




Gecko

Design for *IGA*-type
discretization workflows

Gecko Technical Report 2

DC04 - Juan Ignacio Camarotti



This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101073106
Call: HORIZON-MSCA-2021-DN-01

Funded by the
European Union

Deliverable Information sheet

Version	Date	Author	Document history/approvals
0		Marco Zuñiga	Draft version circulated to partners
1	20.01.2026	Juan I. Camarotti	Adding information
2	30.01.2026	Ricky Aristio	Revising
3	30.01.2026	Juan I. Camarotti	Final document

Executive summary

This technical report summarizes the research activities carried out in 2025 within the framework of the GECKO Project, focusing on the application of high-fidelity isogeometric structural models (IBRA) within partitioned multi-physics co-simulation frameworks. The overarching objective of this work is to enable robust, accurate, and scalable coupling strategies for coupled problems, where different solvers and fidelity levels must interact consistently across non-matching interfaces.

Building upon the foundational work of the previous year, centered on mapping strategies, this year's efforts advanced toward the integration of complete end-to-end co-simulation workflows that combine IBRA-based structural models with different fluid solvers. The emphasis was placed on demonstrating how isogeometric discretizations (e.g., Kirchhoff–Love shell formulations and multipatch representations) introduce specific challenges and opportunities in coupled simulations.

A first major research pillar addressed aeroelastic co-simulation of a simplified wing configuration, combining a structural solver in Kratos Multiphysics (using both classical FEM and IBRA discretizations) with an external aerodynamic solver based on the Vortex Lattice Method (VLM). A partitioned co-simulation setup was constructed where structural displacements are transferred to the aerodynamic model and aerodynamic loads are mapped back to the structure. Special attention was given to the use of radial basis function (RBF) mapping to enable stable transfer across non-matching meshes and to incorporate rotation mapping, which is essential to correctly reflect the elastic contribution to the effective aerodynamic angle of attack. This framework enabled direct comparisons between FEM- and IBRA-based structural predictions in a consistent coupled environment.

A second research pillar focused on partitioned vibro-acoustic co-simulation, coupling a structural harmonic response analysis in Kratos with an external Boundary Element Method (BEM) acoustic solver (implemented in MATLAB and accessed through Python). This workflow targets efficient simulation of fluid-borne noise phenomena governed by the Helmholtz equation while allowing high-fidelity structural representations. The study explored one-way and two-way coupling configurations and examined the impact of mapping and interface treatment on the ability to reproduce resonance peaks and modal behavior, particularly for single and multipatch IBRA models.

Overall, the developments reported here contribute toward a general methodology for multi-fidelity and multi-physics coupling with IBRA, showing that the combination of advanced discretizations and partitioned co-simulation can deliver both accuracy and flexibility, provided that interface mapping and coupling strategies are designed to respect the characteristics of each participating solver.

Table of contents

Deliverable Information sheet	2
Executive summary	3
Table of contents	4
List of figures	5
List of abbreviations	6
Introduction	7
1. Monolithic vs Partitioned Approaches in MP	8
2. Partitioned Simulation of FSI Problems	9
3. Fundamental Concepts of Co-Simulation	10
4. Data Transfer Operators for CoSimulation	12
4.1 Nearest Neighbour Mapper for IGA	12
4.2 Nearest Element Mapper for IGA	13
4.3 Mortar Mapper for Body-Fitted and Unfitted Discretizations	14
4.3.1 1D IBRA-IBRA or IBRA-FEM Interface	15
4.3.2 2D IBRA-FEM Interface	15
4.4 Radial Basis Functions (RBF) Mapper	16
5. Co-simulation Framework for Aeroelastic Analysis of an IBRA Wing (IBRA+VLM)	18
5.1 Aerodynamic Loads Calculation	19
5.2 Aeroelastic Co-simulation Framework	20
5.3 Test Case Set-up	21
5.4 Linear Aeroelastic Analysis	22
5.5. Results: Linear Aeroelastic Analysis of the Minimal Wing	23
5.6 Results: Aeroelastic Lift Coefficient Comparison	24
6. Partitioned Co-simulation Framework for Vibroacoustic Systems	25
6.1 Governing Equations	25
6.2 Monolithic and Partitioned Approaches for Vibroacoustic Systems	27
6.3 Co-simulation Strategies for Partitioned Vibroacoustic Systems	27
6.3.1 Coupling Method	27
6.3.2 Data Exchange between the Solvers	29
6.3.3 Data Mapping	29
6.4 Numerical Results	30
6.4.1 Example 1	32
6.4.2 Example 2	33
6.4.3 Example 3	35
8. CONCLUSIONS	37
9. REFERENCES	38

List of figures

Figura 1: System matrix level of the monolithic approach	8
Figura 2: System matrix level of the partitioned approach	8
Figura 3: Data transfer in partitioned FSI simulations	9
Figura 4: Building blocks of co-simulation	10
Figura 5: Flow of information during data transfer from origin to destination	11
Figura 6: Nearest neighbour mapping at integration-point level	13
Figura 7: Nearest element (closest projection) mapper for IGA	14
Figura 8: Definition of the integration domain for the 1D FEM–FEM interface	15
Figura 9: Definition of the integration domain for the 1D IBRA–IBRA interface	15
Figura 10: Low-order FEM discretization (fluid) and CAD surface (structure)	15
Figura 11: Projection of the low-order element onto the CAD surface.	16
Figura 12: Aircraft design cycle	18
Figura 13: Multi-fidelity modelling spectrum for aerodynamics and structures	18
Figura 14: Wing discretization with vortex lattice method (VLM)	19
Figura 15: Aeroelastic co-simulation framework	20
Figura 16: Wing geometry and material definition	21
Figura 17: Structural load application points and aerodynamic panel discretization	22
Figura 18: Comparison of FEM structural responses under aerodynamic loading	23
Figura 19: Structural deformation of the minimal wing using an IBRA shell model	24
Figura 20: Lift coefficient comparison for rigid and elastic configurations	24
Figura 21: Data exchange architecture in partitioned vibroacoustic co-simulation	29
Figura 22: Simply supported rectangular plate subjected to a uniform surface load	31
Figura 23: SPL frequency sweep: monolithic vs. partitioned one-way coupling	31
Figura 24: First 20 natural frequencies of the rectangular plate	32
Figura 25: Rectangular plate with surface load and acoustic monopole excitation	32
Figura 26: SPL frequency response: weak vs. strong two-way vibroacoustic coupling	33

List of abbreviations

<i>CAD</i>	<i>Computer-aided Design</i>
<i>FEM</i>	<i>Finite Element Method</i>
<i>FSI</i>	<i>Fluid-Structure Interaction</i>
<i>IBRA</i>	<i>Isogeometric B-Rep Analysis</i>
<i>RBF</i>	<i>Radial Basis Functions</i>
<i>SPL</i>	<i>Sound Pressure Level</i>

Introduction

Numerical simulation has become an essential tool in aerospace and mechanical engineering, particularly for lightweight structures operating close to their performance limits. In such applications, the structural response is rarely governed by mechanics alone, but is strongly influenced by external physical fields such as aerodynamic loads or acoustic pressure waves. Reliable prediction therefore requires simulation strategies capable of capturing these multi-physics interactions without incurring prohibitive computational costs.

From a practical engineering perspective, the primary goal is often not to maximize fidelity in every physical domain, but to combine models of different complexity in a consistent and efficient manner. This motivates the use of multi-fidelity simulation frameworks, in which high-resolution structural models are coupled with reduced-order or boundary-based solvers that remain computationally affordable. Such approaches are particularly attractive for parametric studies, early-stage design, and conceptual investigations, where robustness and efficiency are as important as accuracy.

A common strategy to achieve this balance is partitioned co-simulation, where each physical subproblem is solved independently using a dedicated solver, and interaction is enforced through the exchange of interface quantities. Compared to monolithic formulations, partitioned approaches offer clear advantages in terms of modularity, solver reuse, and flexibility, while naturally supporting heterogeneous discretizations and fidelity levels. Their effectiveness, however, depends critically on the stability of the coupling strategy and on the accuracy of the data transfer between non-matching interfaces.

Within this context, particular attention is given to the integration of isogeometric analysis (IGA) for structural modeling using the IBRA framework in Kratos Multiphysics. The use of Kirchhoff–Love shell formulations and multi-patch B-Rep representations enables smooth kinematics and a direct link to CAD geometry, but also introduces additional challenges for interface consistency and data exchange in coupled simulations.

