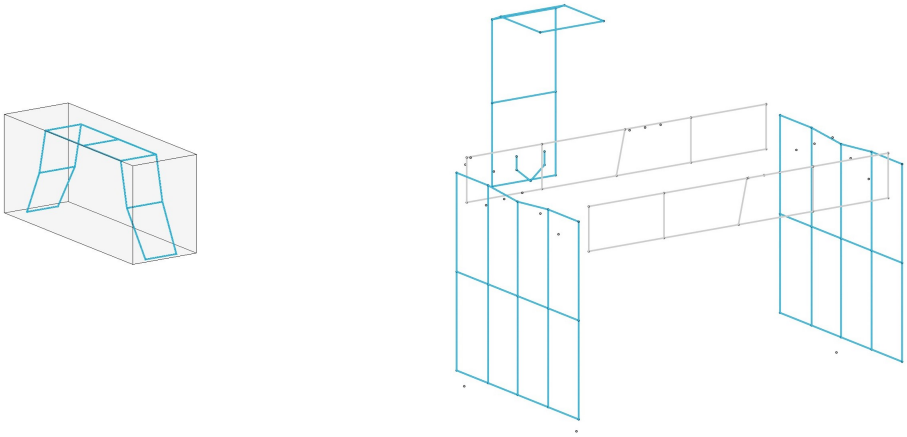

(a) Vibro-acoustic module

(b) Structural assembly

2.4. Planning

Based on the detailed methodology outlined in Section 2.3, the task distribution timeline is summarized in Table 2.2.

2.5. Research Strategy

The research strategy for this literature survey primarily utilized the ScienceDirect and Scopus databases. Boolean operators “AND” and “OR” were employed to refine results and ensure relevance. Initial broad searches with terms like “vibro-acoustic” or “acoustic-structure” were specified further by combining them with “reduced order modelling”, “reduced order models”, or “substructuring”, resulting in six specific search combinations. This search strategy yielded over twenty promising results, of which around half were deemed relevant. References from these were further investigated, especially where citations lacked full explanation or if they seemed pertinent/worthwhile. For several cases, this was an iterative process (looking at the references of the references, and so on). This literature review yielded a vast quantity of different research

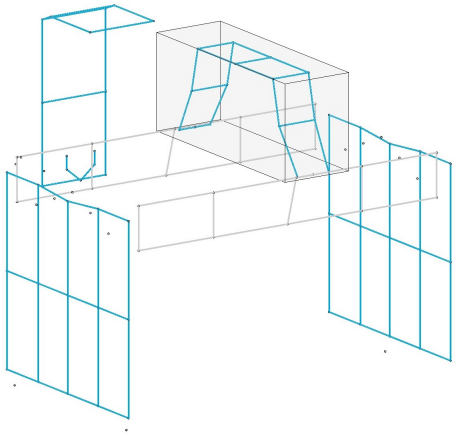

(c) Full system assembly

Figure 2.3: Configuration of the modules for the main study

Table 2.2: Proposed project timeline

Stage	Description	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul
Initial Investigation	Elaborate 2D edge case model								
	Develop Coupling Procedure								
	Reduce edge case model								
	Couple vibro-acoustic to structural modules (full & reduced)								
	Validate model with complete simulation								
Main Study	Define ASML Case Study								
	Create vibro-acoustic and structural SE modules								
	Couple modules								
	Model validation								
Thesis Elaboration	Write Thesis								
	ASML publication board								
	Prepare presentation								

articles, conference proceedings, book chapters and university theses that expose the current state of the art on vibro-acoustic ROM. Note that some literature potentially fit the search criteria, but was not open access or available with the TU Delft or ASML credentials, and therefore was not used.

Part II

Initial Investigation

3

Academic Model

The objective of the Initial Investigation is twofold: determine a methodology to reduce vibro-acoustic models and then assemble these models with other purely structural models. The aim is for this methodology to be compatible with ASML's modular modelling way of work. When validated, these techniques will then be applied in a Case Study that simulates an application of interest to ASML (Chapter 5). This chapter starts with a description of the first model analysed in the Initial Investigation, followed by the necessary steps to prepare this model for reduction, and finally an overview of the reduction techniques used.

3.1. Model Definition

As explained in the Project Proposal, the academic model analysed consists of the assembly of a two-dimensional vibro-acoustic component and a two-dimensional structural component. The created substructures are shown in Figure 3.1(a) and Figure 3.1(b). Additionally, a model that consists of the assembly of these two components was also created, to serve as a reference to the coupling procedure. The chosen software to create the FEM models of these components was COMSOL, due to it being the main commercial package used for acoustic simulations within ASML and for the capabilities of the LiveLink with MATLAB, which will be discussed in the coming sections.

The vibro-acoustic component of Figure 3.1(a) consists of a thin steel beam (1 x 0.02 m) that has an air volume coupled to its top edge (1 x 0.5 m). The beam and the acoustic cavity are coupled together with a FSI, while the remaining edges of the air volume have sound-hard boundaries applied. The structural component of Figure 3.1(b) consists of a steel frame that is fixed to the ground. The steel frame is 1 meter long and 0.5 meter high, with a constant thickness of 0.05 meter. The bottom edge of the steel beam of the vibro-acoustic component is to be coupled together with the top edge of the frame, resulting in the assembled system shown in Figure 3.1(c). Note that the diagrams of the models shown follow the colour convention described with Figure 1.1.

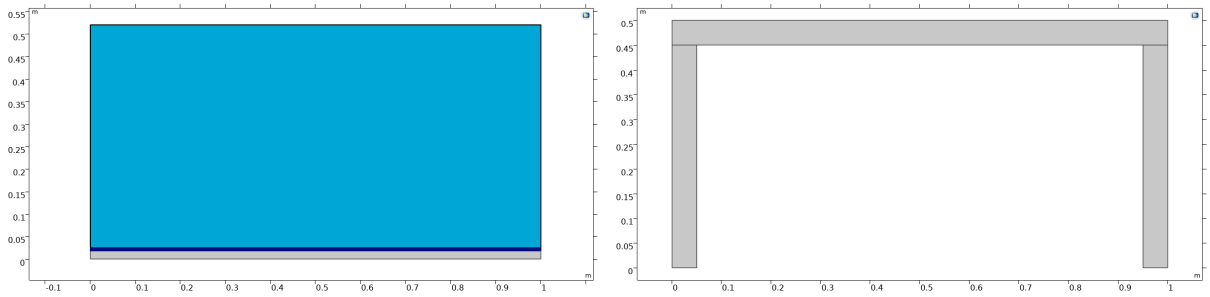

(a) Vibro-acoustic component

(b) Structural component

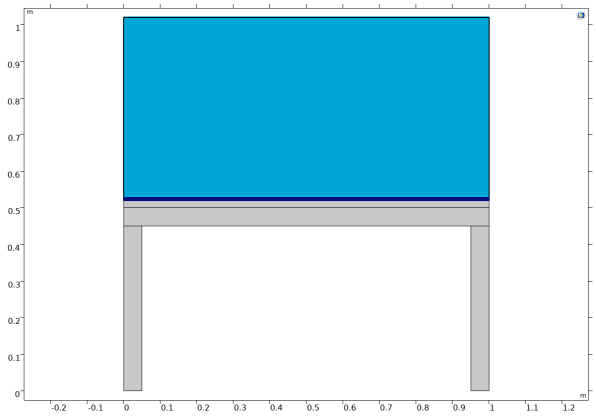

(c) Full system assembly

Figure 3.1: Configuration of the assembly

After the modelling of the components, the remainder of the analysis is done in MATLAB. This switch is required, since several matrix manipulations are required and the reduction bases need these matrices as inputs. For that, the COMSOL model `.mph` files were saved as functions in MATLAB `.m` files.

The COMSOL LiveLink with MATLAB enables these models to be accessed directly in MATLAB and offers non-native functions to manipulate these models. The `mphmatrix` function can be used to obtain the system matrices from the models, as seen in the examples below.

```
M = mphmatrix(model, 'sol1', 'out', {'E'}).E
K = mphmatrix(model, 'sol1', 'out', {'K'}).K
```

For the use of this function, the main input is the generated model function (`model` in this example) and the type of matrix of the output. `E` and `K` correspond to the mass and stiffness matrices respectively. In the case that there are fixed boundary conditions, which is the case for the structural component and the assembly models, it is important to note that this boundary condition is not automatically applied. To enforce it, there are different output matrices to choose, `Ec` and `Kc`, which are the mass and stiffness matrices without the rows and columns corresponding to the nodes that were fixed.

This procedure works, but it does not update the node indices of the matrix equations, which complicates the reordering of matrix equations needed for preconditioning and partitioning. Therefore, an alternative is to export the `E` and `K` and set the columns and rows corresponding to the fixed nodes to zero. The nodes that are fixed can be obtained by the output of the constraint Jacobian (output matrix `N`).

With the system matrices in MATLAB, the components can be prepared for reduction and assembly.

3.2. Model Preparation

3.2.1. Preconditioning

Vibro-acoustic systems, represented by Equation 1.1, are intrinsically ill-conditioned problems. This is due to the different physics having entries in the mass and stiffness matrices that may differ by 10 to 12 orders of magnitude [40]. The ill-conditioned matrices generate several numerical challenges to compute its inverse and to solve eigenvalue problems. These computations are required for the characterisation of the system and to create the reduction bases.

In order to reduce the condition number of the system matrices, and therefore improve the conditioning, it is important to precondition the system. A simple preconditioning method consists of performing a diagonal scaling [7, 40]. Since the initial model analysed does not have damping, Equation 1.1 is reduced to Equation 3.1 below.

$$\begin{bmatrix} \mathbf{M_s} & \mathbf{0} \\ \rho_a \mathbf{K_c}^\intercal & \mathbf{M_a} \end{bmatrix} \begin{Bmatrix} \ddot{\boldsymbol{x}}_\mathbf{s} \\ \ddot{\boldsymbol{x}}_\mathbf{a} \end{Bmatrix} + \begin{bmatrix} \mathbf{K_s} & -\mathbf{K_c} \\ \mathbf{0} & \mathbf{K_a} \end{bmatrix} \begin{Bmatrix} \boldsymbol{x}_\mathbf{s} \\ \boldsymbol{x}_\mathbf{a} \end{Bmatrix} = \begin{Bmatrix} \mathbf{0} \\ \mathbf{0} \end{Bmatrix} \tag{3.1}$$

The aim of the diagonal scaling is to scale the block matrices (structural, acoustic and coupling) in the mass and stiffness matrices such that they have the same Frobenius norm. This is accomplished by premultiplying Equation 3.1 with a matrix **B** and post multiplying by a matrix **A**, as seen in Equation 3.2. The following derivation is presented in the frequency domain for compactness

$$\underbrace{\begin{bmatrix} \mathbf{B}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 \end{bmatrix}}_{\mathbf{B}} \left(\begin{bmatrix} \mathbf{K_s} & -\mathbf{K_c} \\ \mathbf{0} & \mathbf{K_a} \end{bmatrix} - \boldsymbol{\omega}^2 \begin{bmatrix} \mathbf{M_s} & \mathbf{0} \\ \rho_0 \mathbf{K_c}^\intercal & \mathbf{M_a} \end{bmatrix} \right) \underbrace{\begin{bmatrix} \mathbf{A}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_2 \end{bmatrix}}_{\mathbf{A}} \begin{Bmatrix} \hat{\boldsymbol{x}}_\mathbf{s} \\ \hat{\boldsymbol{x}}_\mathbf{a} \end{Bmatrix} = \begin{Bmatrix} \mathbf{0} \\ \mathbf{0} \end{Bmatrix} \tag{3.2}$$

Note that $\mathbf{A}_i$ and $\mathbf{B}_i$ are, respectively, diagonal matrices with a_i and b_i in all entries. The multiplication in Equation 3.2 yields:

$$\left(\begin{bmatrix} a_1 b_1 \mathbf{K_s} & -a_2 b_1 \mathbf{K_c} \\ \mathbf{0} & a_2 b_2 \mathbf{K_a} \end{bmatrix} - \boldsymbol{\omega}^2 \begin{bmatrix} a_1 b_1 \mathbf{M_s} & \mathbf{0} \\ a_1 b_2 \rho_0 \mathbf{K_c}^\intercal & a_2 b_2 \mathbf{M_a} \end{bmatrix} \right) \begin{Bmatrix} \hat{\boldsymbol{x}}_\mathbf{s} \\ \hat{\boldsymbol{x}}_\mathbf{a} \end{Bmatrix} = \begin{Bmatrix} \mathbf{0} \\ \mathbf{0} \end{Bmatrix} \tag{3.3}$$

As discussed previously, the ill-conditioned matrices originate from the fact that the orders of magnitude vary considerably between the entries of each physics. Therefore, the coefficients a_i and b_i are chosen in such a way that ensures that the Frobenius norm of each block matrix is the same, as seen in Equation 3.4.

$$\|a_1 b_1 \mathbf{K_s}\| = \|a_2 b_1 \mathbf{K_c}\| = \|a_2 b_2 \mathbf{K_a}\| \tag{3.4a}$$
$$\|a_1 b_1 \mathbf{M_s}\| = \|a_1 b_2 \mathbf{M_a}\| = \|a_1 b_2 \rho_0 \mathbf{K_c}\| \tag{3.4b}$$

From Equation 3.4 it can be observed that there is no unique solution. Therefore, for simplicity it can be chosen that $a_1 = b_1 = 1$, as done by [40]. By doing so, each line of Equation 3.4 yields one possible value for a_2 and b_2, as seen in Equation 3.5.

$$a_{2_1} = \frac{\|\mathbf{K_s}\|}{\|\mathbf{K_c}\|} \tag{3.5a}$$
$$a_{2_2} = \frac{\|\rho_0 \mathbf{K_c}\|}{\|\mathbf{M_a}\|} \tag{3.5b}$$
$$b_{2_1} = \frac{\|\mathbf{K_c}\|}{\|\mathbf{K_a}\|} \tag{3.5c}$$
$$b_{2_2} = \frac{\|\mathbf{M_s}\|}{\|\rho_0 \mathbf{K_c}\|} \tag{3.5d}$$

Therefore, the final values for a_2 and b_2 can be chosen to be the geometrical average of each value, as seen in Equation 3.6.

$$a_2 = \sqrt{a_{2_1} \cdot a_{2_2}} \tag{3.6a}$$
$$b_2 = \sqrt{b_{2_1} \cdot b_{2_2}} \tag{3.6b}$$

This preconditioning procedure does not affect the eigenvalues and eigenvectors of the system, as only the generalised eigenvalue problem of Equation 3.7 is of interest, not the eigenvalues and eigenvectors of each system matrix individually.

$$(\mathbf{K} - \lambda \mathbf{M})\, \boldsymbol{v} = 0 \tag{3.7}$$

Prepare the vibro-acoustic model for preconditioning

A requirement of this preconditioning method is that the vibro-acoustic system needs to be written out in the $\boldsymbol{u} - \boldsymbol{p}$ formulation, which is not the standard output in most FEM packages. With the COMSOL LiveLink with MATLAB, to obtain the the system matrices in the $\boldsymbol{u} - \boldsymbol{p}$ formulation, the `mphxmeshinfo` function can be used to extract a structure array with information about the extended mesh of the model. Naming this structure `mesh_info`, it can be obtained by applying this function to the MATLAB version of the COMSOL model (`model.m`):

```
mesh_info = mphxmeshinfo(model.m).
```

The different DoFs types that are present in the model have an assigned index that can be checked at `mesh_info.dofs.dofnames`. For the 2D vibro-acoustic model of Figure 3.1(a) this array contains pressure (`comp1.p`) and displacements (`comp1.u` and `comp1.v`) DoFs. Finally, to match the indexes of the DoFs with the matrix equations through `mesh_info.dofs.nameinds`. For the vibro-acoustic module from Figure 3.1(a), Figure 3.2(a) shows the non-zero entries of the original output mass matrix from COMSOL and Figure 3.2(b) shows the transformed matrix in the $\boldsymbol{u} - \boldsymbol{p}$ formulation.

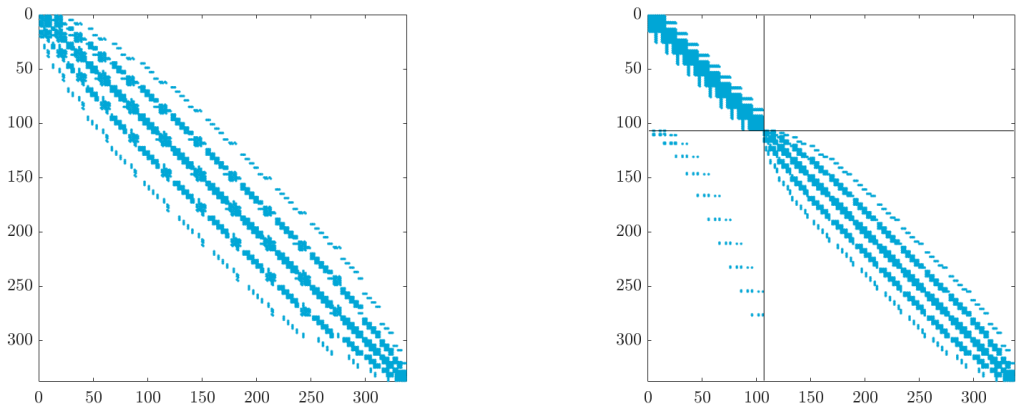

(a) Original mass matrix

(b) Mass matrix in the $\boldsymbol{u} - \boldsymbol{p}$ formulation

Figure 3.2: Manipulation of the mass matrix of the vibro-acoustic module from Figure 3.1(a)

Therefore the systems matrices of the vibro-acoustic component and the assembled system, shown in Figure 3.1(a) and Figure 3.1(c), need to be preconditioned after being imported with the `mphmatrix` function. Once the matrices are preconditioned, they can be used to obtain the natural frequencies and modes with the `eig` function.

3.2.2. Partitioning

Reduction by CMS, as explained in Section 1.3.1, requires the DoFs of the system to be divided into boundary and internal DoFs. This division is called partitioning. Using CMS is very practical

for a modular approach, as the boundary DoFs are maintained in the reduction process, which enables the interfaces of the reduced substructures to be coupled to one another.

Similar to the manipulation of the system matrices performed on the preconditioning to get to the $\boldsymbol{u} - \boldsymbol{p}$ formulation, the order of the matrix equations need to be altered to perform the partitioning. To select the boundary DoFs, it is possible to filter the nodes at specific coordinates with the same structure with extended mesh information (`mesh_info.dofs.coords`), however, this method is not robust for complex interfaces. A solution is to create a named selection of the interface in COMSOL.

There is not a direct way to get the corresponding matrix equation number for the nodes of the named selection with the COMSOL LiveLink with MATLAB. However, it can be done in two steps:

1. Apply the `mpheval` function to get the coordinates of the nodes of the named selection:

   ```
   interface_coords = mpheval(model, {'x','y','z'}, 'selection','sel1','refine',2).p,
   ```

   note that multiple selections can be used for different interfaces, this step would only need to be repeated for each of the n selections (`sel1`,`sel2`,...,`sel`n). Additionally, the `'refine'` number refers to the interpolation between nodes. Note that LiveLink with MATLAB only supports Quadratic Lagrange interpolation. If your model uses Quadratic Serendipity elements, `mpheval` will create the missing interior DOFs for the original selection, which are excluded from the system matrices.

2. After the coordinates of the nodes that belong to the selection are found, MATLAB's `find` and `ismembertol` functions can be used between these coordinates and the full node coordinates (`mesh_info.dofs.coords`). Since the coordinates in `mesh_info.dofs.coords` follow the order of matrix equations, it is possible to determine the node position through the found indices.

Once all the boundary DoFs are found, they can be grouped together to partition the system matrices as seen in Equation 1.5. Figure 3.3 shows the partitioned mass matrix of the vibro-acoustic module from Figure 3.1(a).

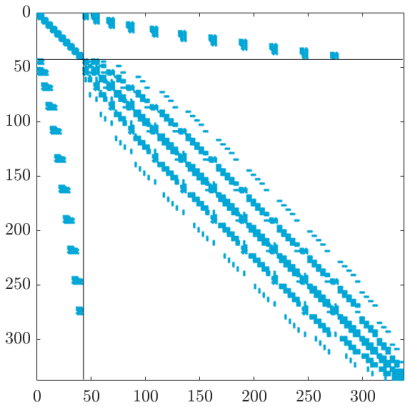

Figure 3.3: Partitioned mass matrix

The boundary set are typically all the DoFs of the nodes at the interfaces between the substructures or modules that are ment to be assembled together. However, any other DoF that is of interest to be maintained in the reduced order model can be included in this set. This is not done for this academic model, but is performed in Chapter 5 in the Case Study.