This report summarizes the developments carried out in 2025 toward complete co-simulation workflows combining IBRA-based structural models with low- and medium-fidelity external solvers. The first part focuses on aeroelastic coupling, combining Kratos (FEM and IBRA) with a Vortex Lattice Method aerodynamic solver. The second part addresses vibroacoustic coupling, linking harmonic structural response analysis in Kratos with an external Boundary Element Method solver for the Helmholtz equation. Together, these studies illustrate the capabilities and limitations of partitioned, multi-fidelity co-simulation strategies for aerospace applications.

1. Monolithic vs Partitioned Approaches in Multiphysics Simulations

In multiphysics problem-solving, two primary approaches are widely used: monolithic and partitioned strategies.

Monolithic Approach

The monolithic approach involves solving all governing equations simultaneously within a unified framework. This method leads to a more stable and robust formulation, as all physical interactions are accounted for within a single system of equations.

However, the monolithic strategy often results in a large, complex global system that can be difficult to solve, especially as the size of the problem increases. As the global system grows, the conditioning of the system tends to worsen, making it more computationally demanding. Despite these challenges, the monolithic approach offers high accuracy and robustness, making it well-suited for certain complex multiphysics problems.




Figure 1: System Matrix Level of the Monolithic Approach

Partitioned Approach

In contrast, the partitioned approach separates the problem into distinct subsystems, solving each independently while coupling them through data exchange, often in the form of boundary conditions.

This approach allows for the reuse of existing single-field solvers, a key advantage in the co-simulation philosophy, where different physical domains are solved independently. The partitioned method is modular, flexible, and scalable, enabling the integration of different solvers for various disciplines. However, it is often less robust than the monolithic approach, as the coupling between subsystems may introduce stability challenges, particularly in highly coupled or stiff problems.

Despite this, partitioned approaches are more flexible and allow for easier integration of specialized solvers, making them valuable for complex multidisciplinary simulations.




Figure 2: System Matrix Level of the Partitioned Approach

2. Partitioned Simulation of FSI Problems

Fluid-structure interaction (FSI) simulations are essential for accurately modeling and predicting the behavior of systems where fluids and structures interact dynamically. It plays a critical role in various engineering fields, including aerospace, biomechanics, and civil engineering, where structural integrity and performance depend on fluid forces. By capturing the mutual influence between fluids and structures, FSI enables the design of safer and more efficient systems, such as aircraft wings, blood flow in arteries, and bridges subjected to wind forces. Without FSI analysis, critical phenomena like vibrations, instabilities, and failure risks may be overlooked, leading to inaccurate predictions and potential structural failures.

FSI, like any multiphysics problem, can be solved using either a monolithic or partitioned approach. Monolithic methods solve the fluid and structural equations simultaneously, ensuring strong coupling but requiring complex solver development. The Partitioned method, on the other hand, is a computational approach where the fluid and structural domains are solved independently using specialized solvers, which are coupled to exchange key data at their interface.

The structural solver provides the displacement field, defining the motion and deformation of the structure, which is then used by the fluid solver to update the fluid domain and enforce the moving boundary conditions. Conversely, the fluid solver calculates the pressure field (and sometimes shear stresses), which acts as the load applied to the structure. This bidirectional exchange of displacements and pressure fields ensures the mutual influence of the fluid and structure, enabling the accurate simulation of their interaction.

Partitioned FSI offers flexibility and modularity by integrating existing solvers without significant modification, but it relies heavily on effective data transfer and coupling algorithms to maintain stability and accuracy, particularly in cases involving strong interactions or large deformations. The data transfer between the fluid and structural domains is accomplished through various data transfer operators, which ensure that critical information, such as displacement and pressure fields, is exchanged correctly. These operators are vital for ensuring the robustness of the coupling, and their efficiency will be addressed shortly.

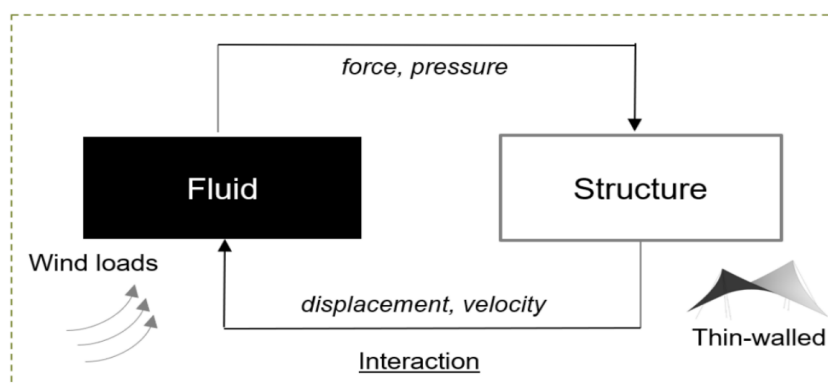


Figure 3: Data Transfer in Partitioned FSI Simulations

3. Fundamental Concepts of CoSimulation

CoSimulation is inherently a partitioned approach. In this context, multiple simulation tools or solvers, each specializing in a specific physical phenomenon, are coupled together to model complex systems involving interactions between different physical domains. Unlike the monolithic approach, where all phenomena are solved together within a single system, CoSimulation allows for the independent solution of each phenomenon.

These individual solvers are then coupled through specialized methods and tools to ensure accurate interaction between the domains. This partitioned nature of CoSimulation enables the reuse of existing solvers, offering flexibility and reducing development time, though it may require careful attention to ensure stability and accuracy, especially in cases involving strong interactions.

3.1 Building Blocks of CoSimulation

CoSimulation using the partitioned approach consists of several fundamental building blocks that are essential for ensuring effective and stable interactions between different solvers. These building blocks are synchronization, data exchange, and data transfer, as illustrated in Figure 4.

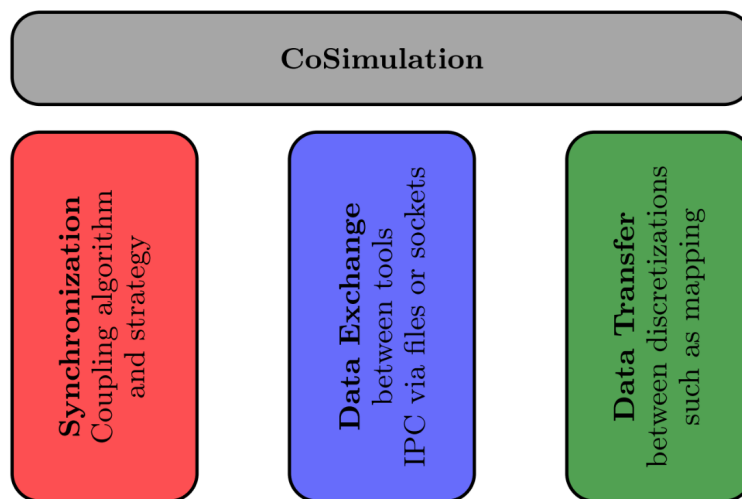


Figure 4: Building Blocks of CoSimulation (Ref. [1])

- 1) **Synchronization and Solution Techniques:** The first building block of CoSimulation, shown in Figure 4, is synchronization, which encompasses the coupling algorithm and strategy. This block determines the sequence in which the participating tools are applied, as well as the incorporation of additional components, such as relaxation techniques. The selection of these procedures greatly impacts the robustness, accuracy, performance, efficiency, and stability of the coupled simulation, making it crucial to choose them based on the specific problem being addressed.
- 2) **Data Exchange and Data Transfer:** Each coupling partner employs its own solution technique, chosen based on what best suits the specific application. Consequently, each

tool handles and stores data in its own format. For CoSimulation, it is essential to transfer data between coupling partners. Often, this requires data exchange.

Additionally, the data must be converted from one format to another, a process referred to as data transfer in this work, as shown in column three of Figure 4. The flow of information proceeds from the origin to the destination, as illustrated in Figure 5.




Figure 5: Flow of information during data transfer between solution techniques, from origin to destination, from the origin to the destination (Ref. [1])

4. Data Transfer Operators for CoSimulation

As part of this year's work, a set of data transfer (mapping) operators specifically targeting IGA-based coupled simulations was implemented and integrated into Kratos Multiphysics. These include standard approaches such as Nearest Neighbour and Nearest Element mapping, as well as higher-order techniques like Radial Basis Function (RBF) interpolation, and the Mortar Mapper.

Data transfer operators are fundamental building blocks of partitioned solution strategies within CoSimulation frameworks, as they enable the consistent exchange of interface quantities (e.g., displacements, rotations, forces, or pressures) between distinct physical domains and non-matching discretizations. In practice, different coupling scenarios and mesh/interface configurations require different mapping choices, and a wide spectrum of techniques is available, ranging from simple proximity-based mappings to variationally consistent formulations.

A critical requirement in co-simulation is to preserve the non-intrusiveness of existing single-field solvers. This means that the coupling procedure should avoid modifications to the internal formulation of each participating code and instead enforce interaction through well-defined interface variables and external communication. In this setting, robust and accurate mapping operators play a central role in achieving stable and reliable coupled simulations.

In the following, the implemented data transfer operators are discussed.

4.1 Nearest Neighbour Mapper for IGA

The Nearest Neighbour (NN) mapper is a simple proximity-based transfer operator in which each destination integration point receives the value of the closest origin integration point. In the context of IGA-based coupling in Kratos, the local neighbour search is performed between integration points located on the coupling interface. However, since IGA solutions are stored at the level of control points, the final mapping operator must ultimately transfer the field between control-point degrees of freedom.

Let $\{x_i^d\}$ with $i = 1, \dots, n_d$ denote the set of destination integration points, and $\{x_j^o\}$ with $j = 1, \dots, n_o$ the set of origin integration points. For each destination point i , the index of the closest origin point is defined as

$$\text{closest}(i) = \arg \min_{j \in \{1, \dots, n_o\}} \|x_i^d - x_j^o\|.$$

The nearest neighbour transfer at the integration-point level is then given by

$$\tilde{u}_i^d = \tilde{u}^o(\text{closest}(i)), i = 1, \dots, n_d.$$

This operation can be written in matrix form as

$$\tilde{u}^d = H_{IP} \tilde{u}^o,$$

where the integration-point mapping matrix H_{IP} is defined component-wise as

$$(H_{IP})_{ij} = \begin{cases} 1, & \text{if } j = \text{closest}(i) \\ 0, & \text{otherwise} \end{cases}$$

In isogeometric analysis, values at integration points are obtained from control-point values

through the shape functions. This relationship can be expressed as

$$\tilde{u}^o = N^o u^o, \tilde{u}^d = N^d u^d.$$

Combining the nearest neighbour mapping with the IGA interpolation yields

$$N^d u^d = H_{IP} N^o u^o.$$

Since N^d is generally not square, the final control-point mapping

$$u^d = H u^o$$

is obtained through a weighted least-squares projection. This formulation shows that, although the proximity search is performed between integration points, the resulting operator transfers quantities consistently at the level of IGA control-point values.




Figure 6: Nearest Neighbour mapping at integration-point level: each destination integration point copies the value of the closest origin integration point (Euclidean distance)

4.2 Nearest Element Mapper for IGA

For IGA-based coupling, the Nearest Element mapper transfers interface quantities by exploiting the geometric representation of the origin IGA patch. In this approach, each destination integration point is mapped to the origin side through a geometric projection procedure; for this reason, the method is also commonly referred to as the closest projection mapper.

Given a destination integration point x_i^d defined in physical space, the algorithm first identifies a set of candidate patches on the origin interface. The destination point is then projected onto the geometric space of the closest origin IGA patch. As part of this procedure, the corresponding parametric coordinates ξ_i^o on the origin patch are computed by solving the inverse mapping from physical space to parametric space.

Once the parametric position ξ_i^o is known, the mapped value at the destination integration point is obtained by evaluating the origin IGA field at this location using the origin shape functions. In practice, the mapped value is computed as

$$\tilde{u}_i^d = N^o(\xi_i^o) u^o,$$

where u^o denotes the vector of origin control-point values and $N^o(\xi_i^o)$ is the vector of origin shape functions evaluated at the projected parametric coordinates.

Consequently, the mapping operation can be written in linear form as

$$\tilde{u}^d = H u^o,$$

where H is the transfer matrix whose entries correspond to the origin shape-function coefficients evaluated at the projection points. Each destination integration-point value is therefore expressed as a weighted linear combination of the origin control-point degrees of freedom.

It is important to note that this mapper naturally provides mapped values at the destination integration points. Depending on the application, these values can either be applied directly at the destination interface or further transferred to the destination control points.

The main advantage of the Nearest Element mapper is that it yields a smooth interpolation of interface quantities while remaining relatively straightforward to assemble once the projection information is available. However, the robustness of the method strongly depends on the success of the projection and inverse mapping procedures. For complex geometries or challenging interface configurations, projections may fail or become ill-conditioned, requiring additional safeguards or fallback strategies to ensure stable and reliable data transfer.

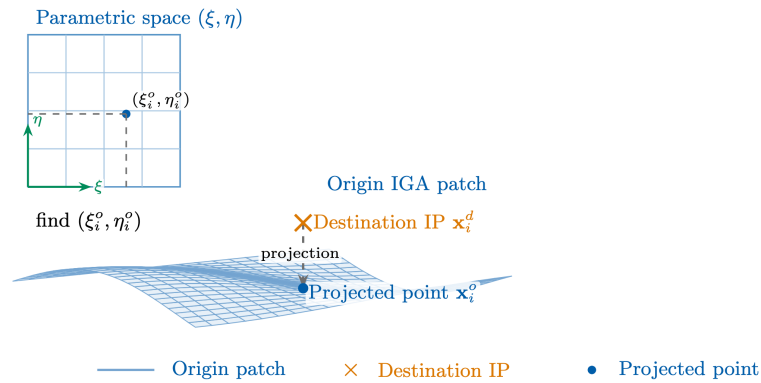


Figure 6: Nearest Element (closest projection) mapper for IGA, the destination integration point is projected onto the origin IGA patch

4.3 Mortar Mapper for Body-Fitted and Unfitted Discretizations

One of the most used data transfer operators in CoSimulation approaches with IGA-FEM is the "Mortar Mapper", which has been the primary focus of our work.

Let $\phi_o(x)$ be a known origin field and $\phi_D(x)$ the destination field. The Mortar Mapper can be derived from the minimization of the following functional $\varphi(\phi_D(x))$ (minimization of the L2-Norm Error between the origin and destination field):

$$\psi(\phi_D) = \frac{1}{2} \int_{\Gamma} (\phi_D(x) - \phi_o(x))^2 d\Gamma$$

As in any FEM-like approach, both fields are represented as a linear combination of nodal values

and their associated shape functions. To minimize the functional, its stationarity with respect to the nodal values of the destination field is imposed. This leads to a minimization problem that can be expressed in the following matrix form:

$$\frac{\partial \varphi(\phi_D)}{\partial \phi_{Dj}} = 0 \Rightarrow \int_{\Gamma_D} N_D N_D^T d\Gamma_D \Phi_D = \int_{\Gamma_{O \rightarrow D}} N_D N_O^T d\Gamma_{OD} \Phi_O$$

$$\mathbf{M}_{DD} \Phi_D = \mathbf{M}_{DO} \Phi_O \Rightarrow \text{Linear system to solve}$$

The primary challenge in the mortar mapper is determining the integration domain Γ_{OD} , which represents the overlapping region between a destination element and a neighboring origin element. The following subsections will outline the procedure for defining this integration domain across various discretization scenarios.

4.3.1 1D IBRA-IBRA or IBRA-FEM Interface

The process of determining the intersective integration domain in the 1D Isogeometric B-Rep Analysis (IBRA) mortar mapper consists of three main steps. First, the knots from the parameter space of the slave patch (destination) are projected onto the parameter space of the master patch (origin) to establish a geometric relationship between the two. Next, integration points are defined within the parameter space of the master patch, ensuring proper numerical integration over the overlap region. Finally, these integration points are mapped back from the master domain's parameter space to the parameter space of the slave domain, effectively defining the intersective integration domain needed for accurate coupling.




Figure 9: Definition of the integration domain for the 1D IBRA-IBRA interface (Ref. [5])

4.3.2 2D IBRA-FEM Interface

A particularly relevant scenario for 3D Fluid-Structure Interaction (FSI) simulations arises when the structural domain is discretized using a high-order NURBS multipatch surface, while the fluid domain is represented by a low-order finite element or finite volume mesh. In such cases, it is crucial to define an appropriate integration domain between the origin (fluid domain) and the destination (structural domain) for the mortar mapper. This process involves several steps to ensure accurate coupling between the two domains.