3.2.3. Assembly of Substructures

Once the system matrices are partitioned they can be assembled together. The assembly procedure is the same for full or reduced matrices, as the coupling is done only at the interfaces. If non interface DoFs are included in the boundary set, as they are to be maintained in the reduced model, they should not be used to construct the Boolean matrices **B** and **L** from Equation 1.6.

To assemble substructures, it is important that, when partitioning, their interface DoFs are grouped together in a matching way. This ensures that the compatibility matrix **B** can be written out as seen in Equation 3.8. Otherwise, the compatibility matrix has to be constructed by finding the nodes of the different substructures that have the same coordinates.

$$\mathbf{B} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{I} & \mathbf{0} \end{bmatrix} \tag{3.8}$$

With the compatibility matrix **B**, a convenient way to obtain the localisation matrix **L** is to find the null space of **B** [14]. To ensure that the outcome is a matrix with only 1's for the non zero entries, it is important to add the rational option in the MATLAB function, as seen below.

```
L = null(B,'rational')
```

Using the localisation matrix, the primal assembly of substructures can now be achieved. Primal assembly is similar to FEM assembly, ensuring compatibility of primal variables between the DoFs at the interfaces, followed by the removal of redundant sets. This technique is preferred for vibro-acoustic systems [7]. To assemble substructures using this method, only the localisation matrix **L** and the block-diagonal system matrices of each substructure are required, as demonstrated in [14]. Equation 3.9 illustrates the procedure for the mass matrix, which is identical for the stiffness and damping matrices.

$$\mathbf{M}_{asm} = \mathbf{L}^{\intercal} \begin{bmatrix} \mathbf{M}_1 & & \\ & \ddots & \\ & & \mathbf{M}_n \end{bmatrix} \mathbf{L} \tag{3.9}$$

In Chapter 5, an alternative coupling method is introduced. This method does not involve the union of different substructures and aligns more closely with ASML's way of work. However, for an initial analysis, the more traditional primal assembly, which has been proven effective for vibro-acoustic systems, is preferred.

3.3. Vibro-Acoustic Reduction

As stated previously, the main challenge when working with vibro-acoustic systems is the computational cost required to solve them. CMS allows for a reduction that maintains a partial physical interpretation of the models, which maintains the modularity required to be implemented at ASML. Three different CMS based vibro-acoustic reduction methods are presented in this section.

3.3.1. Two-Base Reduction

Maess et al. [40] and Herrmann et al. [31, 32] propose and apply a Craig-Bampton based CMS method for vibro-acoustic systems. This method involves creating separate reduction bases for the structural and acoustic elements. The coupling terms are then reduced using both reduction bases. The primary advantage of having separate reduction bases for each physical domain is that the matrices become symmetric, significantly improving the efficiency of computing the required modes. This technique will be referred to as Two-Base Reduction in this thesis. An overview of

this method is presented below, followed by a proposed extension using a Hintz-Herting inspired reduction base.

Craig-Bampton

The classic Craig-Bampton (CB) method uses a reduction base composed of constraint modes $\boldsymbol{\Psi}$ and fixed interface normal modes $\boldsymbol{\Phi}$, therefore the reduction base from Equation 1.4 can be represented as seen in Equation 3.10 below.

$$\begin{aligned} \mathbf{T} &= \begin{bmatrix} \boldsymbol{\Psi} & \boldsymbol{\Phi} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \boldsymbol{\psi} & \boldsymbol{\phi_i} \end{bmatrix} \end{aligned} \tag{3.10}$$

The Two-Base Reduction applies the classic CB method separately to the structural and acoustic DoFs. To do this, the system needs to be written in the $\boldsymbol{u}$-$\boldsymbol{p}$ formulation and the partitioning needs to be done individually for each of the different physics. Therefore, the general partitioning from Equation 1.3 will be updated to the form Equation 3.11 shown below.

$$\begin{Bmatrix} \boldsymbol{x_b} \\ \boldsymbol{x_i} \end{Bmatrix} = \begin{Bmatrix} \boldsymbol{u_b} & \boldsymbol{u_i} & \boldsymbol{p_b} & \boldsymbol{p_i} \end{Bmatrix}^{\mathsf{T}} \tag{3.11}$$

With the updated partitioning, the two reduction bases, $\mathbf{T_s}$ and $\mathbf{T_a}$, can be obtained as shown in Equation 3.12. In this context, $\boldsymbol{u}_b$ and $\boldsymbol{p}_b$ represent the retained boundary structural and acoustic DoFs, respectively. Meanwhile, $\boldsymbol{\eta_i}$ and $\boldsymbol{\zeta_i}$ denote the modal participation factors of the normal modes.

$$\begin{Bmatrix} \boldsymbol{u_b} \\ \boldsymbol{u_i} \end{Bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \boldsymbol{\psi_s} & \boldsymbol{\phi_{i,s}} \end{bmatrix}}_{\mathbf{T_s}} \begin{Bmatrix} \boldsymbol{u}_b \\ \boldsymbol{\eta_i} \end{Bmatrix} \tag{3.12a}$$

$$\begin{Bmatrix} \boldsymbol{p_b} \\ \boldsymbol{p_i} \end{Bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \boldsymbol{\psi_a} & \boldsymbol{\phi_{i,a}} \end{bmatrix}}_{\mathbf{T_a}} \begin{Bmatrix} \boldsymbol{p}_b \\ \boldsymbol{\zeta_i} \end{Bmatrix} \tag{3.12b}$$

The constraint modes and fixed interface normal modes can be calculated as the standard CB, shown in Equation 3.13 below. The full demonstration can be consulted in [11].

$$\boldsymbol{\psi} = -\mathbf{K}_{ii}^{-1}\mathbf{K}_{ib} \tag{3.13a}$$

$$\left(\mathbf{K}_{ii} - \boldsymbol{\omega}_{ii}^2\mathbf{M}_{ii}\right)\boldsymbol{\phi_i} = \mathbf{0} \tag{3.13b}$$

With the two reduction bases, it is now possible to obtain the reduced system matrices by applying Equation 1.5. The reduction of each block of the mass matrix is illustrated in Equation 3.14, and the same logic applies to the damping and stiffness matrices. The structural mass and stiffness are reduced using the structural reduction matrix $\mathbf{T_s}$, while the acoustic mass

and stiffness are reduced using the acoustic reduction matrix $\mathbf{T_a}$. The coupling terms are reduced using both reduction matrices.

$$\tilde{\mathbf{M}}_\mathbf{s} = \mathbf{T_s}^\intercal \mathbf{M_s} \mathbf{T_s} \tag{3.14a}$$

$$\tilde{\mathbf{M}}_\mathbf{a} = \mathbf{T_a}^\intercal \mathbf{M_a} \mathbf{T_a} \tag{3.14b}$$

$$\tilde{\mathbf{K}}_\mathbf{c} = \mathbf{T_a}^\intercal \mathbf{K_c} \mathbf{T_s} \tag{3.14c}$$

Hintz-Herting

The Hintz-Herting reduction is a widely used CMS method within ASML, as it often yields better results for the dynamic analysis of models. This is because HH employs the same constraint modes as the CB method, but replaces the fixed-interface normal modes with a combination of inertial relief modes ($\mathbf{H}$) and free-body modes ($\mathbf{\Phi^*}$). The inertial relief modes are used to mitigate the inertia effects caused by static loads, ensuring that the dynamic analysis focuses on the vibrational characteristics. The classic HH reduction basis is presented in Equation 3.15.

$$\begin{aligned} \mathbf{T} &= \begin{bmatrix} \mathbf{\Psi} & \mathbf{H} & \mathbf{\Phi^*} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{\psi} & \mathbf{H}_{i0} & \boldsymbol{\varphi_i} - \boldsymbol{\psi}\boldsymbol{\varphi_b} \end{bmatrix} \end{aligned} \tag{3.15}$$

Given the potential for improved results in ASML's applications, the Hintz-Herting (HH) reduction basis was integrated into the existing Two-Base Reduction technique developed by Maess et al. [40] and Herrmann et al. [31, 32]. The procedure remains as previously detailed, but now incorporates the different modes specific to the HH reduction basis, as illustrated in Equation 3.16.

$$\begin{Bmatrix} \boldsymbol{u_b} \\ \boldsymbol{u_i} \end{Bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{\psi_s} & \mathbf{H}_{i0,s} & \boldsymbol{\varphi_{i,s}} - \boldsymbol{\psi_s}\boldsymbol{\varphi_{b,s}} \end{bmatrix}}_{\mathbf{T_s}} \begin{Bmatrix} \boldsymbol{u}_b \\ \boldsymbol{\delta_0} \\ \boldsymbol{\delta} \end{Bmatrix} \tag{3.16a}$$

$$\begin{Bmatrix} \boldsymbol{p_b} \\ \boldsymbol{p_i} \end{Bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{\psi_a} & \mathbf{H}_{i0,a} & \boldsymbol{\varphi_{i,a}} - \boldsymbol{\psi_a}\boldsymbol{\varphi_{b,a}} \end{bmatrix}}_{\mathbf{T_s}} \begin{Bmatrix} \boldsymbol{p}_b \\ \boldsymbol{\delta}_\mathbf{0}^* \\ \boldsymbol{\delta}^* \end{Bmatrix} \tag{3.16b}$$

In this case, $\boldsymbol{u_b}$ and $\boldsymbol{p_b}$ remain the retained boundary structural and acoustic DoFs. Furthermore, $\boldsymbol{\delta_0}$ and $\boldsymbol{\delta}$ are modal coordinates that represent, respectively, the relative displacement of the modes and the rigid body acceleration in terms of the free-body modes. The asterisks next to the deltas in the acoustic basis are used to differentiate between the structural and acoustic coordinates. The inertia relief modes can be calculated using the standard HH method, as shown in Equation 3.17 below. The free-body modes are obtained by calculating the eigenvectors $\boldsymbol{\varphi}$ of the system and then separating them into boundary and internal DoFs, $\boldsymbol{\varphi b}$ and $\boldsymbol{\varphi i}$, respectively. The complete demonstration is provided in [12].

$$\mathbf{H}_{i0} = -\mathbf{K}_{ii}^{-1} \left(\mathbf{M}_{ib} + \mathbf{M}_{ii} \boldsymbol{\psi} \right) \boldsymbol{\phi_{b0}} \tag{3.17}$$

Improvement on the Two-Base Reduction

After the preconditioning step, the coupling matrices $\mathbf{K_c}$ differ between the mass and stiffness matrices due to scaling with different parameters. Consequently, an enhancement to the Two-Base Reduction technique involves separately reducing these coupling matrices. Denoting $\mathbf{K_{c_1}}$ and $\mathbf{K_{c_2}}$ as the coupling matrices from the mass matrix and stiffness matrix, respectively. Equation 3.18 illustrates the procedure for computing the reduced coupling matrices independently.

$$\tilde{\mathbf{K}}_{\mathbf{c}1} = \mathbf{T_a}^\intercal \mathbf{K_{c_1}} \mathbf{T_s} \tag{3.18a}$$

$$\tilde{\mathbf{K}}_{\mathbf{c}2} = \mathbf{T_s}^\intercal \mathbf{K_{c_2}} \mathbf{T_a} \tag{3.18b}$$

3.3.2. Nonsymmetric Reduction

A second reduction strategy for vibro-acoustic systems involves directly applying the classic Craig-Bampton and Hintz-Herting CMS reduction techniques. This approach was not found in use during the literature review. For this strategy, the classic ***u*-*p*** formulation is not suitable, as it requires grouping all the boundary displacement and pressure DoFs together, as shown in Equation 3.19. In this equation, the CB (3.19) and HH (3.19a) methods are applied to the partitioned system encompassing both physics. Note that for the case analyzed in Figure 3.1, the ***u*-*p*** formulation is maintained since there are no boundary pressure DoFs.

$$\left\{ \begin{bmatrix} \boldsymbol{u_b} & \boldsymbol{p_b} \end{bmatrix} \quad \begin{bmatrix} \boldsymbol{u_i} & \boldsymbol{p_i} \end{bmatrix} \right\}^\intercal = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \boldsymbol{\psi} & \boldsymbol{\phi_i} \end{bmatrix}_{L,R} \left\{ \boldsymbol{u_b} \quad \boldsymbol{p_b} \quad \boldsymbol{\eta_i} \quad \boldsymbol{\eta_i} \right\}^\intercal \tag{3.19a}$$

$$\left\{ \begin{bmatrix} \boldsymbol{u_b} & \boldsymbol{p_b} \end{bmatrix} \quad \begin{bmatrix} \boldsymbol{u_i} & \boldsymbol{p_i} \end{bmatrix} \right\}^\intercal = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \boldsymbol{\psi} & \mathbf{H}_{i0} & \boldsymbol{\varphi_i} - \boldsymbol{\psi}\boldsymbol{\varphi_b} \end{bmatrix}_{L,R} \left\{ \begin{bmatrix} \boldsymbol{u_b} & \boldsymbol{p_b} \end{bmatrix} \quad \boldsymbol{\delta_0} \quad \boldsymbol{\delta} \right\}^\intercal \tag{3.19b}$$

A notable computational disadvantage of using this approach is that the matrices remain nonsymmetric. Consequently, the eigenvalue problems that need to be solved for calculating the fixed-interface normal modes for the CB method, as well as the inertia relief and free-body modes for the HH method, result in different left and right eigenvectors. Therefore, two separate reductions must be computed: one using the right eigenvectors ($\mathbf{T}_R$) and one using the left eigenvectors ($\mathbf{T}_L$). To obtain the reduced system matrices with these two reduction bases, Equation 1.5 must be modified by pre-multiplying by the transpose of $\mathbf{T}_L$ and post-multiplying by $\mathbf{T}_R$ [23], as shown in Equation 3.20.

$$\tilde{\mathbf{M}} = \mathbf{T}_L^\intercal \, \mathbf{M} \, \mathbf{T}_R \tag{3.20a}$$

$$\tilde{\mathbf{C}} = \mathbf{T}_L^\intercal \, \mathbf{C} \, \mathbf{T}_R \tag{3.20b}$$

$$\tilde{\mathbf{K}} = \mathbf{T}_L^\intercal \, \mathbf{K} \, \mathbf{T}_R \tag{3.20c}$$

3.3.3. Locally Symmetric CB Method

The third and final CMS based technique analysed to reduce vibro-acoustic system is the Locally Symmetric CB method proposed by Tabak [7]. This method consists of performing numerical

operations to convert the originally nonsymmetric vibro-acoustic system matrices from Equation 1.1 into a symmetric form, shown in Equation 3.22. A key benefit of this technique is that it enables the use of symmetric solvers to compute the components of the CB basis at the substructure level. Additionally, the reduced symmetric matrices retain the desirable sparsity patterns typically associated with standard Craig-Bampton reductions.

To transform Equation 1.1 into a symmetric formulation, Tabak [7] proposes two approaches. The first involves applying a scaling matrix to the original system, known as the κ formulation. The second relies on a coordinate transformation, referred to as the τ formulation. This thesis adopts the κ formulation, as it retains the original DoFs of the system. The scaling matrix associated with this formulation is presented in Equation 3.21.

$$\kappa = \begin{bmatrix} \mathbf{I} & \mathbf{K_c}{\mathbf{K_a}}^{-1} \\ \mathbf{0} & \mathbf{M_a}{\mathbf{K_a}}^{-1} \end{bmatrix} \tag{3.21}$$

As previously explained, the scaling matrix is used to pre-multiply the original non-symmetric matrices from Equation 1.1. As a result, the symmetric matrices obtained from the multiplication of Equation 3.21 with Equation 1.1 are presented in Equation 3.22.

$$\begin{bmatrix} \mathbf{M_s} + \mathbf{K_c}{\mathbf{K_a}}^{-1}{\mathbf{K_c}}^{\intercal} & \mathbf{K_c}{\mathbf{K_a}}^{-1}\mathbf{M_a} \\ \mathbf{M_a}{\mathbf{K_a}}^{-1}{\mathbf{K_c}}^{\intercal} & \mathbf{M_a}{\mathbf{K_a}}^{-1}\mathbf{M_a} \end{bmatrix} \begin{Bmatrix} \ddot{\boldsymbol{u}} \\ \ddot{\boldsymbol{p}} \end{Bmatrix} + \begin{bmatrix} \mathbf{K_s} & \mathbf{0} \\ \mathbf{0} & \mathbf{M_a} \end{bmatrix} \begin{Bmatrix} \boldsymbol{u} \\ \boldsymbol{p} \end{Bmatrix} = \kappa \begin{Bmatrix} \mathbf{F_s} \\ \mathbf{F_a} \end{Bmatrix} \tag{3.22}$$

With this symmetric matrix formulation, the CMS reduction base can be computed as:

$$\mathbf{T} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ -\mathbf{K}_{s,ii}^{-1}\mathbf{K}_{s,ib} & \mathbf{0} & \boldsymbol{\Phi}_{sym,ii,s} \\ \mathbf{0} & -\mathbf{M}_{f,ii}^{-1}\mathbf{M}_{f,ib} & \boldsymbol{\Phi}_{sym,ii,s} \end{bmatrix}, \tag{3.23}$$

where the first two columns correspond to the constraint modes and the last column corresponds to the fixed interface normal modes.

3.4. Reduction with Damping

The CMS reduction techniques discussed in Section 3.3 were presented without accounting for damping effects in the system. However, modelling damping in the acoustic domain is particularly important for ASML, as the multiple acoustic reflections inside a confined space are inherently effected by acoustic properties of the boundary surfaces. The need of accounting for a finite impedance is merely driven by the ability to predict the impact of it onto acoustic propagation inside the machine [41].

Damping is typically modelled using an acoustic impedance boundary condition. Acoustic impedance plays a critical role in the transfer of acoustic energy between two media. Sound waves are partially reflected and partially transmitted at boundaries between media with different acoustic properties. When a sound wave encounters such an interface, a portion of the energy is reflected back, while the remainder is transmitted into the second medium. The ratio of reflected to transmitted energy is governed by the acoustic impedances of the two media. Although the wave experiences dissipation due to thermal and viscous losses as it enters the material, these effects are neglected in the present analysis. Figure 3.4 shows an impedance boundary condition applied

to the top edge of the vibro-acoustic system analysed in the Initial Investigation, substituting the existing sound hard boundary.

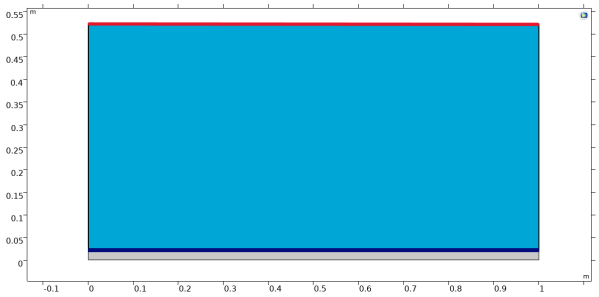

Figure 3.4: Introduction of impedance boundary condition

The impedance value applied to this boundary is set to twice the characteristic impedance of the acoustic domain, as shown in Equation 3.24. Although this value is chosen arbitrarily, it is selected to be high enough to avoid complete transmission of acoustic waves, yet not so high that it causes nearly total reflection.

$$\begin{aligned} Z &= 2 \cdot Z_{\mathrm{a}} \\ &= 2 \cdot \rho_{\mathrm{a}} \cdot c \end{aligned} \tag{3.24}$$

Damping is not applied to the structural component, as ASML typically incorporates modal damping, which can be done at a later stage in the process and ensures that it does not interfere with the reduction procedure. As a result, acoustic damping has no impact on the preconditioning of the system matrices, since the damping matrix in the vibro-acoustic model contains only a single non-zero block associated with the acoustic domain. Figure 3.5 presents the damping matrix in the $\boldsymbol{u}$–$\boldsymbol{p}$ formulation. As shown, non-zero entries are limited to the acoustic block, specifically in the elements that interface with the applied impedance. The matrix is also notably sparse, reflecting the localised nature of the impedance boundary condition applied.

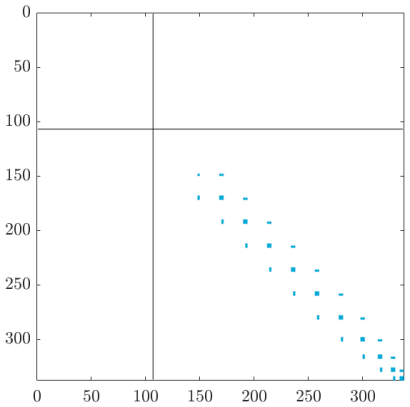

Figure 3.5: Damping matrix in the $\boldsymbol{u} - \boldsymbol{p}$ formulation

3.4.1. Modifications Generated with Damping

To compute the reduction bases discussed in Section 3.3 with the inclusion of damping, a few modifications are required in the process used to calculate the modes that form the reduction

basis.