Figure 10: Low order FEM discretization (fluid) and CAD surface (structure) (Ref. [4])

First, the finite element nodes from the fluid domain are projected onto the parameter space of the NURBS patch, and the corresponding elements are reconstructed. For elements that are only partially projected into the patch, the finite element edges are clipped using the tessellated trimming curves that define the patch's boundary, with the help of some clipping algorithm. Next, the projected finite elements are clipped with the knot lines of the patch's parametric space, followed by a simple triangulation of the resulting geometry. Finally, integration points are placed in the master domain (fluid domain) and then projected back onto the destination finite elements (structural domain) to complete the coupling process. These steps ensure that the integration domain is well-defined for the mortar mapper, enabling an effective exchange of data between the fluid and structural solvers in FSI simulations.

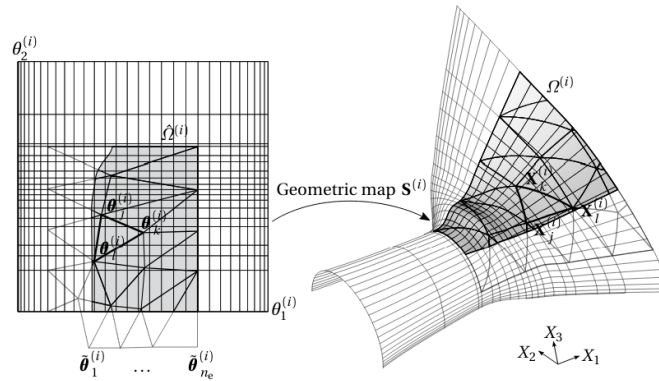


Figure 11: Projection of the low order element to the CAD surface (Ref. [4])

4.4 Radial Basis Function (RBF) Mapper

The Radial Basis Function (RBF) mapper is a smooth interpolation technique used to transfer data between two non-matching discretizations. In contrast to nearest-neighbour or closest-projection approaches, RBF interpolation is inherently more global: the value at each destination point is expressed as a weighted combination of contributions from all origin points. As a result, the method typically yields a smooth mapped field and improved accuracy for irregular or non-matching interfaces, at the expense of a higher setup cost and denser coupling operators.

Let $\{x_j^o\}$, $j = 1, \dots, n_o$, denote the set of origin points with known values u_j^o (e.g. displacements, pressures, or loads). The interpolated field $\hat{u}(x)$ is approximated as

$$\hat{u}(x) = \sum_{j=1}^{n_o} \lambda_j \phi(\|x - x_j^o\|) + p(x)^T \gamma,$$

where $\phi(r)$ is a radial basis function kernel depending only on the distance r , λ_j are the RBF interpolation weights, $p(x)$ is an optional polynomial enrichment basis (e.g. constant or linear terms), and γ contains the corresponding polynomial coefficients.

The unknown coefficients λ and γ are determined by enforcing the interpolation conditions at the origin points,

$$\hat{u}(x_i^o) = u_i^o \quad , \quad i = 1, \dots, n_o.$$

If polynomial enrichment is employed, the standard side constraint

$$P^T \lambda = 0$$

is added to ensure uniqueness of the solution, where P contains the polynomial basis evaluated at the origin point coordinates.

Defining the distance matrix A with entries

$$A_{ij} = \phi(\|x_i^o - x_j^o\|),$$

and the polynomial matrix P with entries $P_{il} = p_l(x_i^o)$, the interpolation coefficients follow from the saddle-point system

$$\begin{bmatrix} A & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} \lambda \\ \gamma \end{bmatrix} = \begin{bmatrix} u^o \\ 0 \end{bmatrix}.$$

Once λ and γ are known, the mapped values at the destination points $\{x_i^d\}$, $i = 1, \dots, n_d$, are obtained by evaluating the same interpolant. Introducing the matrix B with entries

$$B_{ij} = \phi(\|x_i^d - x_j^o\|),$$

and the polynomial matrix P_d at the destination points, with $(P_d)_{il} = p_l(x_i^d)$, the mapping reads

$$u^d = B\lambda + P_d\gamma.$$

Since the coefficients λ and γ depend linearly on the origin values u^o , the transfer can be written in compact form as

$$u^d = H u^o,$$

where H denotes the RBF transfer matrix. In practice, this operator is assembled once during the initialization phase and subsequently reused throughout the partitioned coupling iterations.

For IGA-based interfaces, the RBF interpolation is constructed using the integration points of the coupling surface. However, the coupling quantities are ultimately exchanged at the level of degrees of freedom, namely between IGA control points on one side and FEM nodes on the opposite side. Consequently, the final transfer operator is expressed as a linear mapping between control-point values and nodal values, while integration points are used only internally for assembling the interpolation.

5. Co-simulation Framework for Aeroelastic Analysis of an IBRA Wing (IBRA + VLM)

Aircraft development is inherently multidisciplinary and typically requires combining several analysis tools across different design stages. In early phases, uncertainty is high and computationally inexpensive models are preferred to enable rapid exploration of the design space, whereas in later stages more accurate simulations become necessary to assess structural integrity and performance in detail. This naturally motivates multi-fidelity co-simulation frameworks, where different physical models and fidelity levels can be combined in a modular and non-intrusive manner to balance accuracy and efficiency.

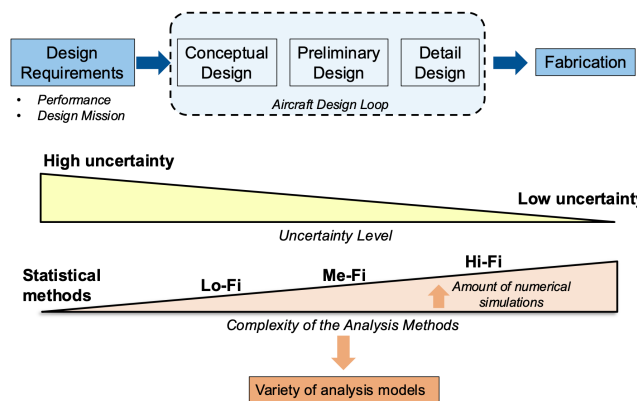


Figure 12: Aircraft design cycle

In this context, the present work considers a partitioned co-simulation workflow for aeroelastic analysis, coupling the structural response of an IBRA-based wing model with a low-fidelity aerodynamic solver based on a vortex-lattice (panel) formulation. The structure is resolved with a high-fidelity shell representation, while the aerodynamic loads are computed at low cost without requiring a full 3D CFD domain. The coupling is performed through a staggered exchange of interface quantities: structural deformations are transferred to the aerodynamic discretization to update the lifting surface geometry, and the resulting aerodynamic forces are mapped back to the structural interface. This strategy enables the prediction of aeroelastic effects such as static deflection and load redistribution under aerodynamic excitation while preserving the non-intrusiveness of both solvers.

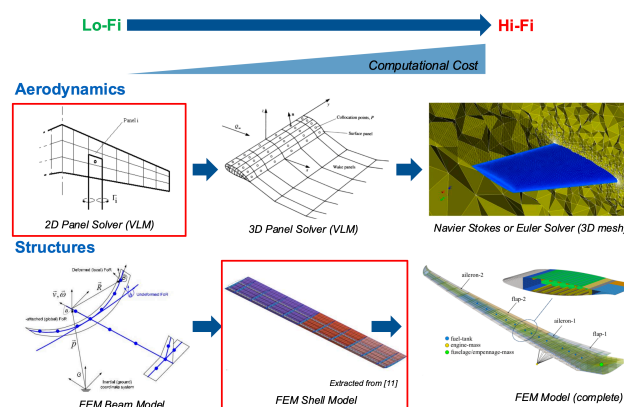


Figure 13: Multi-fidelity modelling spectrum for aerodynamics and structures

5.1 Aerodynamic Loads Calculation

The aerodynamic solver adopted in this work is based on a vortex-lattice (panel) formulation under the classical assumptions of inviscid, incompressible and irrotational flow. Under these conditions, a velocity potential ϕ exists and satisfies Laplace's equation:

$$\nabla^2 \phi = 0, \quad V = \nabla \phi$$

The lifting surface is discretized into N vortex panels of unknown strengths, collected in the vector:

$$\gamma = [\gamma_1 \ \gamma_2 \ \dots \ \gamma_N]^T$$




Figure 14: Wing discretization with VLM

The induced velocity at a spatial location x due to a vortex filament is computed using the Biot-Savart relation:

$$V(x) = (\gamma / 4\pi) \int_{\ell} \frac{(dl \times r)}{r^3}$$

To determine the unknown vortex strengths, the no-penetration boundary condition is enforced at each collocation point i , together with the Kutta condition. This yields:

$$V_{\infty} \cdot n_i + \sum_{j=1}^{\{N\}} AIC_{ij} \gamma_j = 0$$

Equivalently, the system can be written in matrix form as:

$$AIC \Gamma = b_c, \quad b_c = -V_{\infty} \cdot n$$

Once the circulation distribution is obtained, the pressure coefficient on each panel can be computed as:

$$\Delta C_{\{p,i\}} = \left(\frac{2}{C_i}\right) \cdot AIC_{\{ij\}}^{-1} \cdot b_{\{c,j\}} = AIC_{\{cp,ij\}} \cdot b_{\{c,j\}}$$

Finally, the aerodynamic force acting on the structure is assembled from the pressure distribution. In the implementation used here, the aerodynamic load vector can be expressed as:

$$f_A = q A AIC_{\{cp\}} (\alpha_{rigid} + \alpha_{elastic})$$

where q denotes the dynamic pressure, A the panel area matrix (or area scaling), and α_{rigid} and $\alpha_{elastic}$ represent the rigid and elastic contributions to the effective angle of attack.

5.2 Aeroelastic Co-simulation Framework

The aeroelastic co-simulation framework couples a structural IBRA solver with a vortex-lattice aerodynamic solver in a partitioned manner. The structural model provides the elastic deformation of the wing, while the aerodynamic solver computes the corresponding loads on the lifting surface. The coupling is based on a staggered exchange of interface quantities using consistent mapping operators.

On the aerodynamic side, the panel method computes the aerodynamic force vector as

$$f_A = q A AIC_{cp} (\alpha_{rigid} + \alpha_{elastic}),$$

where q is the dynamic pressure, A is the panel area scaling, and AIC_{cp} denotes the aerodynamic influence coefficient operator used to obtain pressure (or force) contributions from the effective angle of attack. The total angle of attack is decomposed into a rigid contribution α_{rigid} (e.g., prescribed flight condition) and an elastic contribution $\alpha_{elastic}$ induced by the structural deformation.

The elastic contribution is obtained from the structural solution by transferring the relevant deformation measure (here the sectional rotation around the y -axis) to the aerodynamic discretization:

$$\alpha_{elastic} = T_{AS\alpha} (\theta_y)_s$$

The operator $T_{AS\alpha}$ represents the mapping from structural degrees of freedom to the aerodynamic evaluation points, ensuring that the deformation-induced change in angle of attack is applied consistently.

After the aerodynamic forces are computed on the panel discretization, they are transferred back to the structural coupling interface as external loads:

$$F_s = T_{ASu}^T F_a.$$

Here, the transpose mapping is used to guarantee energetic consistency, meaning that the work transferred between both solvers is preserved across the interface. This two-way coupling allows the aerodynamic loads to update according to the current wing deformation, enabling an efficient aeroelastic co-simulation without requiring a high-fidelity 3D CFD model.

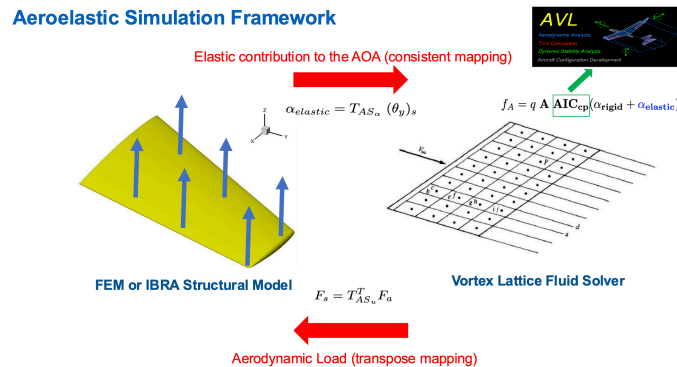


Figure 15: Aeroelastic co-simulation framework

To compute the elastic contribution to the effective angle of attack on the aerodynamic panels, the structural deformation must be transferred not only as displacements, but also as local rotations (or slopes). This is achieved using the same RBF interpolation framework employed for displacement mapping but additionally evaluating the spatial derivatives of the RBF interpolant at the aerodynamic interface points.

For example, using a thin-plate spline (TPS) RBF, the out-of-plane displacement field in the z-direction can be expressed as

$$u_z = w = \gamma^0 + \gamma^1 x + \gamma^2 y + \sum \lambda_i r_i^2 \ln(r_i^2),$$

where r_i denotes the distance between the evaluation point and the RBF support point i , and the coefficients $\{\beta_0, \beta_1, \beta_2, \gamma_i\}$ are obtained from the RBF interpolation system.

The corresponding rotation around the y-axis (slope in the x-direction), which represents the elastic contribution to the local angle of attack, is computed as the derivative of the displacement field:

$$\theta_y = \delta\alpha_{elastic} \approx \partial w / \partial x = \gamma_1 + 2 \sum \lambda_i (1 + \ln(r_i^2)) (x - x_i)$$

In this way, the aerodynamic solver receives a rotation (or angle-of-attack correction) that is consistent with the mapped displacement field, enabling a robust evaluation of aerodynamic loads on the deformed lifting surface.

5.3 Test Case Set-up

The aeroelastic co-simulation workflow is demonstrated on a representative wing test case generated with a parametric geometry definition. A trapezoidal wing planform is defined using a CPACS-based setup, where the main geometric parameters (span, area, aspect ratio, taper ratio and sweep angle) can be easily modified. The airfoil is selected as a symmetric NACA 4-digit profile, enabling a controlled benchmark configuration.

The structural model is represented as a shell-based wing box with uniform material properties. The selected parameters include the shell thickness, density, Young's modulus and Poisson's ratio, which define the stiffness and inertia characteristics of the wing. To represent the internal load-carrying structure, the model includes front and rear spars as well as three ribs located at the root, mid-span and tip. This provides a minimal but meaningful internal layout, allowing bending and torsional effects to develop under aerodynamic loading.

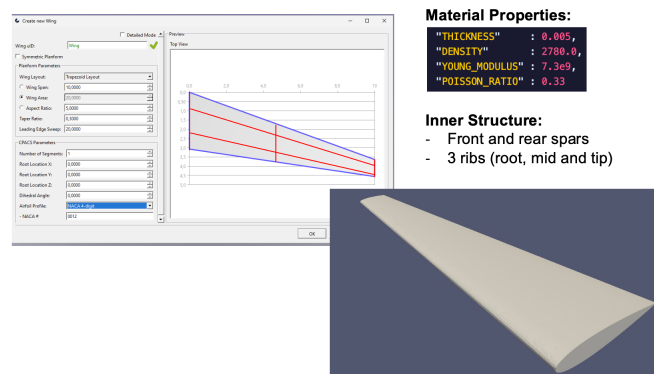


Figure 16: Wing geometry and material definition

5.4 Linear Aeroelastic Analysis

The minimal wing test case is coupled through a partitioned aeroelastic co-simulation workflow, where the structural response is solved in Kratos (IgaApplication), while the aerodynamic loads are evaluated with a vortex-lattice (panel) solver. The coupling is implemented in a non-intrusive manner using CoSimIO, such that both solvers can be executed independently while exchanging interface data at each coupling iteration.

First, the aerodynamic discretization (panel mesh) is exported to Kratos in order to define a compatible mapping interface. In practice, the aerodynamic lifting-surface points are created as a Kratos ModelPart, and the corresponding mesh information is sent through *CoSimIO.ExportMesh()*. This enables the mapping operators in Kratos to establish the transfer matrices between the aerodynamic and structural discretizations.

During the coupled simulation, the aerodynamic solver imports the elastic contribution to the angle of attack from Kratos, which is obtained from the current structural deformation. This is performed via *CoSimIO.ImportData()*, providing the aero solver with the updated elastic AOA values required to account for changes in the lifting surface orientation due to elastic deformation.

With the current flight condition and the imported elastic AOA, the aerodynamic solver then computes the panel forces, i.e. the updated aerodynamic load distribution acting on the lifting surface. These forces are subsequently exported back to Kratos using *CoSimIO.ExportData()*.

On the structural side, the aerodynamic loads are not applied on every structural node. Instead, they are transferred to a selected set of load application points, which are defined as locally stiff regions to ensure a robust and physically meaningful load introduction into the structural model. This avoids artificial local deformations caused by applying concentrated panel forces directly on a highly flexible shell discretization, and provides a stable coupling interface for the aeroelastic iterations.

This workflow results in a staggered coupling loop in which (i) structural deformation is transferred to the aerodynamic model as an elastic AOA update, and (ii) the resulting aerodynamic forces are transferred back to the structure through dedicated load application points, enabling the prediction of static aeroelastic effects such as wing deflection and load redistribution.