- In the two-base reduction approach, the structural reduction basis $\mathbf{T_s}$ is computed in the same way as before. However, for the acoustic reduction basis $\mathbf{T_a}$, while the constraint modes remain unchanged, since they are static approximations that neglect all time derivatives, the additional interface modes must now be computed differently. Specifically, instead of using the standard `eig` function, the `polyeig` function must be used. This change is necessary because the eigenvalue problem has shifted from the linear form shown in Equation 3.7 to the non-linear form presented in Equation 3.25.

$$\left(\lambda^2\mathbf{M} + \lambda\mathbf{C} + \mathbf{K}\right)\boldsymbol{v} = 0 \tag{3.25}$$

- Unlike the `eig` function, the `polyeig` function does not return left eigenvectors as part of its output. Therefore, to obtain the left eigenvectors required for the interface modes in the non-symmetric reduction method, a second eigenvalue problem must be solved. This involves using the transposed system matrices, as shown in Equation 3.26, to compute the left eigenvectors corresponding to the same eigenvalues.

$$\left(\lambda^2\mathbf{M}^\intercal + \lambda\mathbf{C}^\intercal + \mathbf{K}^\intercal\right)\boldsymbol{v} = 0 \tag{3.26}$$

- The `polyeig` function is also used to compute the natural frequencies of the reduced system. Equation 3.27 shows how to convert the resulting eigenvalues into frequencies in hertz. This conversion is necessary because the eigenvalues from Equation 3.7 represent squared frequencies (in radians per second), while the eigenvalues from Equation 3.25 correspond to frequencies (in radians per second) multiplied by the imaginary unit.

$$f_{\text{eig}} = \frac{\sqrt{\lambda}}{2\pi} \tag{3.27a}$$

$$f_{\text{polyeig}} = \frac{-\mathrm{i}\lambda}{2\pi} \tag{3.27b}$$

This academic model will be analysed with and without damping.

3.5. Summary of Procedure

This chapter introduced the initial model used in this study and, more importantly, outlined the essential steps required to prepare the model for reduction. It also presented the three different reduction techniques evaluated for simplifying vibro-acoustic systems. The key steps involved in reducing and coupling a vibro-acoustic substructure with a structural substructure are summarized below.

1. Create COMSOL FE models of each substructure;
2. Export the COMSOL `.mph` files into MATLAB `.m` files;
3. Use the LiveLink to extract the system matrices;
4. Transform the system matrices into the $\boldsymbol{u} - \boldsymbol{p}$ formulation and perform the preconditioning of the matrices;
5. Partition the boundary and internal DoFs of the preconditioned system matrices;
6. Reduce the system with one of the CMS methods;
7. Assemble the substructures with the primal assembly.

4

Preliminary Results

In the previous chapter, the First Model and the various reduction methods to be analysed were introduced. In this chapter, the full substructures will first be coupled using the primal assembly technique to verify the accuracy of the coupling procedure. Subsequently, the reduction techniques will be applied to the individual components, and the results of assembling the reduced components will be compared with those obtained from assembling the corresponding full components. To conclude, the results obtained using the different techniques will be compared, and the key findings of the procedure will be summarized.

4.1. Full Assembly

The first analysis consists of coupling the full substructures, which are presented in Figure 3.1(a) and Figure 3.1(b), using the primal assembly method described in Equation 3.9. The resulting coupled model is compared with the reference model shown in Figure 3.1(c). This comparison is based on the natural frequencies and the mode shapes of the two systems. To assess the natural frequencies, the relative error between the frequency vectors $\boldsymbol{\varepsilon}$ is calculated, as explained in Equation 4.1.

$$\boldsymbol{\varepsilon} = \frac{|\boldsymbol{f} - \boldsymbol{f_{ref}}|}{\boldsymbol{f_{ref}}} \tag{4.1}$$

The mode shapes are compared using the Modal Assurance Criterion, which is a statistical measure that evaluates the similarity between two mode shapes ($\boldsymbol{\chi}_i$ and $\boldsymbol{\chi}_j$) . According to Equation 4.2, the MAC is obtained by dividing the squared magnitude of the dot product of the mode shapes by the product of their squared norms. The result is a value between zero and one, where a value of one indicates perfect correlation, meaning the mode shapes are identical, and a value of zero indicates no correlation.

$$\mathrm{MAC}_{ij} = \frac{\left|\boldsymbol{\chi}_i^\intercal \boldsymbol{\chi}_j\right|^2}{\left(\boldsymbol{\chi}_i^\intercal \boldsymbol{\chi}_i\right)\left(\boldsymbol{\chi}_j^\intercal \boldsymbol{\chi}_j\right)} \tag{4.2}$$

From Equation 4.2, it can be observed that all mode shapes are compared with one another. As a result, the output is a matrix in which each entry represents the MAC value between a pair of modes. Ideally, the main diagonal of this matrix, where the n-th mode of the assembled model is compared with the n-th mode of the reference model, should contain values equal to one, indicating perfect correlation. The other entries in the matrix are not necessarily zero, since there may be some level of correlation between different modes.

After applying the primal assembly technique, the relative error in natural frequencies between the assembled model and the reference model is presented in Figure 4.1.

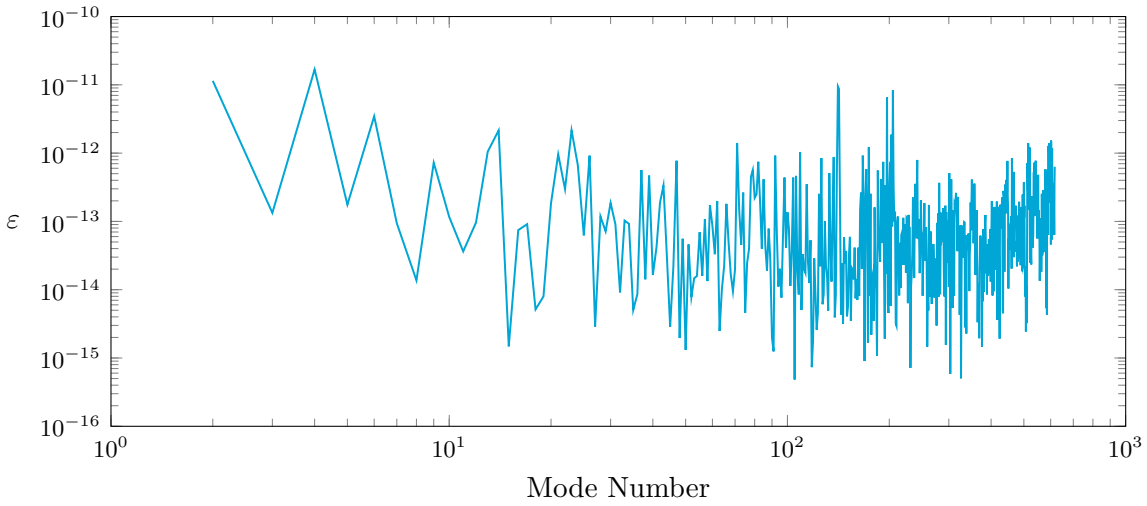

Figure 4.1: Comparison between CB and HH reduction

From Figure 4.1, it can be observed that the assembly of the full substructures was highly effective. The relative error in natural frequencies remains on the order of 10^{-13}, even for the highest frequency considered, corresponding to an error of approximately $10^{-11}\%$. It is important to note that the first frequency shown corresponds to the second mode shape. Although the structural component is fully fixed and does not exhibit any rigid-body modes, the vibro-acoustic component includes a constant pressure mode. This mode is the acoustic equivalent of a rigid-body mode and is not considered in the comparison, as frequencies near zero can lead to larger discrepancies and less meaningful error values.

The MAC values comparing the assembled model with the reference model are presented in Figure 4.2 below.

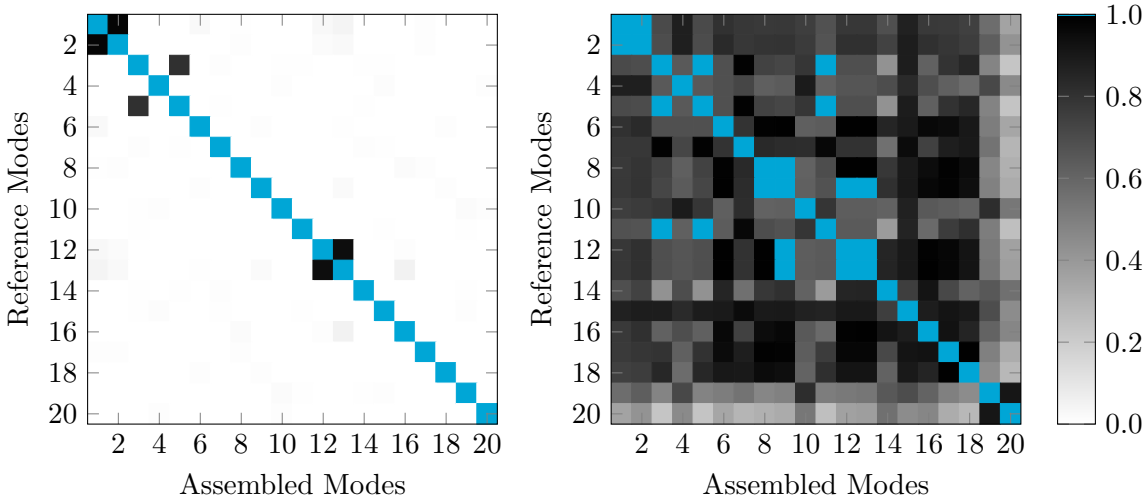

(a) MAC of the acoustic DoFs

(b) MAC of the structural DoFs

Figure 4.2: MAC between the full assembly and reference model

Note that the MAC was separated according to the degrees of freedom (DoFs) associated with the structural and acoustic components. This distinction was necessary because the scaling of the acoustic and structural contributions to the mode shapes was inconsistent between the assembled system and the reference model, as shown in the two modes shapes in Figure A.1 from Section A.1. Specifically, the mode shapes in the assembled system were normalised with respect to the structural DoFs, whereas in the reference model, they were normalised with respect to the acoustic DoFs. Despite this discrepancy, Figure 4.2 demonstrates that the assembled system accurately represents the reference model, as evidenced by the strong correlation of corresponding modes along the main diagonal.

The implemented MAC function was designed to round values of 0.99 or higher up to 1.0, indicating full correlation, as reflected in the colour-bar of the graph. This explains the presence of several off-diagonal entries marked as one in Section 4.1. Nevertheless, the figure reveals a generally high degree of similarity among the structural contributions of the different modes, as indicated by the consistently high MAC values across much of the matrix.

4.2. Reduction Results

This section presents and analyses the results of the vibro-acoustic reduction techniques. As previously mentioned, the purely structural component (Figure 3.1(b)) is reduced using the standard Craig-Bampton and Hintz-Herting methods, while the vibro-acoustic component (Figure 3.1(a)) is reduced using the various methods described in Section 3.3. The reduction basis includes the full set of constraint modes, which are necessary to preserve the complete interface, and the number of interface modes is varied to assess its influence on the results. Furthermore, if a CB based method is applied to the vibro-acoustic substructure, it is also applied to the structural substructure in the assembly. The same applies when using the HH method.

Two-Base Reduction

The first method analysed was the Two-Base Reduction. This approach was implemented using the existing CB framework, which was then further extended to support the HH reduction base. Figure 4.3 presents the results of both methods for three different numbers of internal modes: 8, 16, and 32. The results presented correspond to the relative frequency error. The MAC values are not shown, as it was verified that whenever the relative frequency error remained below 1%, the MAC between the reduced-order model and the full model indicated a complete correlation.

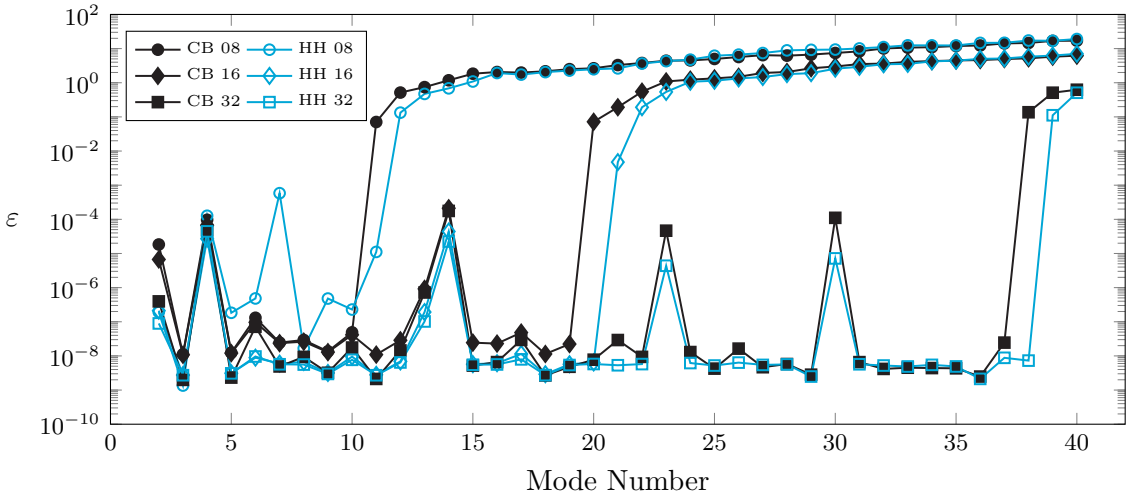

Figure 4.3: Two-Base Reduction: CB x HH

An important clarification is that since separate reduction bases are used for the structural and acoustic components, each configuration includes internal modes for both domains. For instance, selecting 16 internal modes means that both the structural transformation matrix $\mathbf{T_s}$ and the acoustic transformation matrix $\mathbf{T_a}$ include 16 internal modes each.

Figure 4.3 shows the performance of both reduction methods in producing reduced order models of the vibro-acoustic substructures. When these reduced models are assembled with the reduced structural substructure, the resulting natural frequencies closely match those of the full reference model. The level of accuracy is influenced by the number of internal modes included in the reduction bases. As expected, increasing the number of internal modes improves the accuracy over a wider frequency range. An important observation is that adding more internal modes does not significantly improve accuracy at lower frequencies. This behavior differs from what is typically observed in traditional CMS methods. Furthermore, unlike traditional approaches where the error tends to increase gradually with frequency, the error in this case shows a sudden jump once the number of internal modes becomes insufficient to represent the system dynamics.

The extension of the original CB framework to switch to the HH framework produced results that were slightly better in this model. Specifically, the HH approach performed better by one additional mode for the same number of internal modes and generally resulted in slightly lower relative frequency errors. However, it is important to note that, due to the absence of boundary pressure terms, the acoustic reduction basis does not include inertial relief modes. As a result, the constant pressure mode is not represented in the reduced model.

The Two-Base reduction approach omits certain coupling terms in the representation of the constraint modes due to the use of two separate bases. Despite this simplification, the method still employs coupled fixed interface modes to characterize the dynamic behavior of the reduced system. As a result, the reduction does not provide a statically complete representation, as noted in [7]. However, as demonstrated in the results, this limitation did not appear to negatively impact the accuracy or performance of the reduced model.

In the initial analysis, the number of internal modes was kept equal in both the structural and acoustic reduction bases. Figure 4.4 explores the impact of different distributions of internal modes while keeping the total number of modes constant across the two bases. Two additional test cases are introduced: in the first, the internal modes are distributed proportionally to the number of structural and acoustic degrees of freedom; in the second, the distribution is chosen so that the highest retained mode in each base corresponds to the same frequency.

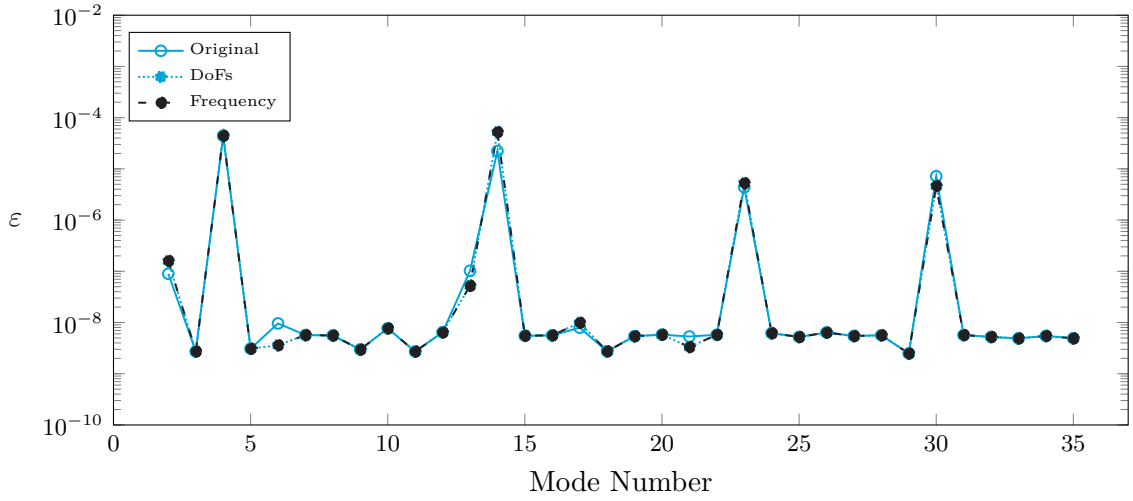

Figure 4.4: Influence of the distribution of structural and acoustic internal modes

As shown in Figure 4.4, optimizing the distribution of structural and acoustic internal modes does not lead to improved results in the reduction procedure. This may suggest that fewer modes could be sufficient for constructing the reduction basis. However, there appears to be no definitive best practice for determining the optimal number of modes to include from each domain. While it may be possible to define a cost function to minimize the error and identify an optimal combination of modes, this approach falls outside the scope of the current project. Therefore, the number of acoustic and structural modes will be kept equal in the subsequent analysis.

Nonsymmetric Reduction

The next reduction method analysed is the Nonsymmetric Reduction, which involves applying the traditional CMS procedure to a multi-physics system. As previously discussed, the primary drawback of this approach is the increased computational cost associated with solving eigenvalue problems involving nonsymmetric matrices. The results of applying the CB and HH reduction to the nonsymmetric problem are shown in Figure 4.5.

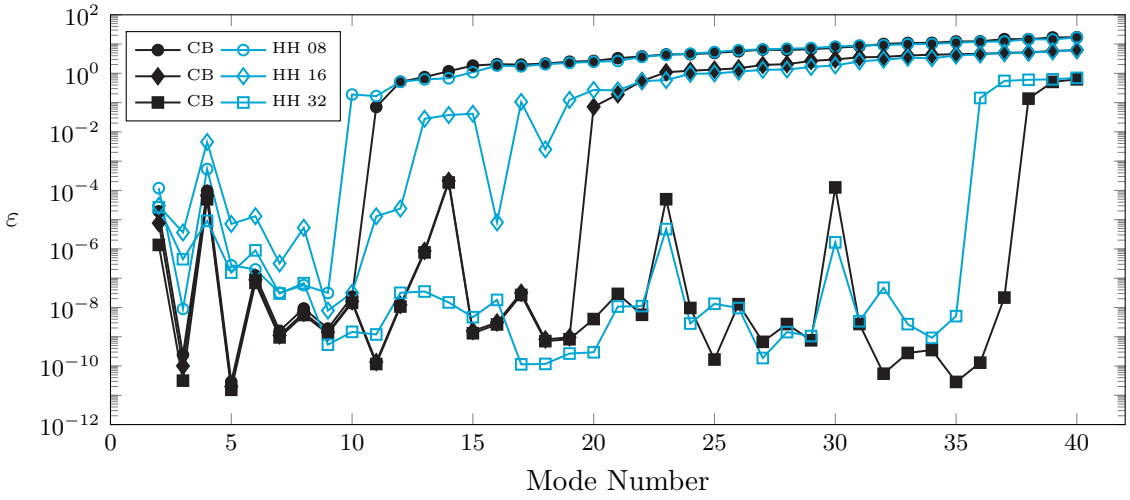

Figure 4.5: Nonsymmetric Reduction: CB x HH

As shown in Figure 4.5, the Nonsymmetric Reduction successfully produces a reduced-order model of the vibro-acoustic system, evidenced by the low relative error in the assembled eigenfrequencies. However, unlike the Two-Base Reduction, the CB method outperforms the HH approach in this case. An additional advantage of this method is its relatively straightforward implementation, requiring only the partitioning of degrees of freedom (DoFs) into boundary and internal groups. A similar sudden increase in error is observed with this technique, where the relative error rises sharply at a specific frequency.