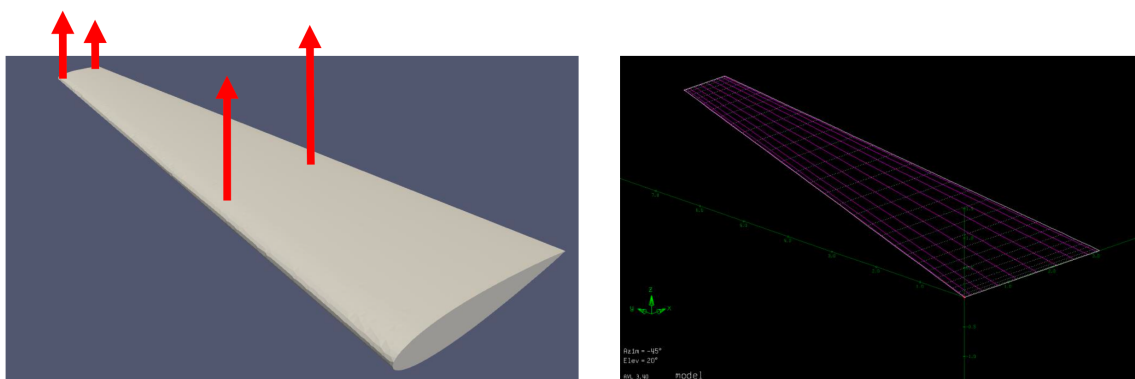


Figure 17: Structural load application points and corresponding aerodynamic panel discretization

5.5 Results: Linear Aeroelastic Response of the Minimal Wing

As a first step, the aeroelastic response of the minimal wing is compared using two independent finite element solvers: the Kratos StructuralMechanicsApplication and the Airbus in-house Lagrange solver, both using shell element. In both cases, the structural models are subjected to the same aerodynamic loads computed by the vortex-lattice (panel) solver and transferred to the structure through the coupling interface.

As shown in Fig. 18, the resulting displacement fields show very good agreement between the two solvers in terms of overall deformation shape and displacement levels. The spanwise bending behavior and the location of the maximum out-of-plane displacement near the wing tip are consistently captured. Small differences are mainly related to differences in mesh resolution and shell element formulations.




Figure 18: Comparison of FEM structural responses under aerodynamic loading, left is Kratos solution (26850 ShellThinCorotational3D3N elements) and right is Lagrange (18516 CQUAD4 and CTRIA3 elements)

This comparison validates the Kratos FEM solution under panel-based aerodynamic loading and serves as a reference for the subsequent structural analysis using IGA shell elements.

The aeroelastic response of the minimal wing structure was also computed using an Isogeometric Analysis (IGA) formulation based on an IBRA shell discretization. The structural model is represented by a Kirchhoff–Love shell formulation.

The IGA model is constructed as a multipatch NURBS surface representation, allowing a modular description of the wing topology. Interface continuity (C0 and C1) between patches is enforced using a penalty-based coupling approach, and the wing root support condition is also imposed via a penalty method. The aerodynamic loads applied to the IGA model originate from the vortex-lattice panel solver, mapped onto the structural interface using the previously defined transfer operators, ensuring consistency with the FEM-based reference analysis.

As seen in Fig. 19, the resulting displacement field exhibits the expected global bending and torsional deformation patterns, showing qualitative agreement with the FEM solutions obtained using both Kratos and the LAGRANGE in-house code.

In this example, the multipatch formulation required a careful fine-tuning of the penalty parameters to achieve numerical stability and accurate inter-patch continuity. This tuning process proved to be time-consuming and problem-dependent, highlighting a known limitation of penalty-based coupling strategies in complex multipatch configurations. While satisfactory results were obtained, this motivates the investigation of more robust and parameter-free coupling approaches, such as Nitsche-based coupling methods, which are expected to improve numerical robustness, reduce

user intervention, and enhance performance and scalability in future large-scale aeroelastic simulations.

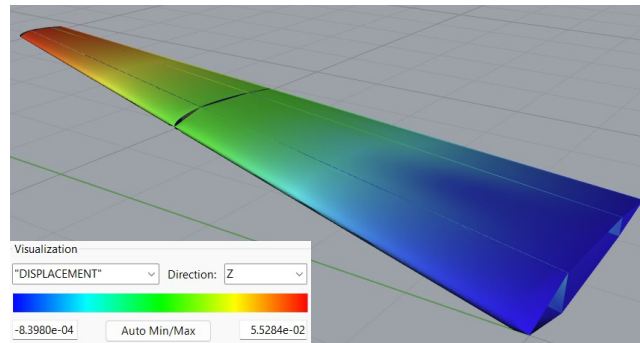


Figure 19: Structural Deformation of the Minimal Wing Using an IBRA Shell Model (9184 CPs)

5.6 Results: Aeroelastic Lift Coefficient Comparison

To assess the accuracy of the proposed aeroelastic coupling framework, the rigid and elastic lift coefficient obtained with different structural discretizations is compared for the minimal wing test case. The comparison includes a Lagrange-based reference solution, a Kratos FEM shell model, and a Kratos IBRA shell model.

In all cases, the aerodynamic loads are computed using the same vortex-lattice (panel) solver, and the same flight conditions are applied. This ensures that any differences in the resulting lift coefficients are exclusively attributable to the structural formulation and the mapping strategy employed in the coupled analysis.

	Lagrange	Kratos FEM	Kratos IBRA
CL rigid	0.4270073	0.427088	0.427088
CL elastic	0.4308836	0.43011	0.430351
Delta CL elastic [%]	-	0.18%	0.086%

Figure 20: Lift Coefficient Comparison

For the rigid configuration, all three approaches yield virtually identical lift coefficients, confirming the consistency of the aerodynamic model and load transfer. When elastic effects are included, a small increase in lift is observed due to the elastic contribution to the effective angle of attack, which is consistently transferred from the structural solver to the aerodynamic solver within the co-simulation framework.

The elastic lift increment remains below 0.2% for both Kratos FEM and Kratos IBRA, with the IBRA formulation exhibiting a slightly smaller deviation from the reference solution. This level of agreement demonstrates that the IBRA-based structural model is capable of accurately capturing the aeroelastic response while delivering results fully comparable to classical FEM formulations. Overall, these results validate the correctness of the coupling strategy and establish the FEM-based solution as a reliable reference for the subsequent aeroelastic analyses using IGA shell elements.

6. Partitioned Co-simulation Framework for Vibroacoustic Systems

Vibroacoustic problems involve the coupled interaction between a vibrating structure and the surrounding acoustic medium. To efficiently address this multi-physics interaction while preserving solver modularity, a partitioned co-simulation approach is adopted. Within this framework, the structural subproblem is solved using the Kratos *StructuralMechanicsApplication*, while the acoustic field governed by the Helmholtz equation is resolved using a MATLAB-based boundary element solver. Both subproblems are treated independently and coupled through the exchange of interface quantities in the frequency domain, enabling a flexible and non-intrusive vibroacoustic simulation strategy.

6.1 Governing Equations

Structural Subproblem

The structural dynamics of the elastic body are governed by the semi-discrete equations of motion

$$\mathbf{M}\ddot{\mathbf{u}}(t) + \mathbf{C}\dot{\mathbf{u}}(t) + \mathbf{K}\mathbf{u}(t) = \mathbf{f}_m(t) + \mathbf{f}_a(t),$$

where \mathbf{M} , \mathbf{C} , and \mathbf{K} denote the mass, damping, and stiffness matrices, respectively, $\mathbf{u}(t)$ collects the structural degrees of freedom, $\mathbf{f}_m(t)$ represents mechanical loads, and $\mathbf{f}_a(t)$ corresponds to the acoustic pressure-induced forces acting on the vibroacoustic interface.

Assuming time-harmonic excitation at angular frequency ω , the structural response is expressed in the frequency domain as

$$(-\omega^2\mathbf{M} + i\omega\mathbf{C} + \mathbf{K})\hat{\mathbf{u}}(\omega) = \hat{\mathbf{f}}_m(\omega) + \hat{\mathbf{f}}_a(\omega).$$

To enable efficient coupling with the acoustic solver, a modal reduction strategy is employed. The structural displacement field is approximated using a truncated expansion of mass-normalized eigenmodes,

$$\hat{\mathbf{u}}(\omega) \approx \sum_{n=1}^{N_m} \boldsymbol{\phi}_n \hat{q}_n(\omega),$$

where N_m denotes the number of retained modes and $\hat{q}_n(\omega)$ are the modal amplitudes. This reduced-order representation significantly decreases the computational cost while retaining the dominant dynamic behavior in the frequency range of interest.

Acoustic Subproblem

The acoustic field is assumed to be linear, inviscid, and time-harmonic, and is governed by the Helmholtz equation in the acoustic domain,

$$\nabla^2 \hat{p}_a(\mathbf{r}, \omega) + k^2 \hat{p}_a(\mathbf{r}, \omega) = -i\rho_a \omega \hat{q}_a(\mathbf{r}, \omega),$$

where \hat{p}_a denotes the complex acoustic pressure, $k = \omega/c$ is the acoustic wavenumber, c is the

speed of sound, and ρ_a is the fluid density.

To efficiently treat unbounded domains, the acoustic problem is solved using a boundary element method (BEM) formulation. The structural motion enters the acoustic problem through Neumann boundary conditions prescribed on the vibroacoustic interface, where the normal fluid particle velocity is imposed by the structural velocity,

$$\hat{v}_n = i\omega \hat{u}_n \text{ on } \Gamma_{sa}.$$

Once the acoustic problem is solved, the resulting pressure field is evaluated on the interface and converted into equivalent nodal forces acting on the structure.

Vibroacoustic Coupling Conditions

Let Ω_s be the structural domain and Ω_a the acoustic domain, sharing the interface $\Gamma = \partial\Omega_s \cap \partial\Omega_a$. Denote by \mathbf{n} the unit normal pointing from the structure into the acoustic domain.

1) Kinematic condition (velocity continuity)

The normal component of the acoustic particle velocity equals the structural normal velocity on Γ :

$$v_n^a(\mathbf{x}, \omega) = v_n^s(\mathbf{x}, \omega) \quad \forall \mathbf{x} \in \Gamma.$$

With time-harmonic convention $e^{i\omega t}$, the structural normal velocity is

$$v_n^s(\mathbf{x}, \omega) = i\omega \mathbf{u}(\mathbf{x}, \omega) \cdot \mathbf{n}.$$

For a pressure-based Helmholtz formulation, the corresponding Neumann boundary condition is

$$\frac{\partial p}{\partial n}(\mathbf{x}, \omega) = -i\omega \rho_a v_n^s(\mathbf{x}, \omega) = \omega^2 \rho_a \mathbf{u}(\mathbf{x}, \omega) \cdot \mathbf{n}, \quad \forall \mathbf{x} \in \Gamma,$$

where ρ_a is the acoustic medium density.

2) Dynamic condition (traction equilibrium)

The structure is loaded by the acoustic pressure acting as a normal traction:

$$\boldsymbol{\sigma}(\mathbf{u}) \mathbf{n} = -p \mathbf{n} \text{ on } \Gamma,$$

i.e. the external interface traction is

$$\mathbf{t}_\Gamma(\mathbf{x}, \omega) = -p(\mathbf{x}, \omega) \mathbf{n}.$$

In discrete form, the equivalent structural load vector is

$$\mathbf{f}_\Gamma(\omega) = \int_\Gamma \mathbf{N}^\top(\mathbf{x}) \mathbf{t}_\Gamma(\mathbf{x}, \omega) d\Gamma = - \int_\Gamma \mathbf{N}^\top(\mathbf{x}) p(\mathbf{x}, \omega) \mathbf{n} d\Gamma.$$

6.2 Monolithic and Partitioned Approaches for Vibroacoustic Systems

Monolithic Vibroacoustic Formulation

In a monolithic approach, the structural and acoustic governing equations are assembled into a single coupled system. The structural displacements and the acoustic pressure are treated as unknowns of one global problem, and the coupling conditions at the fluid–structure interface are enforced implicitly within the system matrix.

After spatial discretization and assuming time-harmonic behavior, the coupled system can be written as:

$$\begin{bmatrix} \mathbf{K}_s - \omega^2 \mathbf{M}_s & \mathbf{C}_{sa} \\ \mathbf{C}_{as} & \mathbf{K}_a - \omega^2 \mathbf{M}_a \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_s \\ \mathbf{f}_a \end{bmatrix},$$

where \mathbf{u} denotes the structural displacement vector, \mathbf{p} the acoustic pressure vector, and \mathbf{C}_{sa} , \mathbf{C}_{as} represent the coupling operators enforcing velocity continuity and pressure loading at the interface.

The main advantage of the monolithic approach is its numerical robustness and unconditional stability, since the coupling is resolved implicitly. However, it requires a single solver framework, leads to large coupled systems, and significantly reduces flexibility. As a result, it is often impractical when different specialized solvers or software environments are used for the structural and acoustic domains.

Partitioned Vibroacoustic Formulation

In contrast, the partitioned approach solves the structural and acoustic subproblems independently, using dedicated solvers, and enforces the coupling through an exchange of interface quantities. The structure and the acoustic field remain modular, and the interaction is handled through boundary conditions at the vibroacoustic interface.

In operator form, the partitioned coupling can be expressed as:

$$p = \mathcal{A}(v_n), \quad v_n = \mathcal{S}(p),$$

where \mathcal{A} and \mathcal{S} denote the acoustic and structural solution operators, respectively.

The partitioned approach can be implemented using:

- Weak (explicit) coupling, where the interface quantities are exchanged once per frequency.
- Strong (iterative) coupling, where the solvers are iterated until convergence of the interface residual.

The main advantages of the partitioned strategy are solver modularity, non-intrusive coupling, and scalability, which make it particularly suitable for multi-physics and multi-fidelity workflows. The main drawback is the potential loss of stability for strongly coupled problems, which may require iterative coupling and relaxation techniques.

6.3 Co-simulation Strategies for Partitioned Vibroacoustic Systems

6.3.1 Coupling Methods

Let \mathcal{S} denote the structural solution operator that maps acoustic pressure acting on the interface to the resulting structural normal velocity, and let \mathcal{A} denote the acoustic solution operator that

maps a prescribed structural normal velocity to the resulting acoustic pressure on the interface. For a given angular frequency ω , the coupled vibroacoustic problem can be written as the fixed-point system

$$v_n = \mathcal{S}(p), p = \mathcal{A}(v_n),$$

or equivalently,

$$v_n = \mathcal{S}(\mathcal{A}(v_n)).$$

1) Weak (Explicit) Coupling

In a weakly coupled (explicit) partitioned approach, the interface quantities are exchanged once per frequency without iteration. Given an initial estimate of the structural state (e.g. a purely mechanical solution or zero pressure loading), the structural normal velocity is computed and passed to the acoustic solver:

$$p^{(0)} = \mathcal{A}(v_n^{(0)}).$$

The resulting acoustic pressure is then applied to the structure to compute the updated structural response:

$$v_n^{(1)} = \mathcal{S}(p^{(0)}).$$

The procedure stops after this single exchange. While this approach is computationally inexpensive and straightforward to implement, it neglects the feedback loop between the structure and the acoustic field. Therefore, weak coupling is only suitable for problems with weak fluid–structure interaction, where the acoustic loading has a limited influence on the structural response.

2) Strong (Iterative) Coupling

In a strongly coupled (iterative) partitioned approach, the fixed-point problem is solved iteratively until convergence of the interface quantities is achieved. Starting from an initial guess $v_n^{(0)}$, the following iterations are performed:

$$\begin{aligned} p^{(k)} &= \mathcal{A}(v_n^{(k)}), \\ v_n^{(k+1)} &= \mathcal{S}(p^{(k)}), \end{aligned}$$

for $k = 0, 1, 2, \dots$, until the interface residual satisfies a prescribed tolerance. A common convergence criterion is based on the relative change of the normal velocity:

$$\frac{\|v_n^{(k+1)} - v_n^{(k)}\|}{\|v_n^{(k)}\|} < \varepsilon.$$

Strong coupling accounts for the mutual interaction between structure and acoustic field and ensures consistency of the interface conditions. Although it requires additional solver calls and increased computational effort, it provides improved stability and accuracy, especially for problems with strong vibroacoustic coupling.

6.3.2 Data Exchange between the Solvers

In the partitioned vibroacoustic framework, the structural and acoustic subproblems are solved by independent solvers running as separate processes. As a consequence, an explicit interprocess communication (IPC) mechanism is required to exchange interface quantities during the co-simulation. In this work, data transfer is handled using CoSimIO, which is integrated into the Kratos CoSimulationApplication and provides a lightweight, solver-agnostic communication layer. CoSimIO supports different IPC backends.