The key consideration is whether the increased computational cost of this method is justified by the improved accuracy. Since all coupling terms are now captured within the constraint modes, the reduction yields a statically complete solution. A comprehensive comparison of this technique with others is provided in Section 4.3.

Locally Symmetric CB Method

The final CMS based vibro-acoustic reduction technique analysed is the Locally Symmetric CB Method, proposed by Tabak [7]. However, when attempting to implement the method as described, the resulting reduced-order model of the vibro-acoustic substructure (Figure 3.1(a)) failed to retain all the modes present in the full component. Specifically, all acoustically dominant modes, including the constant pressure mode, were absent from the reduced model.

As a result, the ROM generated using this technique preserved only the structurally dominant modes of the system, including the rigid-body modes. The first four flexible modes of the system are shown in Section A.2. An analysis of these modes reveals that they clearly correspond to the natural bending modes of the structural beam. Within the frequency range illustrated in Figure A.2, sixteen acoustically dominant modes are missing, excluding the constant pressure mode.

A possible explanation for this outcome is that the reduction technique may only be valid for components that feature both structural and acoustic elements at the interface (retained DoFs). This assumption aligns with the models studied by Tabak [7], whereas the academic model used in this project includes only structural elements at the interface. Further investigation is necessary to validate this hypothesis. For the Case Study this technique will not be applied.

4.3. Comparison of the Reduction Techniques

To evaluate which reduction technique produced the best results, the outcomes from the Two-Base and Nonsymmetric Reduction methods are compared in Figure 4.6. For clarity, only the results obtained using the CB method are shown.

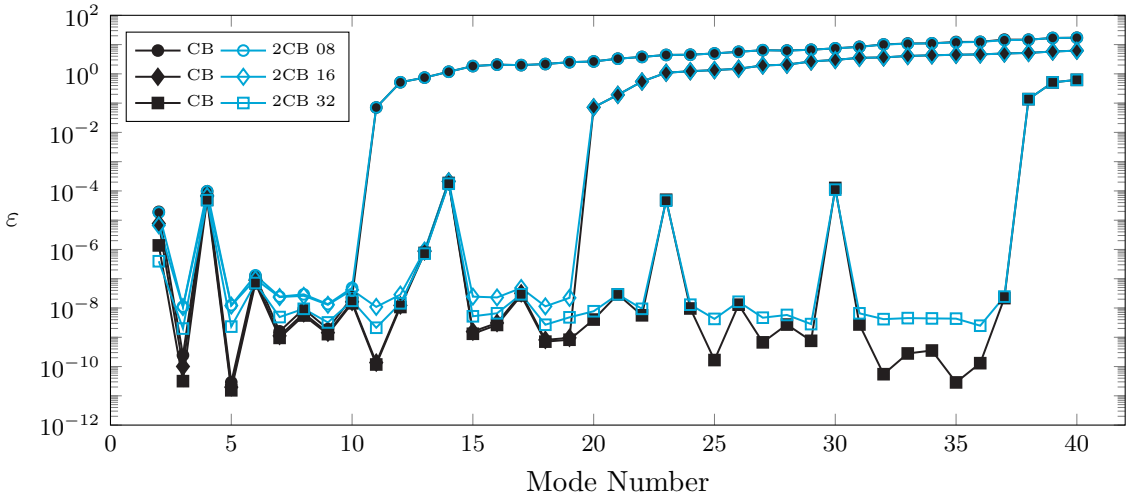

Figure 4.6: CB comparison of Two-Base (blue) and Nonsymmetric (black) Reduction techniques

From Figure 4.6, it can be observed that both methods are valid within the same frequency range, as the sudden jumps in error occur at identical frequencies. Moreover, in the frequency regions where the reduced models are accurate, the Nonsymmetric Reduction either outperforms or at least matches the performance of the Two-Base Reduction. However, this improved performance comes at the cost of increased computational effort required to generate the models, as shown in Table 4.1.

Table 4.1: Computation Time for Generating Reduced Models

Reduction Technique	Method	Time [s]
Two-Base	CB	0.039327
	HH	0.014474
Nonsymmetric	CB	0.148607
	HH	0.123205

These computations were performed on a local machine equipped with an Intel i9−13950HX processor featuring 24 cores, 32 threads, a base clock speed of 2.2 GHz and turbo boost up to 5.5 GHz, along with an NVIDIA RTX 3500 Ada graphics card and 32 GB of RAM.

The need to solve a larger generalised eigenvalue problem, combined with the fact that the matrices are nonsymmetric, results in a substantial increase in computational time when generating and applying the reduction bases to the vibro-acoustic models. Table 4.1 shows that by choosing the Nonsymmetric Reduction over the Two-Base Reduction leads to a 277.88% increase in computation time for the CB method and a 751.22% increase for the HH method. These times were computed with reduction bases with the constraint modes and 64 internal modes for each base. While this increase may be negligible for the academic model, it could pose a significant limitation for industry scale models.

Additionally, the comparison presented in Table 4.1 and Figure 4.6 is not entirely fair. This is because the Two-Base Reduction used twice as many internal modes as the Nonsymmetric Reduction, since separate bases were computed for the structural and acoustic components. To address this imbalance, Figure 4.7 presents a revised comparison in which both techniques use the same total number of internal modes. This was achieved by halving the number of internal modes in each base of the Two-Base Reduction.

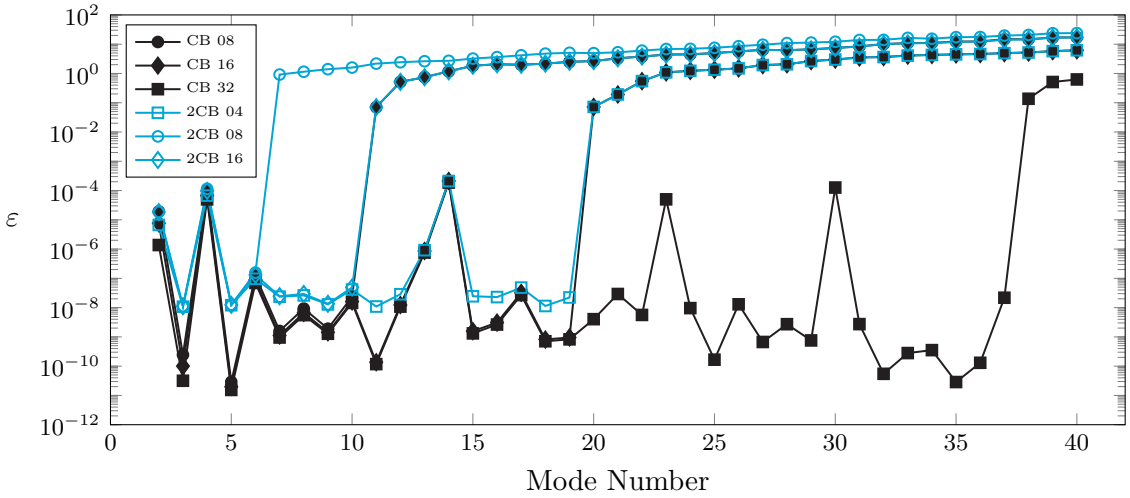

Figure 4.7: CB comparison - same Mmtrix sizes

Figure 4.7 shows that the relative errors for the Nonsymmetric Reduction are significantly lower than those of the Two-Base Reduction when both methods use matrices of the same size. This means that the same total number of internal modes was used in the reduction process. However, while the difference in computational effort was already considerable in Table 4.1, it becomes even more pronounced when the Two-Base Reduction requires fewer modes to achieve the same matrix size.

4.4. Results with Damping

In Section 4.2, the assembly of the reduced substructures was evaluated against the full reference model in the undamped case. This section extends that analysis by introducing an impedance boundary condition on the vibro-acoustic substructure, as described in Section 3.4.

With the introduction of damping, the generalised eigenvalue problem becomes a non-linear eigenvalue problem, as shown in Equation 3.25, which must be solved to obtain the system's frequencies. As a result, the frequencies of the damped model appears as complex conjugate pairs.

Since MATLAB does not consistently order these complex pairs in the frequency vector (i.e. the order of the positive and negative components may vary), the relative frequency error, originally denoted as $\boldsymbol{\varepsilon}$, is adjusted to account for these discrepancies. The modified error metric, denoted as $\boldsymbol{\epsilon}$, is defined in Equation 4.3.

$$\boldsymbol{\epsilon} = \frac{\big||\boldsymbol{f}| - |\boldsymbol{f_{ref}}|\big|}{|\boldsymbol{f_{ref}}|} \tag{4.3}$$

To obtain the reduced models, the same procedure used for the undamped case was followed. The only modification involved the computation of the internal modes that form the reduction basis. The constraint modes remained unchanged, as they are based on a static approximation and are therefore unaffected by damping. Figure 4.8 presents the results for the Two-Base Reduction, while Figure 4.9 compares the CB results of the Two-Base and Nonsymmetric Reductions. The Locally Symmetric Reduction produced the same outcome as in the undamped case, retaining only the structurally dominant modes in the reduced system.

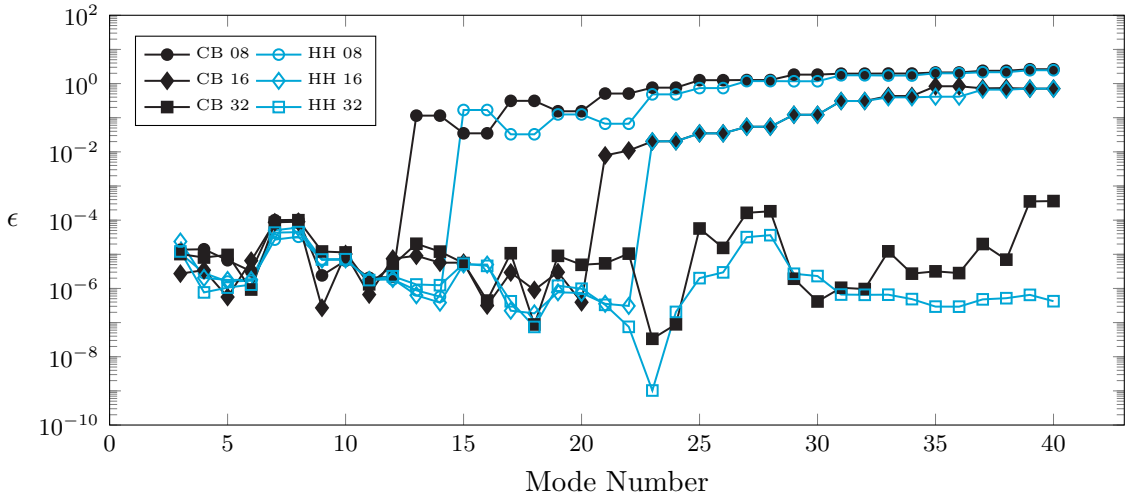

Figure 4.8: Damped Two-Base Reduction

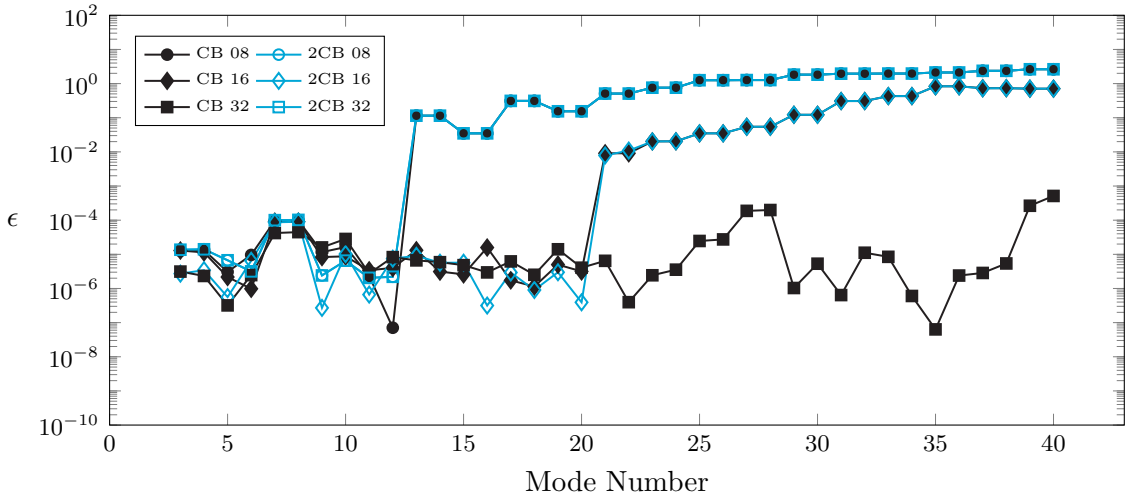

Figure 4.9: Damped CB reduction: Two-Base x Nonsymmetric

Figure 4.8 and Figure 4.9 clearly show that the reduction techniques analysed remain effective when damping is present in the system. As in the undamped case, the HH method outperforms CB within the Two Base Reduction, while the Nonsymmetric Reduction delivers better performance than the Two Base approach when the same total number of internal modes is used. Note that ϵ begins from the third mode number, as the system now includes a pair of constant pressure modes.

It is important to highlight that, with damping, using the same number of modes as in the undamped case reduces the frequency range over which the reduced model remains valid. This is because the frequencies now appear in complex conjugate pairs, meaning that for the same number of modes, fewer distinct frequencies are effectively captured. Since the accuracy of the reduced-order model depends on the presence of both complex conjugate modes, omitting one member of each complex pair can degrade the accuracy of the reduced model.

4.5. Main Takeaways

Of the three CMS based reduction techniques analysed, two demonstrated potential for generating reduced-order models of vibro-acoustic systems. Based on the various analyses conducted, several key observations have been made:

- Unlike traditional CMS methods, the methods analysed, when applied to vibro-acoustic systems, the relative error shows a sharp threshold beyond which the reduced model no longer accurately captures the system's dynamics.
- Also differing from traditional CMS behaviour, increasing the number of internal modes does not significantly improve accuracy at lower frequencies, it instead extends the frequency range over which the reduced-order model remains reliable.
- Similar to rigid-body modes, constant pressure modes theoretically have a natural frequency of zero. However, in vibro-acoustic models, the constant pressure mode consistently appears at a non-zero frequency of several hertz.
- The Two Base Reduction requires a greater number of internal modes to achieve accuracy comparable to the Nonsymmetric Reduction, which can reach similar performance with significantly lower computational effort.
- The Locally Symmetric Reduction was only able to produce a reduced-order model that captured the structurally dominant modes of the full system. This may be due to the fact that the analysed model had only structural elements at the boundary. Further investigation is needed to confirm this.
- The techniques are compatible with damping, provided that both modes of each complex conjugate pair are included in the reduction bases.

Part III

Case Study

5

Industry Size Model

The primary aim of the Case Study is to apply the methodology developed and validated during the Initial Investigation, originally tested on an academic model, to a model relevant to ASML. This chapter begins by describing the models used in the Case Study. It then outlines the necessary modifications to the original methodology to handle models with millions of DoFs and to align with ASML's way of working. Next, the updated reduction and coupling procedures are presented. Finally, the performance of the reduced models is compared to that of the full models, with the aim of quantifying the computational savings achieved.

5.1. Testbench Model

The model selected for analysis in the Case Study is a variation of the Testbench Model shown in Figure 5.1(b). The version depicted in Figure 5.1(a) is a well-established model within ASML, commonly used to evaluate new modelling techniques. It was chosen for this study due to its structural similarity to models of interest and, more importantly, to ensure the thesis remains fully publishable without confidentiality restrictions. This choice allows the full methodology and results to be published without restrictions, ensuring the work remains accessible and complete to all.

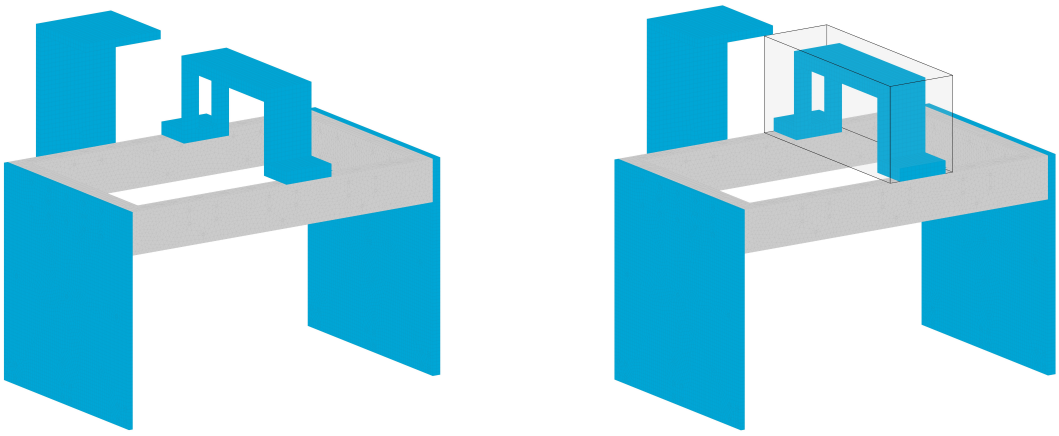

(a) Original Testbench model

(b) Modified Testbench model

Figure 5.1: Addition of an acoustic cavity to the Testbench model

The Testbench model consists of five components: the frame, thin plate, thick plate, L-bracket, and bridge. With the component names and Figure 5.1, it is easy to identify each part. It is worth noting that the thin plate is positioned on the L-bracket side, while the thick plate is

located on the bridge side. The only modification made to the original Testbench model is to the bridge component, where an acoustic cavity has been added around it. By introducing the acoustic cavity, the previously structural-only component is now modelled as a vibro-acoustic component, as further detailed in Section 5.2.

The primary objective of this Case Study is to integrate the newly developed vibro-acoustic model of the bridge together with the existing, purely structural, models of the frame, thin plate, thick plate, and L-bracket. This would showcase that ASML's modular modelling methodology has been extended to account for the vibro-acoustic phenomena.

5.1.1. Modified Bridge

This vibro-acoustic component is the bridge surrounded by air, as seen in the in the diagram of the FSI in Figure 5.2. Note that the acoustic-structure coupling will only be applied between the air and the bridge, the air will not be coupled to the other purely structural components.

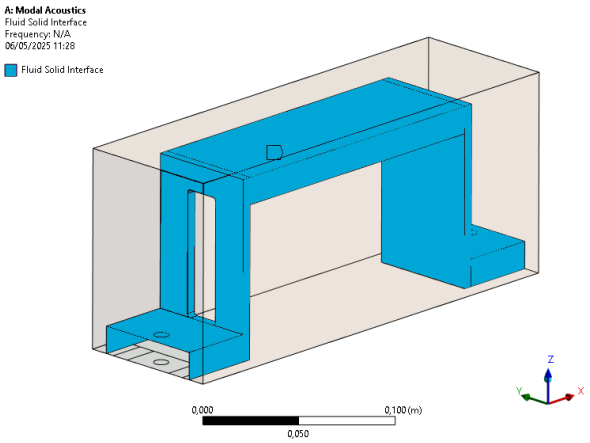

Figure 5.2: FSI of the vibro-acoustic bridge substructure

This vibro-acoustic model serves as the basis for applying the reduction techniques under analysis. The objective is to enable its integration into a broader system-level simulation, thereby extending ASML's modular modelling methodology to incorporate vibro-acoustic phenomena. It is worth noting that the acoustic cavity used in the final model was extended from the original 0.25 meters (as shown in Figure 5.2) to a length of three meters. This modification was made to lower the natural frequencies of the acoustically dominant modes, allowing them to appear within the same frequency range as the first few structurally dominant modes in the frequency response functions, as shown in Section 6.3.

5.1.2. Purely Structural Components

The frame, L-bracket, thin plate, and thick plate components remain unchanged from the original Testbench component models. All are modelled using the same steel material as the structural part of the bridge. Since the focus is solely on vibro-acoustic reduction, these components will be reduced using the standard Hintz-Herting CMS method, as it demonstrated better performance than the CB method. For all analyses, the reduction will include the constraint modes along with a fixed set of 64 internal modes (inertial relief and free-body modes).

5.2. Modifications to the Methodology

This section outlines the modifications necessary to handle larger models and to adapt the presented reduction and assembly methodology to align with ASML's way of working. Managing larger models requires the use of specialised numerical methods capable of efficiently handling and operating on large matrices. To ensure compatibility with ASML's modular approach, the assembly process must reflect how substructures (typically machine components) are physically connected, often using bolts and rivets. In ASML's modelling practice, these connections are represented using Remote Points, which are linked via spring elements to simulate the mechanical behavior of the joints. The following subsections provide a detailed discussion of these procedures.

5.2.1. Remote Points

As briefly introduced in Section 1.3.3, ASML's modular approach relies on the use of Remote Points (RPs). A Remote Point is an additional node, referred to as the master node, that is integrated into the system and connected to a selected group of component nodes. These connections can be either rigid or flexible and are implemented via Multipoint Constraint (MPC) equations [42]. A Remote Point replaces all the nodes it was originally connected to in the system matrix equations.

MPCs provide a framework for relating degrees of freedom within a model. In the case of rigid connections, the system is constrained to enforce a one-to-one linear relationship in both displacement and rotation between the master node and the connected nodes, artificially increasing the system's stiffness. In contrast, flexible connections allow the motion of the master node to be a weighted average of the motions of the connected nodes, enabling more compliant behavior.

The choice between rigid and flexible connections depends on the specific application. At ASML, a general rule of thumb is to use a rigid connection when the master node is linked to a small number of nodes, and a flexible connection when it is connected to a larger number of nodes. Remote points are typically added at component interfaces, as well as at any other locations that need to be preserved in the reduced-order models of the component—such as sensor and actuator positions.

Figure 5.3 illustrates the placement of remote points at the interface of the bridge. This is a conceptual representation, in order to make the visualization of the constraints between the master node and the selected surface nodes (shown in red) more intuitive. By convention, the remote point is positioned at the geometric center of the associated nodes.

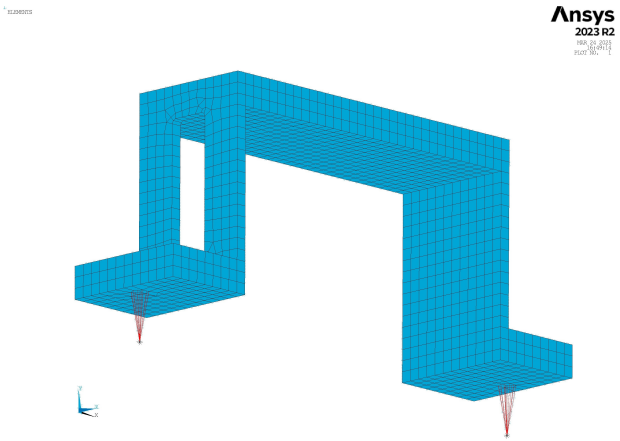

Figure 5.3: Visualisation of the Remote Points

Since the use of Remote Points is essential to align with ASML's modelling approach,

the modelling software must be changed from COMSOL to ANSYS. This is because COMSOL lacks a direct equivalent to Remote Points. While COMSOL does allow the creation of virtual nodes, these are not treated as physical nodes within the system and therefore cannot replicate the functionality required. The main benefits of using the remote points are outlined below.

Interface Improvement

One of the main limitations of CMS as a reduction technique is its inefficiency for systems with large interfaces. This is because CMS requires one constraint mode for each DoF at the interface nodes, which can lead to a significant number of modes and reduced computational efficiency. By introducing Remote Points, the interface DoFs can be condensed into a single master node with only six DoFs, three translational and three rotational. This approach allows the entire interface to be represented using just six constraint modes, significantly reducing the complexity of the reduced model.

Reducing the number of constraint modes that need to be computed becomes increasingly important for larger models. The modes from Equation 3.13 can be computed using matrix left division, which is significantly more efficient when the number of columns in $\mathbf{K}_{ib}$ is reduced. In MATLAB, matrix left division is preferred over explicitly computing the matrix inverse and then multiplying, as it is both more robust and numerically stable. MATLAB internally uses Gaussian elimination (or similar factorization methods) to solve the system without explicitly forming the inverse, which improves both performance and accuracy.

Remote Points of the Testbench Components

As seen in Figure 1.3(a), the Testbench model has four main interfaces (Int.1 - Int. 4), which corresponds to the coupling of four components to the frame. Additionally, both plates are fixed to the ground. To perform the coupling, Remote Points were added to the interfaces of these components. The number of RPs used for each interface is shown in Table 5.1 below. Additionally, remote points are also used for the ground connections.

Table 5.1: Number of Remote Points per Interface

Component	Int. 1	Int. 2	Int. 3	Int. 4	Ground
Bridge	2	0	0	0	0
Frame	2	5	5	5	0
Thin	0	5	0	0	2
Thick	0	0	5	0	1
L-bracket	0	0	0	5	0

These are the only nodes retained in the reduced models, as the Remote Points are designated as the boundary nodes of the system. However, since pressure needs to be evaluated at a specific location within the acoustic cavity (as discussed in Chapter 6), an additional pressure DoF node (not a Remote Point) must also be included in the boundary node set of the bridge.

Assembly of Substructures

As previously mentioned, the various substructures are typically connected using bolts or rivets. These connectors can be modelled as joints with defined stiffness properties in all six degrees of freedom. Each type of joint imposes different constraints on these DoFs. This modelling approach allows the individual components to be meshed independently, as it eliminates the need for matching mesh nodes at the interfaces, unlike in a primal assembly. Figure 5.4 shows a

conceptual representation of how the assembly is constructed between the bridge and the frame.

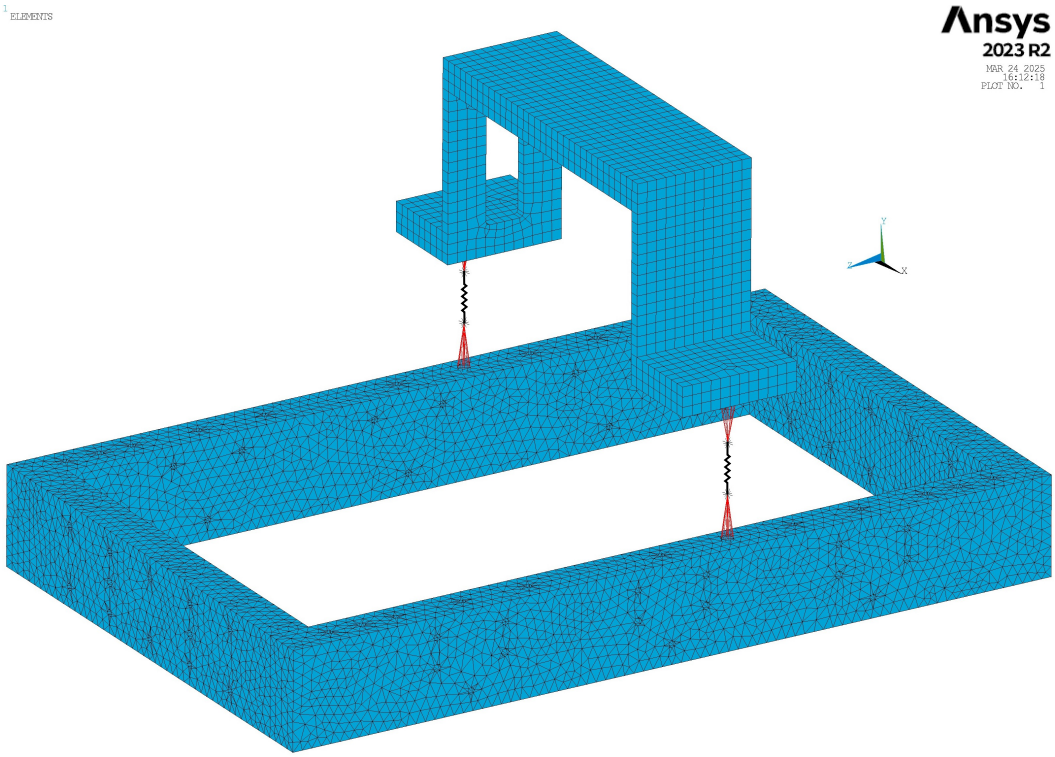

Figure 5.4: Representation of spring connection between Remote Points

As indicated, Figure 5.4 serves purely as a visual aid of how these connections are established. For the joint connection to be performed, the linked remote points must share the same coordinates in the global coordinate system, within a specified tolerance. Equation 5.1 demonstrates how the coupling between structures is achieved by introducing a connecting stiffness between two remote points, p_1 and p_2.

$$K(p_1, p_2) = \text{diag}\left(\begin{bmatrix} k_i & -k_i \\ -k_i & k_i \end{bmatrix}_i\right) \quad \forall \quad i = \begin{bmatrix} x, & y, & z, & rx, & ry, & rz \end{bmatrix} \tag{5.1}$$

The values chosen for all the translational stiffness (k_x, k_y and k_z) and the rotational stiffness (k_{rx}, k_{ry} and k_{rz}) are 10^{11} $\mathrm{N\,m^{-1}}$ and 10^{9} $\mathrm{N\,rad^{-1}}$, respectively. These values are approximate representations of the stiffness typically associated with bolted connections.

To assemble the different substructures, a spreadsheet is required. This spreadsheet includes, in its first two columns, the names of the Remote Points that are to be coupled. The following twelve columns specify the connection properties: six stiffness coefficients and six damping coefficients. If a Remote Point is to be connected to the ground, the second Remote Point name is left blank. ASML uses a proprietary function that takes a cell array containing all the substructures along with this spreadsheet as input. The function then assembles the components accordingly to generate the complete system model.

5.2.2. Numerical Considerations

Working with models containing millions of DoFs presents significant numerical challenges due to the size of the system matrices. These matrices are typically sparse, meaning most of their entries are zero, because of the way finite element models are structured. Storing matrices in sparse format omits zero entries, allowing them to be handled on local machines with limited memory. To efficiently work with sparse matrices, optimized algorithms are required. One such example is MATLAB's `eigs` function, which is the sparse-matrix counterpart to the standard `eig` function. It solves the generalized eigenvalue problem (Equation 3.7) but computes only a selected subset of eigenvalues, making it suitable for large-scale problems.

When computing the modes for the reduction bases, simply requesting the smallest-magnitude eigenvalues was not sufficient to ensure convergence. Instead, a scalar `sigma` must be specified in the function, guiding the algorithm to search for eigenvalues near a given value. For vibro-acoustic systems, which always includes a constant pressure mode near zero hertz, a small value for `sigma` can be used to target low-frequency modes. The algorithm showed low sensitivity to the exact value of `sigma`.

The way matrix operations are performed becomes critical to computational efficiency when working with large matrices. It is essential to minimise the number of computationally intensive operations and to structure matrix groupings so that operations are carried out on the smallest possible matrices. For example, when computing the inertia relief modes from Equation 3.17, the MATLAB operation shown in Equation 5.2a is approximately 48% more efficient than the alternative in Equation 5.2b, despite producing the same result.

$$\mathbf{H}_{i0} = -\mathbf{K}_{ii} \backslash \left[\left(\mathbf{M}_{ib} + \mathbf{M}_{ii} \boldsymbol{\psi} \right) \boldsymbol{\phi}_{\boldsymbol{b0}} \right] \tag{5.2a}$$

$$\mathbf{H}_{i0} = -\left[\mathbf{K}_{ii} \backslash \left(\mathbf{M}_{ib} \boldsymbol{\phi}_{\boldsymbol{b0}} \right) \right] - \left[\mathbf{K}_{ii} \backslash \left(\mathbf{M}_{ii} \boldsymbol{\psi} \boldsymbol{\phi}_{\boldsymbol{b0}} \right) \right] \tag{5.2b}$$

Although primal assembly is no longer required for the case study under analysis, it is worth noting that computing the localisation matrix $\mathbf{L}$ using the null space of the compatibility matrix $\mathbf{B}$, as discussed in Section 3.2.3, is not feasible. This is because MATLAB's built-in `null` function does not support sparse matrices. As a result, the localisation matrix must be computed manually. Fortunately, this process is relatively straightforward to automate.

Additionally, the `polyeig` function is also not compatible with sparse matrices. Working with the full matrices is not feasible, as storing a single matrix of that size would require in the order of 1 TB of RAM. Therefore, to compute the modes necessary for reducing large-scale systems when there is acoustic damping, an alternative approach is required. The quadratic eigenvalue problem described in Equation 3.25 can be reformulated as a linear eigenvalue problem using the state-space method proposed by Gruber et al. [43]. This approach introduces a redundant equation, Equation 5.3, which is added to the original equation of motion in Equation 1.2.

$$\mathbf{M}\dot{x} - \mathbf{M}\dot{x} = 0 \tag{5.3}$$

By combining Equation 1.2 with Equation 5.3, the system can be reformulated into the following state-space representation. The state vector $\boldsymbol{x}$ represents the original DoFs vector, comprising both pressure and displacement components.

$$\mathcal{A} \begin{bmatrix} \dot{\boldsymbol{x}} \\ \ddot{\boldsymbol{x}} \end{bmatrix} + \mathcal{B} \begin{bmatrix} \boldsymbol{x} \\ \dot{\boldsymbol{x}} \end{bmatrix} = \mathcal{F} \tag{5.4}$$

The expanded $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{F}$ matrices are shown in Equations 5.5-5.7. Note that the equation corresponding to the first line of the matrices corresponds to Equation 1.2, while the second line corresponds to Equation 5.3.

$$\mathcal{A} = \begin{bmatrix} \mathbf{C} & \mathbf{M} \\ \mathbf{M} & \mathbf{0} \end{bmatrix} \tag{5.5}$$

$$\mathcal{B} = \begin{bmatrix} \mathbf{K} & \mathbf{0} \\ \mathbf{0} & -\mathbf{M} \end{bmatrix} \tag{5.6}$$

$$\mathcal{F} = \begin{bmatrix} \boldsymbol{f} \\ \mathbf{0} \end{bmatrix} \tag{5.7}$$

With the state-space formulation presented in Equation 5.4, the generalised eigenvalue problem:

$$(\mathcal{B} - \boldsymbol{\lambda}\mathcal{A})\,\hat{\boldsymbol{x}} = \mathbf{0} \tag{5.8}$$

can be solved with MATLAB with the normal `eigs` function. An alternative to the state-space formulation is the use of the non-native `polyeigs` function, a modified version of MATLAB's built-in `polyeig` function that supports sparse matrices. This version was found in the MATLAB forum and preliminary tests showed similar efficiency as the state-space approach. However, since it is not officially supported by MATLAB, its robustness across different applications cannot be guaranteed.

5.3. Updated Reduction Procedure

5.3.1. FEM Package

As previously explained, the use of Remote Points requires switching from COMSOL to ANSYS as the FEM modelling tool. This transition calls for several adjustments to the existing methodology. Since ANSYS does not offer a LiveLink with MATLAB, system matrices must be extracted using the code below directly in the ANSYS APDL command prompt. Alternatively, if ANSYS Workbench is being used, the same commands must be inserted into the solution branch as a command snippet.

```
hbmat,%outFile%,txt,,ascii,MASS,no,yes
hbmat,%outFile%,txt,,ascii,DAMP,no,yes
hbmat,%outFile%,txt,,ascii,STIFF,no,yes
```

Once exported, these files can be imported into MATLAB using a conversion function that transforms HB-format files into MATLAB sparse matrices. Additionally, a mapping file is generated, which links the matrix equations to the corresponding node numbers and DoFs types (pressure, translational displacements (pressure, x, y, z-displacements and x, y, z-rotations). This mapping is essential for converting the system matrices into the ***u-p*** formulation required for preconditioning. With the system matrices and mapping files now available, all operations previously carried out using LiveLink on the COMSOL model can be reproduced for ANSYS models.

5.3.2. Preconditioning

In ANSYS, the vibro-acoustic formulation follows a convention that differs from the one commonly found in literature. Specifically, the coupling terms (off-diagonal blocks) in Equation 3.1 are reversed. As a result, the mass matrix in ANSYS takes the form of an upper triangular block matrix, while the stiffness matrix is structured as a lower triangular block matrix. This difference does not affect the reduction procedure itself but does influence the computation of the preconditioning coefficients.

By following the same logic outlined in Subsection 3.2.1, the parameters from Equation (3.5) are updated accordingly to reflect this formulation:

$$a_{2_1} = \frac{\|\mathbf{K_c}\|}{\|\mathbf{K_a}\|} \tag{5.9a}$$

$$a_{2_2} = \frac{\|\mathbf{M_s}\|}{\|\rho_0\mathbf{K_c}\|} \tag{5.9b}$$

$$b_{2_1} = \frac{\|\mathbf{K_s}\|}{\|\mathbf{K_c}\|} \tag{5.9c}$$

$$b_{2_2} = \frac{\|\rho_0\mathbf{K_c}\|}{\|\mathbf{M_a}\|} \tag{5.9d}$$

5.3.3. Compatibility with ASML's Way of Working

The core idea behind the modifications disclosed in this chapter is to use Remote Points to reduce the number of interface degrees of freedom, and then connect the resulting substructures via master nodes using springs defined in a spreadsheet. However, simply including Remote Points in the model is not sufficient to ensure compatibility with ASML's existing methodology. Additional requirements must be met. Specifically, the reduced model must be stored in a MATLAB structure with standardised field names. This structure should include, for example, the reduced matrices, the names of the master nodes, and their coordinates.

Although the exact format of this structure cannot be disclosed due to confidentiality, it is designed to be compatible with ASML's proprietary functions. These functions rely on the structure to interpret the model and couple it with other components. One of these functions, for example, identifies matching node names between models based on the spreadsheet, verifies that the corresponding nodes share the same coordinates, and then applies the coupling stiffness terms (Equation 5.1) to establish the connection. In this way, ASML's modular structural modelling way of work is extended to support vibro-acoustic substructures.

5.4. Applying the Reduced Model

The final task of this case study is to evaluate the computational efficiency gained by using the reduced models. The reduced models, represented schematically in Figure 5.5(b), will be used in a straightforward analysis to compare their performance against the full model (Figure 5.5(a)). The objective is to assess whether the computational speed-up justifies the overhead of generating the reduced models.

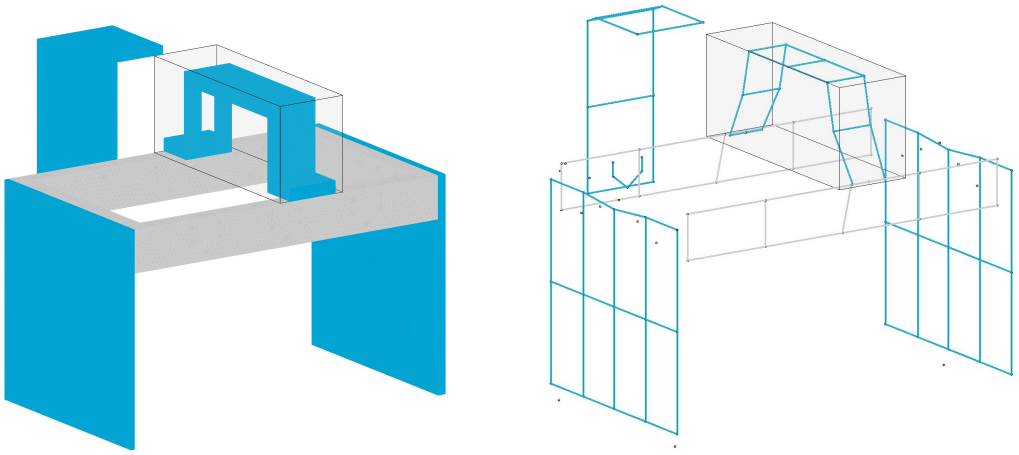

(a) Full Testbench model

(b) Reduced Testbench model

Figure 5.5: Addition of an acoustic cavity to the Testbench model

5.4.1. Harmonic Excitation

For this analysis, harmonic excitation will be introduced into the model in two distinct ways: as a unit point source within the acoustic cavity and as a unit point force applied to the structure. Each excitation will be applied in separate simulations, and the resulting displacement and pressure will be measured at specific locations.

In the first simulation, the acoustic point source is activated, and the resulting structural displacement is measured as a function of frequency, as illustrated in Figure 5.6(a). In the subsequent simulation, the excitation is shifted to one of the previously measured structural points, where a harmonic unit force is applied, as shown in Figure 5.6(b). In this second case, the pressure response is measured at the location where the acoustic point source was originally applied. Figure 5.6 illustrates the locations of the sensors and excitation points on the full model. These same positions will also be used in the reduced models.

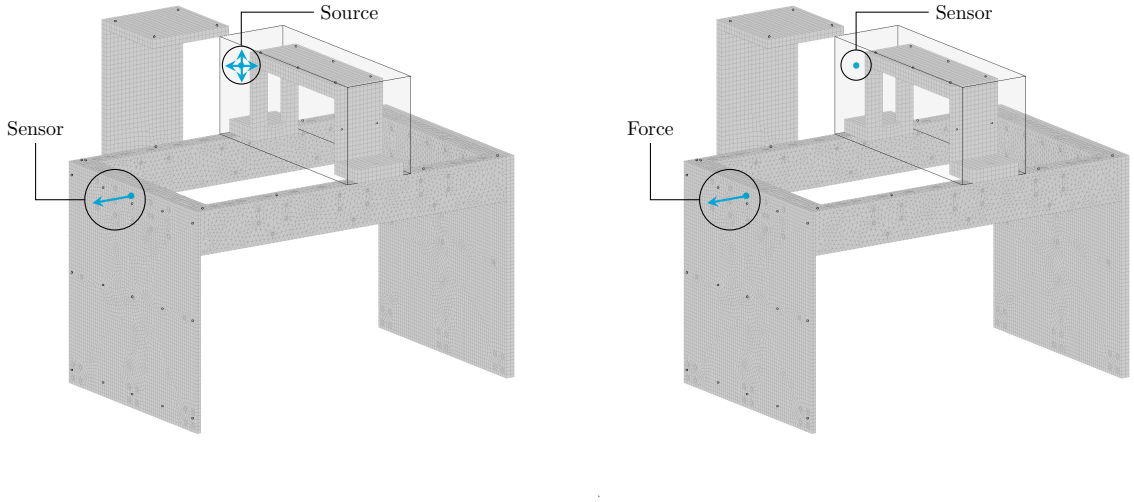

(a) Position of the point source

(b) Position of the point force

Figure 5.6: Position of the different excitations

The system's response to harmonic excitation is computed by multiplying the inverse of the mechanical impedance matrix, $\mathbf{Z}$, also known as the dynamic stiffness matrix, with the force vector $\boldsymbol{f}$. This matrix $\mathbf{Z}$ is derived by applying a Fourier transform to the displacement and force terms in Equation 1.1, thereby expressing the system in the frequency domain [17]. Since time derivatives in the frequency domain correspond to multiplication by $\mathrm{i}\omega$, the left-hand side terms can be consolidated into a single frequency-dependent matrix $\mathbf{Z}$, as shown in Equation 5.10.

$$\begin{aligned}(-\omega^2\,\mathbf{M} + \mathrm{i}\omega\,\mathbf{C} + \mathbf{K})\hat{\boldsymbol{x}}(\omega) &= \boldsymbol{f}(\omega),\\ \mathbf{Z}(\omega)\,\hat{\boldsymbol{x}}(\omega) &= \boldsymbol{f}(\omega).\end{aligned} \tag{5.10}$$

Note that when applying a force in the reduced models, the node at which the excitation is applied does not need to be included in the boundary set of nodes ($\boldsymbol{x_b}$) during the reduction process. This is because the reduction basis can be applied to the full excitation vector, allowing the force to be distributed across the internal modes, as shown on the right-hand side of Equation 1.5.

However, since displacement and pressure need to be measured at specific points, the corresponding nodes must be included in the boundary set. These nodes must be retained in the reduced models to ensure they remain accessible for sensing. Therefore, the boundary set for this problem includes all interface nodes (Remote Points) as well as the nodes corresponding to the DoFs at the sensor locations ($\boldsymbol{x_{sp}}$), as defined in Equation 5.11.

$$\boldsymbol{x} = \begin{bmatrix} \boldsymbol{x_b} \\ \boldsymbol{x_i} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \boldsymbol{x_{rp1}} \\ \boldsymbol{x_{rp2}} \\ \vdots \\ \boldsymbol{x_{sp}} \end{bmatrix}_b \\ \boldsymbol{x_i} \end{bmatrix} \tag{5.11}$$

The computational efficiency will be evaluated by comparing the total time required to solve the full model with the combined time needed to solve the reduced models and to generate them. This includes the time to compute the modes for the reduction basis and to generate the reduced system matrices, as described in Equation 1.5.

6

Results Analysis

In the previous chapter, the Case Study was introduced, along with the key considerations required to implement the methodologies previously discussed on an industry-scale model. These included the use of Remote Points to condense interfaces into single nodes, the coupling procedure in which joints are applied to the corresponding Remote Points, and various computational techniques used to efficiently manipulate large system matrices.

In this chapter, the Testbench model will be analysed. The bridge component will be reduced using the proposed reduction methods, while the remaining purely structural components will be modelled using ASML's standard CMS approach. The assembly of these reduced components will then be compared to the full system model. Finally, both the Reduced Order Model (ROM) and the Full Order Model (FOM) will be evaluated in a harmonic excitation scenario to assess the computational advantages of using a ROM over a FOM.

6.1. Reduction Results

This section presents and analyses the results of the vibro-acoustic reduction techniques applied to the Case Study. In these analyses, all purely structural components are reduced using the original Hintz Herting method, each retaining a total of 64 internal modes (inertia relief and free-body modes). The vibro-acoustic bridge, in contrast, is reduced using the various proposed methods, with different numbers of internal modes considered.

Figure 6.1 illustrates the relative frequency error of the assembled system when the bridge is reduced using the Two-Base method.

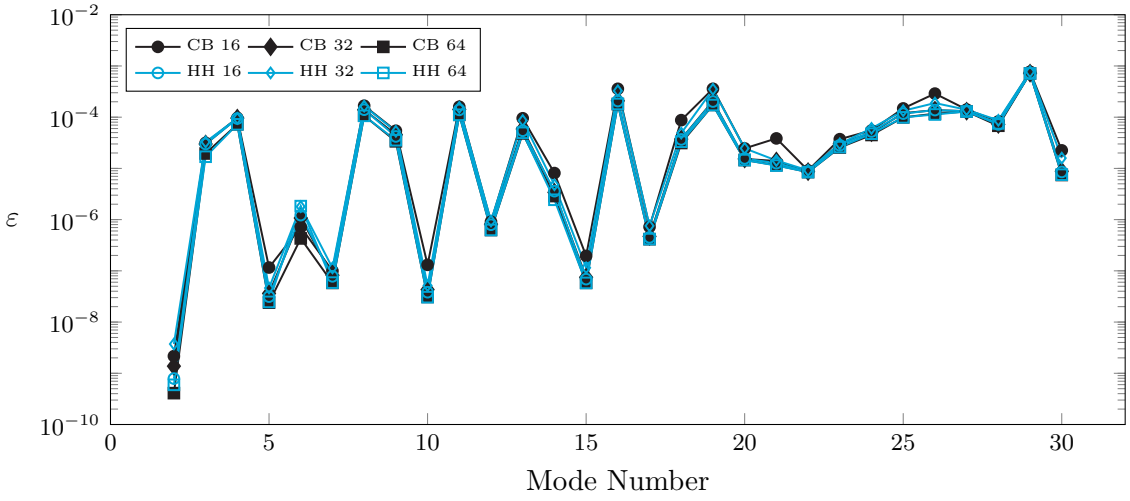

Figure 6.1: Testbench Two-Base Reduction: CB x HH

Figure 6.1 compares the performance of the CB and HH methods for the Two-Base reduction approach. When these reduced models are assembled with the reduced structural substructure, the resulting natural frequencies show strong agreement with those of the full reference model. Increasing the number of internal modes does not significantly reduce the frequency error. This may be due to the use of the same reduced versions of the purely structural components in all analyses, which could have a dominant influence on the overall assembly frequencies.

The gradual increase in error across the frequency range, without the sudden spikes observed in Chapter 4, supports the assumption that the dominant source of error appears to be the purely structural components. This pattern is typical of classical CMS methods. Furthermore, the difference in performance between the CB and HH methods is relatively small, although HH still appears to slightly outperform CB. A computational effort analysis of generating these reduction bases is presented in Section 6.4, and therefore the CB method is only preferable if it offers a significant computational advantage.

The Nonsymmetric Reduction method was then applied to reduce the vibro-acoustic bridge. However, this approach did not produce a ROM that accurately captured the dynamic characteristics of the full model. Both the natural frequencies and mode shapes showed significant discrepancies. The same number of internal modes used in the successful Two-Base reduction was applied here, as previously validated in Section 4.2. This outcome may be due to numerical instabilities when computing the eigenvalues and eigenvectors of large sparse nonsymmetric matrices. Alternatively, it is possible that a greater number of internal modes is required for this model, which would significantly increase the computational cost. Further investigation is needed to determine the exact cause.

To advance the reduction analysis, the Two Base reduction method is now evaluated with damping included in the system. The same impedance boundary condition described in Section 3.4 is applied to the top boundary of the acoustic domain of the vibro-acoustic bridge. The results, based on the modified error metric defined in Equation 4.3, are presented in Figure 6.2 below.

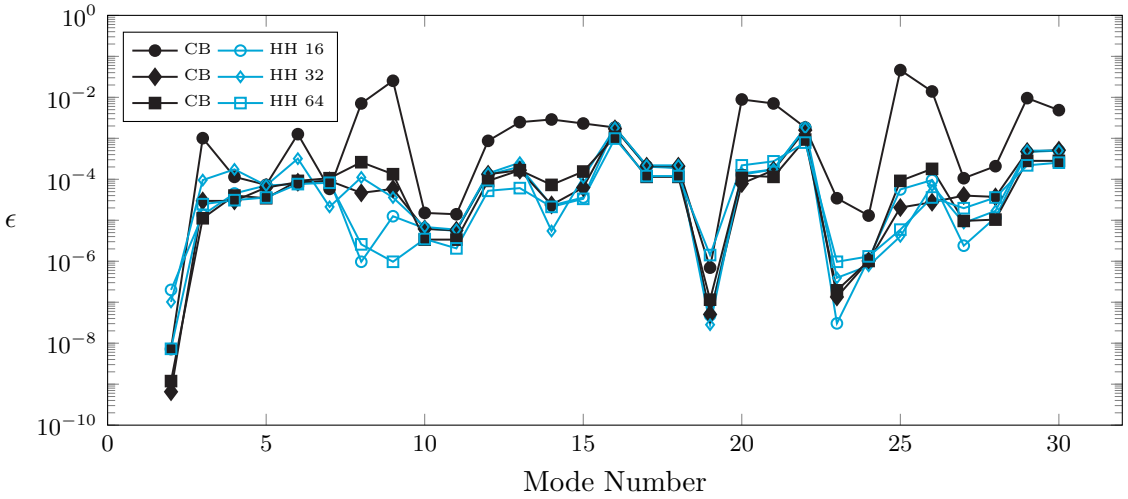

Figure 6.2: Damped Testbench Two-Base Reduction

From Figure 6.2, it can be observed that the Two-Base reduction method was successfully applied to a large scale model, where the reduction basis modes were computed using the state-space approach described in Equation 5.4. In the damped case, the HH method clearly outperformed the CB method. Moreover, aside from the CB case with 16 internal modes, the number of internal

modes had little impact on the overall accuracy. This suggests that the error is, once again, likely limited by the purely structural components, which have a stronger influence on the assembly's natural frequencies.

6.2. Coupling Strength Analysis

The acoustic domain conditions used so far have been representative of typical scenarios relevant to ASML, where the fluid is modelled as air. This choice reflects the fact that, as explained in Section 1.2, these machines do not operate in a vacuum. Various disturbances and operating conditions can generate pressure fields within the machine, which may interact with structural components (acoustic-structure interaction).

However, the interaction between air and metal is generally considered a weak coupling between the acoustic and structural domains [3, 4]. To investigate whether the effectiveness of the Two-Base reduction method is limited by this coupling strength, a limiting case is proposed. In this case, air is replaced by a fuel oil vapour, which has approximately ten times the density and three times the propagating speed of sound compared to air. The comparison of the relative frequency error for the assembly using the two different fluids is shown in Figure 6.3. Both models were reduced using the HH method with 64 internal modes.

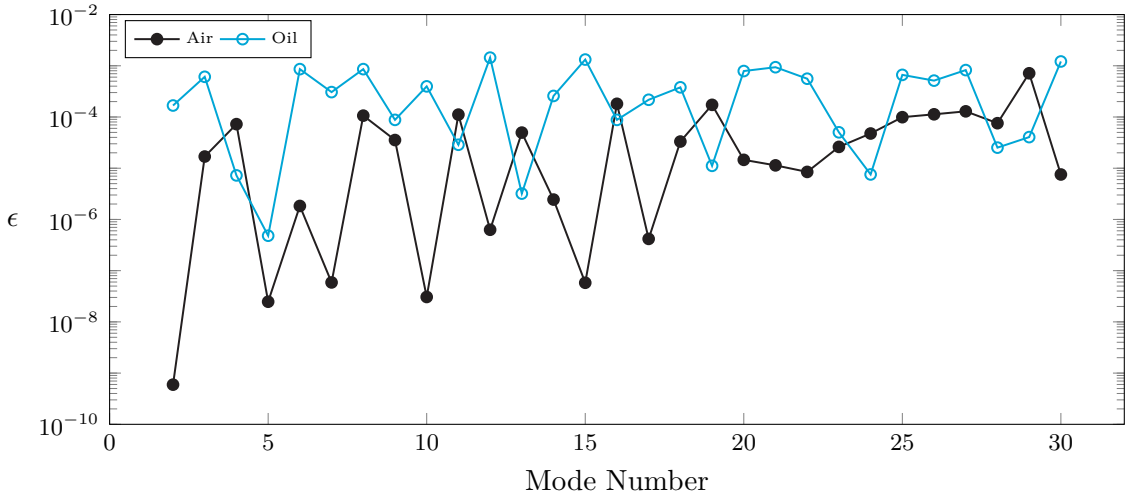

Figure 6.3: Coupling strength analysis

It can be observed that even under higher coupling conditions, the Two-Base Reduction method successfully generated a ROM that faithfully approximates the FOM. However, as shown in Figure 6.3, the accuracy is slightly reduced compared to the air-filled case, which exhibited a generally lower error. Furthermore, the eigenvalue solver exhibited increased difficulty in computing the modes for the reduction bases, as reflected in the computational cost analysis presented in Section 6.4.

6.3. Harmonic Excitation

Up to this point, the results, both for the Academic model and the Case Study, have focused on how accurately the ROM replicates the FOM. The next step is to evaluate the ROM's performance in the context of specific analyses, directly comparing it with the FOM. This comparison will demonstrate the practical value of using ROMs for vibro-acoustics, and in the context of integrating with ASML's modular modelling framework.

In all the plots presented in this section, the purely structural substructures are the same as those used in Section 6.1, while the vibro-acoustic bridge was reduced using the Two-Base HH reduction method. The CB alternative was only used for the computational effort analysis discussed in Section 6.4. As explained in Section 5.4, this analysis involves applying a harmonic excitation and observing a specific system response. To avoid explicitly inverting the full dynamic impedance matrix $\mathbf{Z}$, MATLAB's matrix left division operator is used, as shown below.

$$\hat{\boldsymbol{x}} = \mathbf{Z} \setminus \boldsymbol{f} \tag{6.1}$$

For a given harmonic excitation vector $\boldsymbol{f}$, Equation 6.1 calculates the response across all degrees of freedom (DoFs). In this case, a unit harmonic point source is applied at the node `p_bridge_1`, and the resulting displacements are measured at the nodes `uy_thinplate_py_4` and `ux_lbracket_if5`. These nodes correspond to Remote Points, ensuring they are present in both the ROMs and FOMs. Figure 6.4 illustrates the location of the excitation source and the two measurement points on the full model.

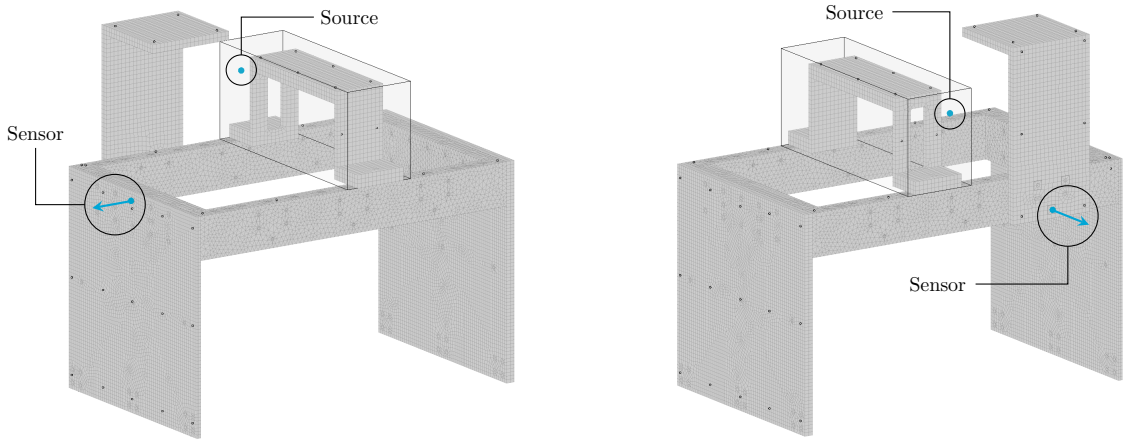

(a) Thin-plate: y-displacement

(b) L-bracket: x-displacement

Figure 6.4: Sensing positions for the point source

The response vector computed using Equation 6.1 is frequency-dependent, meaning that a harmonic excitation at each drive frequency yields a distinct system response. To examine how the response evolves across frequencies, a range from 50 to 250 Hz is analysed. This range is discretized into 630 equally spaced frequency points, at which the system response is evaluated. To ensure that both the dominant acoustic and structural modes fall within this frequency range, the length of the acoustic domain was extended to three meters, as described in Section 5.1.1. This adjustment allows the analysis to capture both acoustic and structural dominant modes within the specified frequency range.

The same frequency steps are used for both the reduced-order model (ROM) and the full-order model (FOM). For the point source excitation, the y-displacement of the top-center Remote Point on the thin plate (`uy_thinplate_py_4`) and the x-displacement of the bottom-center Remote Point on the L-bracket (`ux_lbracket_if5`) are plotted in the frequency response

functions (FRFs) shown in Figure 6.5 and Figure 6.6, respectively. The total computation time required for this analysis is recorded and discussed in Section 6.4.

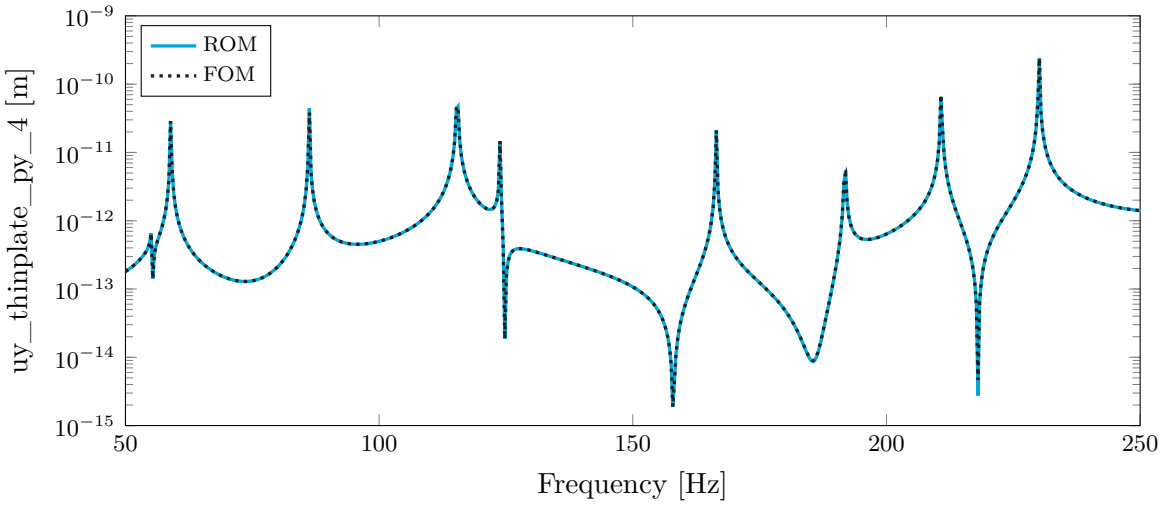

Figure 6.5: Thin-plate response to point source excitation

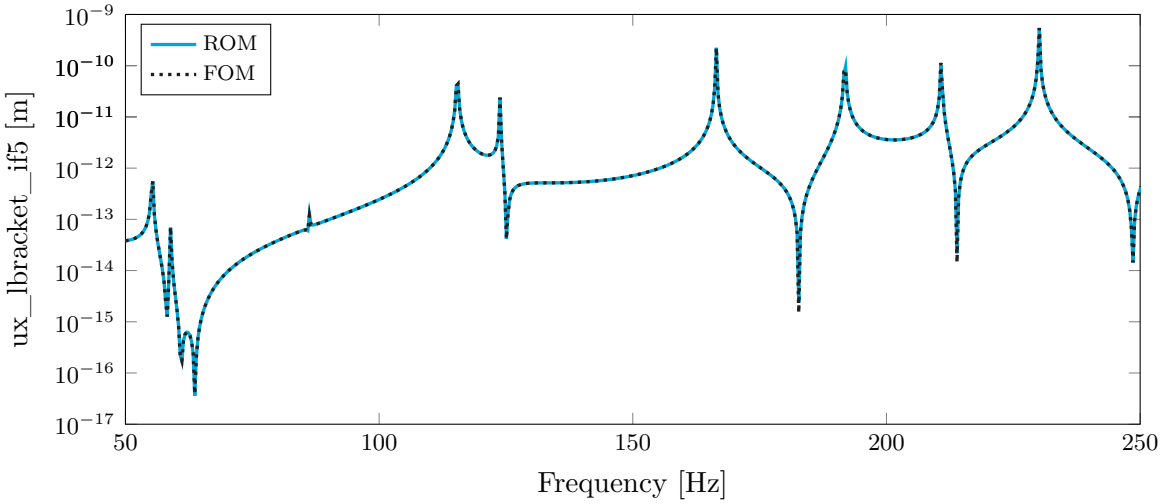

Figure 6.6: L-bracket response to point source excitation

Figure 6.5 and Figure 6.6 clearly demonstrate that the ROM in the presence of an acoustic source accurately determines the generated displacements at various points. Although only two responses are shown, the ROM results are consistent with the FOM across all Remote Points. While the FOM requires significantly more computation time, it provides responses for the full set of DoFs. In contrast, the ROM only retains data for the predefined Remote Points. Therefore, it is essential to include all nodes of interest in the boundary set during model reduction, to ensure they remain accessible for such analyses.

The presence of dominant acoustic modes, corresponding to the first, fourth, sixth, and ninth peaks in Figure 6.5 and Figure 6.6, confirms that the vibro-acoustic substructure was successfully coupled to the larger structural system. These peaks indicate that disturbances in the acoustic domain can influence structural components that are not directly part of the acoustic medium.

For a fair comparison between the ROM and the FOM, it is essential to use the same frequency vector discretisation. This is because small variations near resonance peaks can lead to significant differences in response magnitude. This effect is illustrated in Figure 6.8, where the FOM was evaluated using a coarser frequency vector, denoted as $\boldsymbol{\omega_2}$, which contains only 315 samples within the same frequency range. This reduced discretization was chosen to decrease computation time. Figure 6.8 shows the same response at the `uy_thinplate_py_4` Remote Point, but this time resulting from a point force applied to the bridge, as illustrated in Figure 6.7.

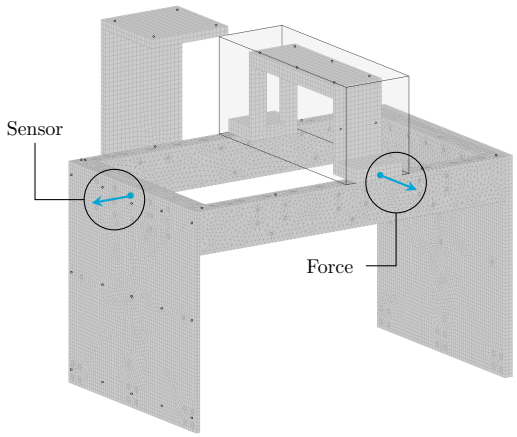

Figure 6.7: Different excitation on the Testbench model

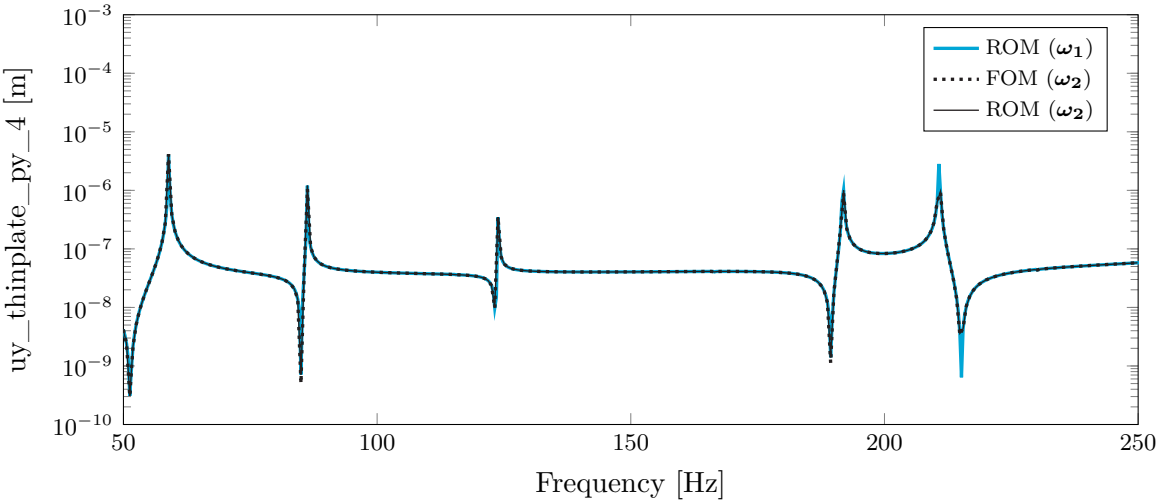

Figure 6.8: Point force excitation

From Figure 6.8, it can be observed that when comparing the ROM evaluated with 630 frequency samples ($\boldsymbol{\omega_1}$) to the FOM evaluated with the coarser $\boldsymbol{\omega_2}$ vector, the overall responses align well. However, near the resonance peaks at 191 Hz and 210 Hz, the ROM appears to exhibit higher amplitudes. This discrepancy is not due to model inaccuracy but rather a discretization effect. When the ROM is also evaluated using the $\boldsymbol{\omega_2}$ frequency vector, the responses match more closely, confirming that the difference is purely due to frequency resolution. Note that this specific

point force excitation does not activate acoustic modes.

Coupling Strength

In Section 6.2, it was demonstrated that replacing the fluid in the acoustic domain with one that increases the coupling strength in the acoustic-structure interaction results in a ROM with higher error compared to the case using air. To extend this analysis, Figure 6.9 shows the response of the thin plate to a point source, depicted in Figure 6.4(a), when fuel oil vapour is used as the fluid.

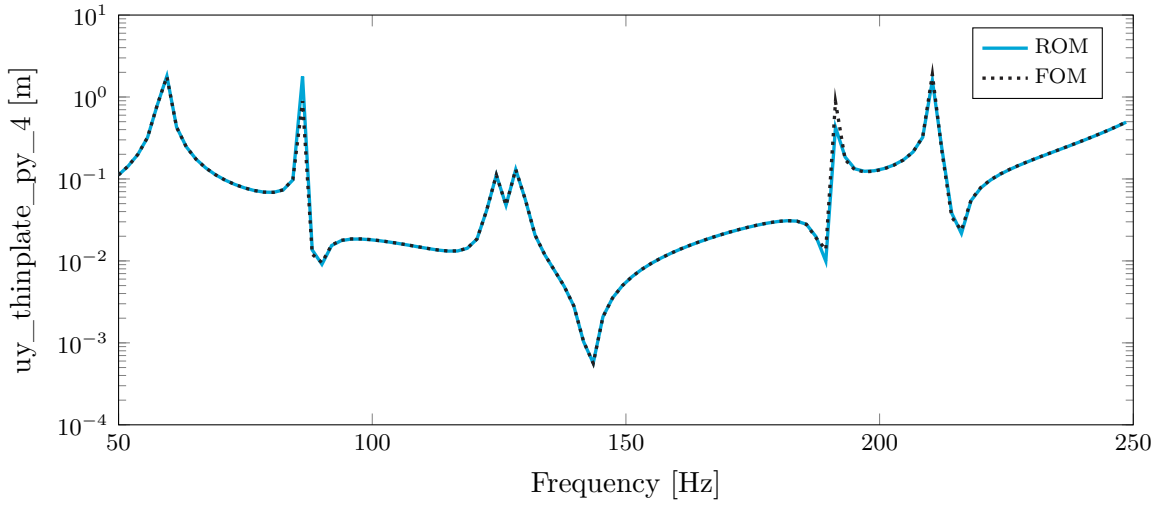

Figure 6.9: Point force excitation

As shown in Figure 6.9, the FRF appears less smooth than those obtained in the previous analysis. This is due to the significantly higher computational time required to solve Equation 6.1 at each frequency when using fuel oil vapour, compared to air. To manage computational resources, the number of frequency samples was reduced to 105 points over the same frequency range. As detailed in Section 6.4, the computational effort for these 105 samples was comparable to that required for 630 frequency samples in the air case. This highlights the importance of the ROM in cases involving strong vibro-acoustic coupling.

Although the stronger coupling and increased computational effort present challenges, they do not compromise the accuracy of the solution. As shown in Figure 6.9, the ROM accurately captures the expected response of the FOM. Moreover, it can be observed that, under the same unit source, the displacement magnitude was considerably higher when fuel vapour was used instead of air. Some discrepancies appear near the resonance peaks, which are attributed to discretisation errors. This occurs because there is a higher error between the resonance frequencies of the ROM and FOM, as illustrated in Figure 6.3. Therefore, with the same frequency vector, some frequency points may fall closer to the true resonance peaks than others, leading to localised differences in accuracy.

Reciprocity

A common routine for validating FEM models is reciprocity analysis, as it helps identify issues such as incorrect model implementation or poor mesh quality. This process involves applying an excitation at node A while measuring the response at node B. The procedure is then repeated with the excitation applied at node B and the response measured at node A. The results are compared to assess consistency. In the case shown in Figure 5.6(a), the roles of excitation and

sensing are reversed. A unit force is applied to the thin-plate, and the pressure is measured at the node in the acoustic cavity where the original point source was located.

To retrieve the pressure at the original source location (`p_bridge_1`), the DoF corresponding to the sensor position in the acoustic domain ($\boldsymbol{x_p}$) must be included in the set of boundary DoFs, as shown in Equation 5.11. For the vibro-acoustic bridge, this includes all DoFs associated with both interface Remote Points ($\boldsymbol{x_{rp1}}$ and $\boldsymbol{x_{rp2}}$), as well as the sensor position $\boldsymbol{x_p}$. This ensures it is retained in the ROM. It is possible to apply a force to a DoF that is not retained in the CMS method using the transformation $\mathbf{T}^\top \boldsymbol{f}$ (see Equation (1.5)). However, it is not possible to obtain the response at a DoF that was excluded from the reduced set. Figure 6.10 shows the pressure at `p_bridge_1` for the harmonic point force.

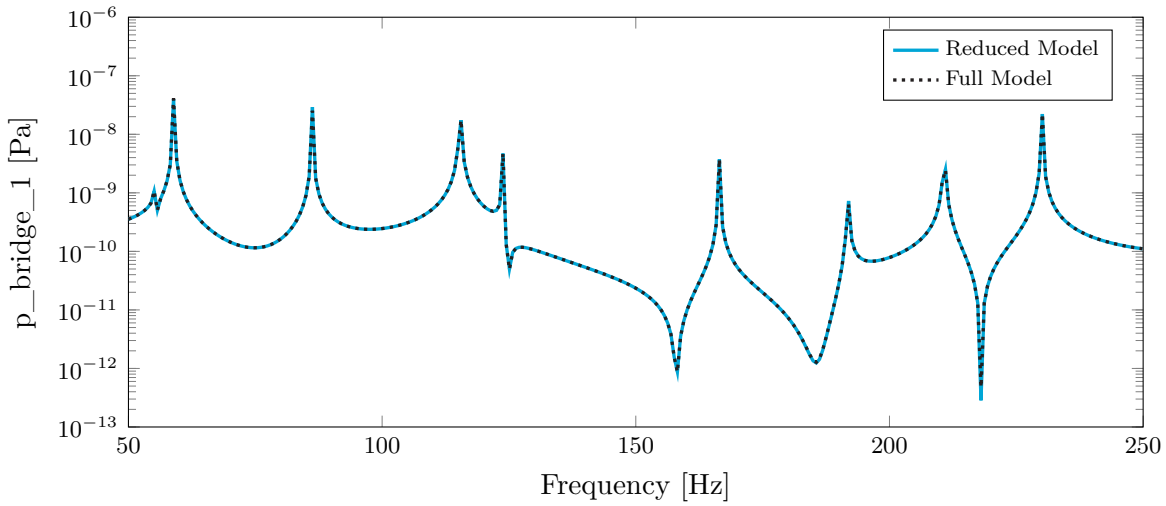

Figure 6.10: Point Source Excitation

As anticipated from Figure 6.5, the ROM also closely matches the FOM in Figure 6.10. However, to evaluate reciprocity, the displacement response of the thin-plate at `uy_thinplate_4` for the point source must be compared with the pressure response of the bridge at `p_bridge_1` for the point force. Figure 6.11 presents the ROM FRFs for these two cases.

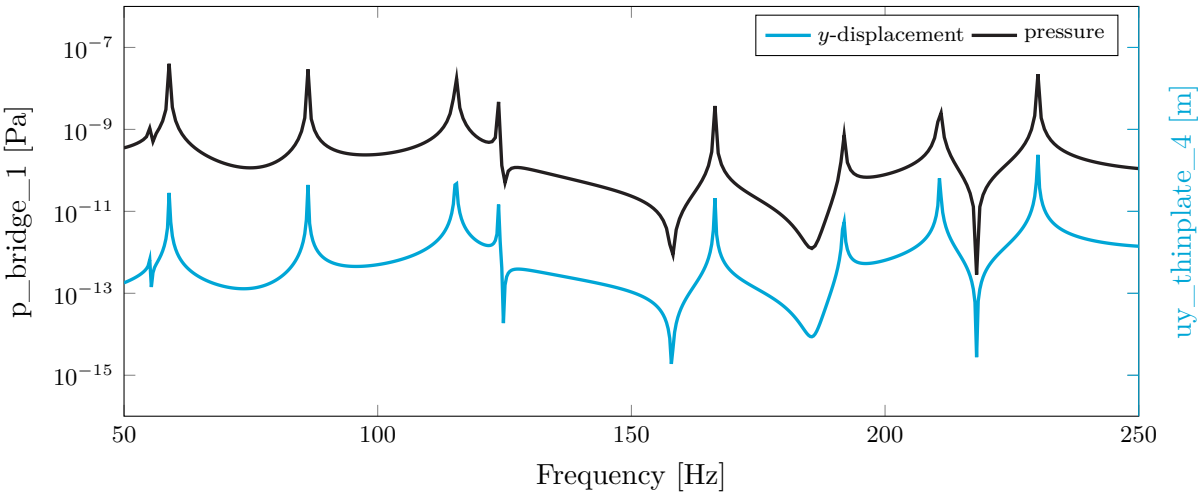

Figure 6.11: Reciprocity

Figure 6.11 shows that the responses do not exactly match, which is expected in a multi-physics model like the Testbench. In such cases, the FRFs represent different physical quantities, displacement and pressure, and therefore have different units, making direct comparison invalid. However, Figure 6.11 also demonstrates that the overall trends across the frequency range are very similar. The FRFs exhibit comparable behavior in terms of shape and resonance patterns, differing only in magnitude.

Since reciprocity analysis is not ideally suited for comparing responses across different physical domains (i.e., in unsymmetric multi-physics systems), it can instead be verified using two purely structural nodes: `uy_thinplate_4` and `ux_lbracket_if5`. Figure 6.12 illustrates the reciprocity between these two structural points on the thin plate and the L-bracket.

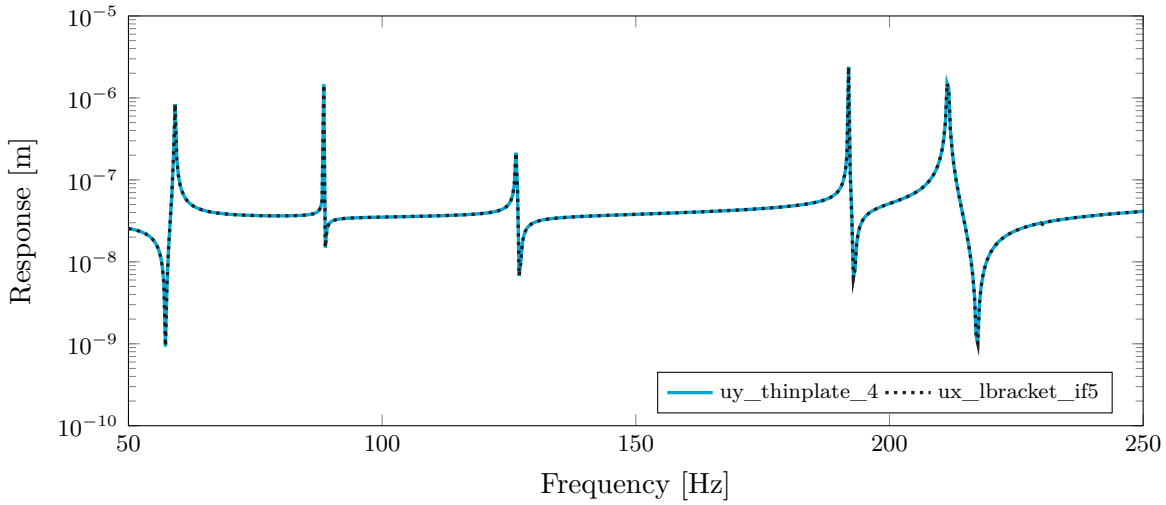

Figure 6.12: Reciprocity: Thin-plate & L-bracket

Now that the comparison is made between FRFs with matching physical units, both representing structural displacements, Figure 6.12 confirms that the model exhibits the expected reciprocity. By focusing on two structural nodes, the analysis avoids the complications introduced by comparing across different physical domains. Figure 6.12 shows that the responses at these two points are symmetric when the excitation and measurement locations are switched.

6.4. Computational Effort

This section provides a quantitative comparison of the computational effort required by various Reduced Order Models (ROMs) in relation to the Full Order Model (FOM). The analysis focuses on the total computing time needed to obtain the harmonic response, while also accounting for the cost associated with generating each ROM. Additionally, the performance of the different reduction techniques is evaluated relative to one another. The objective is to assess the overall efficiency gains achieved through model reduction and to offer insights into the compromise between computational cost and model accuracy.

For this analysis, all reduced models were constructed using the full set of constraint modes along with 64 internal modes. As outlined in Section 5.1.2, the purely structural substructures were consistently reduced using the traditional HH method across all cases. The first scenario examined corresponds to the point source excitation depicted in Figure 6.4(a). The computational times required to compute the harmonic responses across 630 frequency samples are presented in Table 6.1.

Table 6.1: Computational effort: harmonic response to point source, 630 samples

Model	Frequency Discretisation	Computational Effort		File Size
		Reduction	Calculation	
Reduced CB	630	51.55 s	7.31 s	3.02 MB
Reduced HH		69.07 s	7.37 s	3.01 MB
Full Model		∼19 h		5.79 GB

Table 6.1 clearly demonstrates the importance of using ROMs in vibro-acoustic analyses. For the ROMs, the harmonic response was computed in just around a minute, including the time required to generate the ROMs. In contrast, the FOM required over nineteen hours to complete the same analysis. It is important to note that this extended runtime for the FOM applies specifically to the MATLAB implementation using exported system matrices. If the harmonic analysis were performed directly in ANSYS, the computation time would be significantly lower due to ANSYS's proprietary and highly efficient matrix reordering algorithms. However, to protect their intellectual property, ANSYS does not apply these optimisations to the exported matrices, resulting in less efficient performance in external environments.

Nevertheless, performing the harmonic analysis in MATLAB is necessary to maintain consistency with ASML's way of working. The mechanical analysts create the models in ANSYS, but these models are later used by the System Dynamics group for further analysis. This includes evaluating FRFs, verifying set-points, analysing responses to disturbances, and incorporating control mechanisms. Since the System Dynamics group works exclusively in MATLAB, ensuring compatibility with the MATLAB environment is essential for effective collaboration and analysis.

The reduction time reported in Table 6.1 corresponds to the generation of the ROM for the vibro-acoustic bridge using the Two-Base Reduction technique. The reduction times for the purely structural components are not included, as these are significantly more efficient to compute due to the symmetry of their system matrices and it is assumed that these models are already readily available. Assembly time is also excluded, as it is negligible in comparison ($\ll 1$ second). An additional advantage of using ROMs is the substantial reduction in storage requirements. In this case, the ROM solution file is roughly two thousand times smaller than the FOM equivalent, as it stores results only at points of interest rather than at all DoFs.

For the vibro-acoustic bridge, it was observed that the HH Two-Base Reduction method was slightly more computationally expensive to generate than the CB alternative. This result contrasts with the findings for the Academic Model presented in Table 4.1, where the HH method was more efficient. However, the difference in generation time is not prohibitive, and the improved performance of the HH method justifies its use despite the slightly higher cost. As a result, all subsequent analyses were conducted using the HH-based ROM. The computational effort required to obtain the harmonic response to the point force excitation, illustrated in Figure 6.7, is summarized in Table 6.2.

Table 6.2: Computational effort: harmonic response to point force, 315 samples

Model	Frequency Discretisation	Computational Effort		File Size
		Reduction	Calculation	
Reduced Model	315	69.07 s	4 s	1.47 MB
Full Model		∼8 h		2.90 GB

It can be observed that the point force excitation required less computational effort compared to the point source case. As explained in Section 6.3, the response was computed for only half the number of frequency samples used in Table 6.1. This further highlights the advantage of using ROMs, which can reduce simulations that would otherwise take several hours to just around a minute. The reduction in solution file size was comparable to that observed in Table 6.1.

The final case analysed involves the high coupling strength model, in which the air in the acoustic domain was replaced with fuel oil vapour. This change significantly increased the computational effort required to compute the response to the point source excitation, as shown in Table 6.3. Due to this increased cost, the number of frequency samples was further reduced to 105.

Table 6.3: Computational effort for the fuel oil vapour case, 105 samples

Model	**Frequency Discretisation**	**Computational Effort**		**File Size**
		Reduction	Calculation	
Reduced Model	105	386.4 s	1.17 s	513 KB
Full Model		∼14 h		987 MB

This change in fluid significantly increases the computational cost of the vibro-acoustic model for both the ROM and the FOM. For the FOM, the solve time with fuel oil vapour was similar to that of the air case, even though only one-sixth of the frequency samples were used. Additionally, the time required to generate the ROM increased by a factor of approximately 5.5. However, the total time to obtain the ROM solution, including the time needed for model reduction, remains much lower than that of the FOM. This highlights that in cases with strong fluid-structure coupling, the use of reduced order models becomes even more critical for maintaining computational efficiency.

All simulations were performed on the same local machine described in Section 4.3.

Part IV

Conclusion

7

Project Outcome

This thesis has presented a comprehensive investigation into the development and application of reduced order modelling techniques for vibro-acoustic systems, with a particular focus on making sure these models can be integrated into ASML's modular modelling framework. The motivation stemmed from the increasing importance of accurately capturing acoustic-structure interactions in high-precision lithography machines, where traditional one-way coupling approaches may fall short.

To address this, three CMS-based reduction techniques were evaluated: the Two-Base Reduction, the Nonsymmetric Reduction, and the Locally Symmetric CB Reduction. Their effectiveness in reducing the models analysed in this thesis is summarised in Table 7.1. These methods were first tested on a simplified academic model to validate their theoretical foundations and practical implementation. The Two-Base Reduction, particularly when extended with the Hintz-Herting basis, demonstrated strong performance and computational efficiency. The Nonsymmetric Reduction, while effective in the academic case, encountered limitations when applied to the industry-scale model, likely due to numerical challenges associated with large, unsymmetric matrices. The Locally Symmetric CB method, although promising in theory, failed to capture the full vibro-acoustic behaviour in the tested configurations, suggesting that further refinement or specific boundary conditions may be required.

Table 7.1: Efficacy of the analysed reduction techniques

Reduction Technique	Academic Model	Case Study
Two-Base	✓	✓
Nonsymmetric	✓	✗
Locally Symmetric	✗	✗

The transition from the Academic Model to the Testbench Model in the case study also necessitated a shift from COMSOL to ANSYS as the software used to develop the full FEM models of the vibro-acoustic substructure. This change was essential to ensure compatibility with ASML's existing modular modelling way of work, primarily due to the use of Remote Points. Consequently, several adjustments had to be made to the procedure originally used for the Academic Model.

By incorporating Remote Points and adapting the reduction and coupling procedures to align with ASML's established practices, the revised methodology was successfully applied to a representative, industry-scale case study: the Testbench Model. In this model, the reduced vibro-acoustic representation of the bridge was seamlessly integrated with existing purely structural components. The resulting reduced-order models achieved significant reductions in computational time, transforming harmonic excitation simulations from tasks that previously took hours into ones

completed in under a minute. The methodology demonstrated robustness even under conditions of strong coupling.

In conclusion, this work has effectively extended ASML's modular modelling capabilities to encompass two-way coupled acoustic-structural systems. The developed approach not only improves simulation efficiency but also paves the way for more accurate system-level analyses and design optimisations. With further refinement and validation against experimental data, this methodology holds considerable promise for future integration into ASML's engineering workflows and beyond.

8

Recommendations

This chapter presents a concise overview of the key recommendations for the continued development of this research project.

Further investigation is recommended to validate the assumption that structural and acoustic elements must be present at the boundaries for the Locally Symmetric CB Reduction technique to effectively reduce a vibro-acoustic substructure. Therefore, a starting point could be to replicate the results from the academic model presented in [7], to assess whether this assumption holds true in practice.

The Nonsymmetric Reduction method was successfully implemented in the two-dimensional academic model analysed in this thesis. However, its application to the industry-scale case study was unsuccessful. Further research is needed to determine whether this limitation stems from the use of unsymmetric solvers to compute the the modes in the reduction base, or if the method inherently requires a larger number of internal modes to accurately capture the system dynamics.

Now that the Two-Base Reduction technique has been validated, it would be valuable to apply it to the reduction of an actual machine substructure within ASML. This would enable a direct comparison between the two-way coupled approach and the currently used one-way coupled method, using experimental data for validation. Additionally, incorporating a method to visualize the mode shapes of the reduced model would enhance interpretability. This could be achieved by adding extra boundary nodes throughout the structure, retained solely for visualization purposes. Further consideration is needed to develop a method for displaying the acoustic mode shapes in the reduced model effectively.

References

[1] VIBES Technology. *What is NVH?* URL: https://www.vibestechnology.com/academy/methodology/what-is-nvh/#:~:text=NVH%20engineers%20work%20on%20a,with%20sound%20and%20vibration%20engineering. (visited on 26/11/2024).

[2] Robert Benjamin Davis. 'Techniques to Assess Acoustic-Structure Interaction in Liquid Rocket Engines'. PhD thesis. Duke University, 2008.

[3] Wim Desmet. 'A wave-based prediction technique for coupled vibro- acoustic analysis'. PhD thesis. Katholieke Universiteit Leuven, 1998.

[4] Axel van de Walle. 'The Power of Model Order Reduction in Vibroacoustics and its Applications in Model-based Sensing'. PhD thesis. Katholieke Universiteit Leuven, 2018.

[5] A. Craggs. 'The transient response of a coupled plate-acoustic system using plate and acoustic finite elements'. In: *Journal of Sound and Vibration* 15.4 (1971), pp. 509–528. DOI: https://doi.org/10.1016/0022-460X(71)90408-1.

[6] O.C. Zienkiewicz et al. 'Coupled vibrations of a structure submerged in a compressible fluid'. In: *Symposium on Finite Element Techniques*. 1969.

[7] Umut Tabak. 'Methods for efficient analysis of vibro-acoustic problems'. PhD thesis. Delft University of Technology, 2020.

[8] P.L.C van der Valk. 'Model Reduction & Interface Modeling in Dynamic Substructuring'. MSc thesis. Delft University of Technology, 2010.

[9] J.J van Steen. 'Comparison of model order reduction techniques for interface dynamics'. MSc thesis. Eindhoven University of Technology, 2020.

[10] Robert J. Guyan. 'Reduction of Stiffness and Mass Matrices'. In: *AIAA Journal* 3.2 (1965), pp. 380–380. DOI: 10.2514/3.2874.

[11] Roy R. Craig et al. 'Coupling of Substructures for Dynamic Analyses'. In: *AIAA Journal* 6.7 (1968), pp. 1313–1319. DOI: https://doi.org/10.1016/0022-460X(71)90408-1.

[12] D.N. Herting. 'A general purpose, multi-stage, component modal synthesis method'. In: *Finite Elements in Analysis and Design* 1.2 (1985), pp. 153–164. DOI: https://doi.org/10.1016/0168-874X(85)90025-3. URL: https://www.sciencedirect.com/science/article/pii/0168874X85900253.

[13] B. Besselink et al. 'A comparison of model reduction techniques from structural dynamics, numerical mathematics and systems and control'. In: *Journal of Sound and Vibration* 332.19 (2013), pp. 4403–4422. DOI: https://doi.org/10.1016/j.jsv.2013.03.025. URL: https://www.sciencedirect.com/science/article/pii/S0022460X1300285X.

[14] Dennis Klerk et al. 'General Framework for Dynamic Substructuring: History, Review, and Classification of Techniques'. In: *Aiaa Journal - AIAA J* 46 (May 2008), pp. 1169–1181. DOI: 10.2514/1.33274.

[15] Rafael Engelmann. 'An investigation into low-frequency aero-vibro-acoustic mechanisms in vehicles'. PhD thesis. Technischen Universitat Wien, 2023.

[16] Matthew S. Allen et al. *Substructuring in Engineering Dynamics Emerging Numerical and Experimental Techniques*. CISM International Centre for Mechanical Sciences. Springer Cham, 2020. DOI: https://doi.org/10.1007/978-3-030-25532-9.

[17] M.V van der Seijs. 'Experimental Dynamic Substructuring'. PhD thesis. Delft University of Technology, 2016.

[18] Fabian Gruber et al. 'Comparison between primal and dual Craig-Bampton substructure reduction techniques'. In: Oct. 2015.

[19] Daniel Rixen et al. 'The Frequency Based Substructuring Method reformulated according to the Dual Domain Decomposition Method'. In: Jan. 2006.

[20] R S O Dias et al. 'Lagrange Multiplier State-Space Substructuring'. In: *Journal of Physics: Conference Series* 2041.1 (Oct. 2021), p. 012016. DOI: 10.1088/1742-6596/2041/1/012016. URL: https://dx.doi.org/10.1088/1742-6596/2041/1/012016.

[21] R.S.O. Dias et al. 'On the use of Lagrange Multiplier State-Space Substructuring in dynamic substructuring analysis'. In: *Mechanical Systems and Signal Processing* 180 (2022), p. 109419. DOI: https://doi.org/10.1016/j.ymssp.2022.109419. URL: https://www.sciencedirect.com/science/article/pii/S0888327022005416.

[22] Axel Van de Walle et al. 'Stability-preserving model order reduction for time-domain simulation of vibro-acoustic FE models: Stability-preserving model order reduction for time-domain simulation of vibro-acoustic FE models'. In: *International Journal for Numerical Methods in Engineering* 109 (June 2016). DOI: 10.1002/nme.5323.

[23] Leopoldo de Oliveira et al. 'A CAE modeling approach for the analysis of vibro-acoustic systems with distributed ASAC control'. In: *Proceedings of ISMA2006: International Conference on Noise and Vibration Engineering* 1 (Jan. 2006).

[24] Leopoldo de Oliveira et al. 'A State-Space Modeling Approach for Active Structural Acoustic Control'. In: *Shock and Vibration* 16 (Jan. 2009), pp. 607–621. DOI: 10.3233/SAV-2009-0492.

[25] Soo Min Kim et al. 'A strongly coupled model reduction of vibro-acoustic interaction'. In: *Computer Methods in Applied Mechanics and Engineering* 347 (2019), pp. 495–516. DOI: https://doi.org/10.1016/j.cma.2018.12.029. URL: https://www.sciencedirect.com/science/article/pii/S0045782518306285.

[26] R. Puri et al. 'Two-way coupled structural acoustic optimization via model order reduction (MOR)'. In: *Proceedings of ISMA2006: International Conference on Noise and Vibration Engineering* 8 (Jan. 2006).

[27] R. Srinivasan Puri et al. 'Reduced order fully coupled structural–acoustic analysis via implicit moment matching'. In: *Applied Mathematical Modelling* 33.11 (2009), pp. 4097–4119. DOI: https://doi.org/10.1016/j.apm.2009.02.016. URL: https://www.sciencedirect.com/science/article/pii/S0307904X09000420.

[28] R. Puri et al. 'Compact structural-acoustic coupled models via model order reduction (MOR)'. In: 4 (Jan. 2006).

[29] Stijn Donders et al. 'CAE technologies for efficient vibro-Acosutic vehicle design modification and optimization'. In: *23rd International Conference on Noise and Vibration Engineering 2008, ISMA 2008* 7 (Jan. 2008).

[30] Umut Tabak et al. 'Globally enriched substructuring techniques for vibro-acoustic simulation'. In: *Linking Models and Experiments, Volume 2*. Ed. by Tom Proulx. New York, NY: Springer New York, 2011, pp. 263–280.

[31] Jan Herrmann et al. 'Substructuring including interface reduction for the efficient vibro-acoustic simulation of fluid-filled piping systems'. In: *Mechanical Systems and Signal Processing* 24.1 (2010), pp. 153–163. DOI: https://doi.org/10.1016/j.ymssp.2009.05.003. URL: https://www.sciencedirect.com/science/article/pii/S0888327009001435.

[32] Jan Herrmann et al. 'Model Reduction and Substructuring Techniques for the Vibro-Acoustic Simulation of Automotive Piping and Exhaust Systems'. In: *Structural Dynamics, Volume 3*. Ed. by Tom Proulx. New York, NY: Springer New York, 2011, pp. 1055–1064.

[33] Eskil Lindberg et al. 'Component Mode Synthesis Using Undeformed Interface Coupling Modes to Connect Soft and Stiff Substructures'. In: *Shock and Vibration* 20 (Jan. 2013). DOI: 10.3233/SAV-2012-0734.

[34] Eskil Lindberg et al. 'A vibro-acoustic reduced order model using undeformed coupling interface substructuring – Application to rubber bushing isolation in vehicle suspension systems'. In: *Applied Acoustics* 78 (Apr. 2014), pp. 43–50. DOI: 10.1016/j.apacoust.2013.11.001.

[35] Valentin Meyer et al. 'A condensed transfer function method as a tool for solving vibroacoustic problems'. In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 230.6 (2016), pp. 928–938. DOI: 10.1177/0954406215615627.

[36] Valentin Meyer et al. 'A substructuring approach based on mechanical admittances to solve vibro-acoustic problems in the mid-frequency range'. In: *The Journal of the Acoustical Society of America* 146 (Oct. 2019), pp. 3006–3006. DOI: 10.1121/1.5137407.

[37] F. Dumortier et al. 'Vibroacoustic subtractive modeling using a reverse condensed transfer function approach'. In: *Journal of Sound and Vibration* 499 (2021), p. 115982. DOI: https://doi.org/10.1016/j.jsv.2021.115982. URL: https://www.sciencedirect.com/science/article/pii/S0022460X21000547.

[38] A.A. Lakis et al. 'Prediction of the response of a cylindrical shell to arbitrary or boundary-layer-induced random pressure fields'. In: *Journal of Sound and Vibration* 25.1 (1972), pp. 1–27. DOI: https://doi.org/10.1016/0022-460X(72)90592-5.

[39] Michael J. Allen et al. 'Integration of finite element and boundary element methods for calculating the radiated sound from a randomly excited structure'. In: *Computers & Structures* 77.2 (2000), pp. 155–169. DOI: https://doi.org/10.1016/S0045-7949(99)00208-4.

[40] Matthias Maess et al. 'Substructuring and model reduction of pipe components interacting with acoustic fluids'. In: *Mechanical Systems and Signal Processing* 20.1 (2006), pp. 45–64. DOI: https://doi.org/10.1016/j.ymssp.2005.02.008. URL: https://www.sciencedirect.com/science/article/pii/S0888327005000270.

[41] Paul Regtien et al. '9 - Acoustic sensors'. In: *Sensors for Mechatronics (Second Edition)*. Ed. by Paul Regtien et al. Second Edition. Elsevier, 2018, pp. 267–303. DOI: https://doi.org/10.1016/B978-0-12-813810-6.00009-4. URL: https://www.sciencedirect.com/science/article/pii/B9780128138106000094.

[42] ANSYS. *Mechanical Users Guide, Chapter 12: Remote Points*. URL: https://ansyshelp.ansys.com/public/account/secured?returnurl=/Views/Secured/corp/v242/en/wb_sim/ds_remote_point.html (visited on 19/05/2025).

[43] Fabian M. Gruber et al. 'Comparison of Craig-Bampton Approaches for Systems with Arbitrary Viscous Damping in Dynamic Substructuring'. In: *Dynamics of Coupled Structures, Volume 4*. Ed. by Andreas Linderholt et al. Cham: Springer International Publishing, 2018, pp. 35–49.

A

Appendix

A.1. Mode Shape Visualisation

In Figure A.1, the mode shapes corresponding to the sixth natural frequency are shown for both the assembly of the two full substructures and the reference model from Figure 3.1. The colour maps reveal that the scaling of the acoustic and structural components is not consistent between the two systems. This discrepancy highlights the need to compute the MAC separately for each physical domain, as explained in Section 4.1.

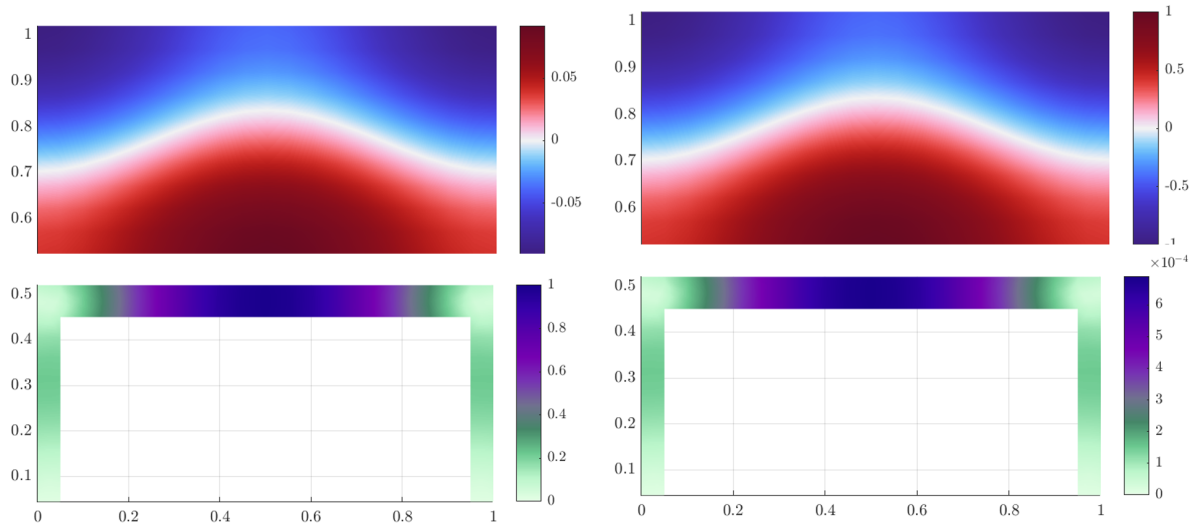

(a) Assembly of substructures

(b) Reference

Figure A.1: Sixth mode shape of the full assembly

A.2. Locally Symmetric CB Modes

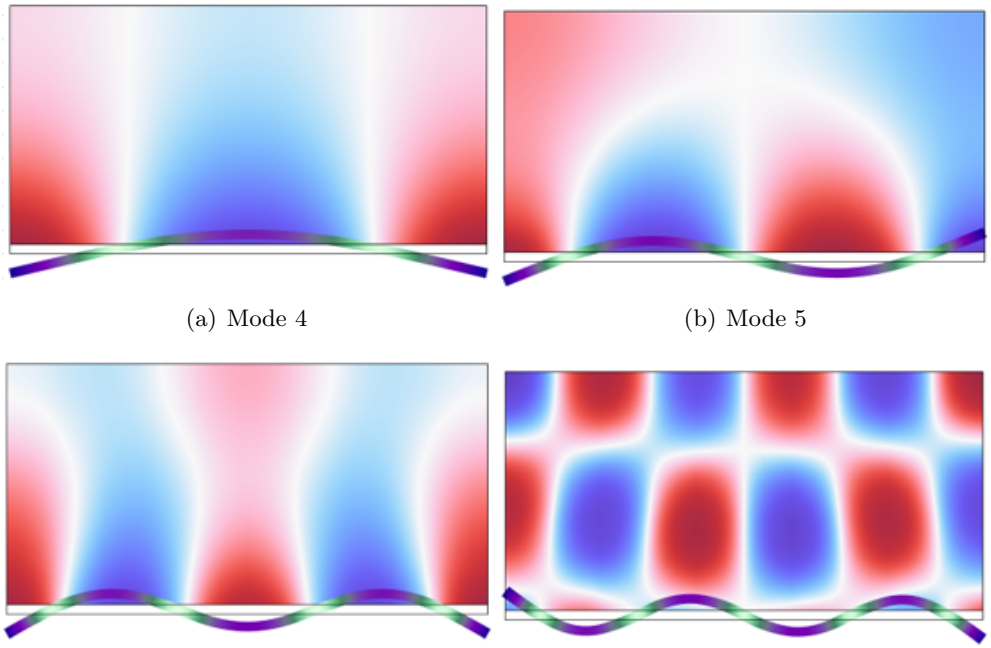

(a) Mode 4 (b) Mode 5

(c) Mode 6 (d) Mode 7

Figure A.2: First non rigid-body modes obtained with Locally Symmetric Reduction