In the present implementation, a file-based IPC strategy is employed for the communication between the structural solver in Kratos and the external MATLAB BEM acoustic solver. In this mode, interface data are exchanged through files written to a dedicated communication directory. Synchronization is ensured by auxiliary control files, guaranteeing that data are accessed only after being fully written by the sending solver. This approach avoids race conditions and allows the solvers to remain loosely coupled in time.




Figure 21: Data exchange architecture in the partitioned vibroacoustic co-simulation

From a practical standpoint, file-based IPC offers several advantages for this application. It is robust, easy to set up, and well suited for coupling solvers implemented in different languages and execution environments. Moreover, it allows straightforward inspection of the exchanged data, which is particularly useful during development and debugging. While this communication method introduces a higher latency compared to shared-memory or socket-based approaches, its overhead is acceptable for the frequency-domain vibroacoustic analyses considered here.

During each coupling step, the structural solver exports the relevant interface quantities (e.g., displacements at the vibroacoustic interface) using CoSimIO. The acoustic solver then imports this data, solves the Helmholtz problem, and returns the resulting acoustic loads, which are subsequently applied as external loads in the structural analysis. In this way, CoSimIO enables a clean and modular data exchange without requiring intrusive modifications of either solver.

6.3.3 Data Mapping

In partitioned vibroacoustic simulations, the structural and acoustic solvers typically employ different discretizations at the coupling interface. As a result, a data mapping step is required to consistently transfer interface quantities between the non-matching meshes used by the two solvers. This mapping operation constitutes a key component of the co-simulation workflow, as it directly affects the accuracy of the coupled solution.

In the present work, data transfer across the vibroacoustic interface is performed using mapping operators provided by the Kratos MappingApplication. For the vibroacoustic test cases considered, two mapping strategies are employed: the nearest neighbour mapper and the nearest element mapper. Both approaches have been introduced and discussed in detail in the previous sections of this report.

6.4 Numerical Results

This section presents a set of numerical examples designed to evaluate the proposed partitioned vibroacoustic co-simulation framework across increasing levels of complexity. The test cases are deliberately organized in a progressive manner, allowing the performance, robustness, and consistency of the framework to be assessed systematically. Three representative benchmark problems are considered.

For all examples, the acoustic response is evaluated and presented in terms of the Sound Pressure Level (SPL), which provides a standard logarithmic measure of the radiated acoustic field. The SPL is defined as

$$\text{SPL} = 20 \log_{10} \left(\frac{p_{\text{rms}}}{p_{\text{ref}}} \right) \quad [\text{dB}],$$

where p_{rms} denotes the root-mean-square acoustic pressure and $p_{\text{ref}} = 2 \times 10^{-5}$ Pa is the reference pressure in air. Expressing the results in terms of SPL enables a direct assessment of how structural vibration characteristics, and in particular resonance phenomena, influence the acoustic field over the investigated frequency range.

The first example addresses a one-way coupled vibroacoustic problem. A simply supported rectangular plate is subjected to a prescribed surface load, and the resulting structural vibration is transferred to the acoustic solver. No acoustic source is present, and no feedback from the acoustic domain to the structure is considered. This configuration serves as a baseline case for validating the structural-acoustic data transfer and for analyzing the impact of structural resonances on the resulting sound pressure levels.

The second example extends the previous configuration to a two-way coupled problem, where the same simply supported rectangular plate is excited by a single acoustic monopole. In this case, the acoustic pressure field induces structural vibrations, which in turn modify the acoustic response. This setup introduces bidirectional coupling and increased numerical complexity, allowing the stability and consistency of the partitioned coupling strategy to be assessed.

Finally, a simply supported hemispherical structure subjected to an incident plane acoustic wave is considered. This example combines a curved structural interface with a distributed acoustic excitation in a fully three-dimensional setting. It represents the most challenging configuration and is used to evaluate the robustness of the framework with respect to geometric complexity and strongly coupled vibroacoustic interactions.

6.4.1 Example 1: One-Way Coupled Vibroacoustic Analysis of a Simply Supported Rectangular Plate subject to a Surface Load

The first numerical example considers a one-way coupled vibroacoustic problem involving a simply supported rectangular plate subjected to a uniform surface load. The plate has length $L = 1$ and width $w = 0.5$, and is simply supported along all edges, as illustrated in Fig. 22. A constant harmonic surface load of unit amplitude $q = 1$ is applied normal to the plate surface. The structure is modeled as a thin, homogeneous, isotropic elastic shell with thickness $t = 5 \times 10^{-3} \text{ m}$, density $\rho = 2780 \text{ kg/m}^3$, Young's modulus $E = 7.3 \times 10^{10} \text{ Pa}$, and Poisson's ratio $\nu = 0.33$.




Figure 22: Rectangular plate simply supported along all four edges and subjected to a uniform surface load

In this configuration, only the structural-to-acoustic coupling is considered. The structural response is computed in the frequency domain, and the resulting displacements at the plate surface are transferred to the acoustic solver as boundary conditions. No acoustic source is present, and feedback from the acoustic domain to the structure is neglected. This setup allows the isolated effect of structural resonances on the acoustic field to be analyzed.

The radiated acoustic response is evaluated by computing the sound pressure level (SPL) at a fixed field point located at

$$\mathbf{x}_{\text{obs}} = [0.5, 0.25, 1],$$

measured from the plate reference frame. The SPL is evaluated over a prescribed frequency range in order to identify peaks associated with the structural resonance frequencies.

The following image compares the monolithic and partitioned solutions. The results show very good agreement across the investigated frequency range, demonstrating that the partitioned approach reliably captures the structural resonance effects on the acoustic sound pressure level.

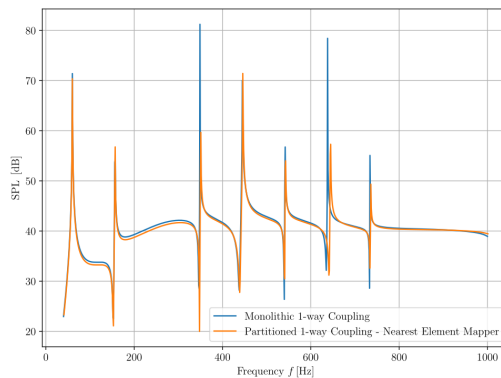


Figure 23: SPL Frequency Sweep: Monolithic vs. Partitioned One-Way Coupling

The first 20 structural natural frequencies of the simply supported rectangular plate are listed in Fig. 24. By analyzing the SPL frequency sweep, it can be clearly observed that pronounced peaks occur at frequencies that closely match these structural eigenfrequencies. This indicates that the applied surface load efficiently excites specific structural modes, which in turn dominate the radiated acoustic response. In particular, modes whose modal shapes are compatible with the spatial distribution of the surface load lead to strong structural vibrations and a corresponding amplification of the sound pressure level at the observation point. This confirms that the acoustic field is primarily driven by structurally induced resonances in the considered one-way coupled configuration.

$f_1 = 60.18 \text{ Hz}$	$f_{11} = 482.36 \text{ Hz}$
$f_2 = 96.31 \text{ Hz}$	$f_{12} = 484.96 \text{ Hz}$
$f_3 = 156.62 \text{ Hz}$	$f_{13} = 494.73 \text{ Hz}$
$f_4 = 204.77 \text{ Hz}$	$f_{14} = 542.24 \text{ Hz}$
$f_5 = 240.86 \text{ Hz}$	$f_{15} = 626.73 \text{ Hz}$
$f_6 = 241.31 \text{ Hz}$	$f_{16} = 628.56 \text{ Hz}$
$f_7 = 301.13 \text{ Hz}$	$f_{17} = 644.41 \text{ Hz}$
$f_8 = 350.58 \text{ Hz}$	$f_{18} = 735.67 \text{ Hz}$
$f_9 = 385.67 \text{ Hz}$	$f_{19} = 785.95 \text{ Hz}$
$f_{10} = 446.31 \text{ Hz}$	$f_{20} = 787.73 \text{ Hz}$

Figure 24: First 20 natural frequencies of the considered structure

6.4.2 Example 2: Two-Way Vibroacoustic Coupling - Rectangular Plate with Surface Load and Monopole Excitation

The second numerical example considers the same simply supported rectangular plate introduced in Example 1 (dimensions $L = 1$, $w = 0.5$, simply supported along all four edges). In contrast to the previous one-way configuration, this case introduces an acoustic excitation and a two-way vibroacoustic coupling, such that both solvers influence each other through the interface conditions.

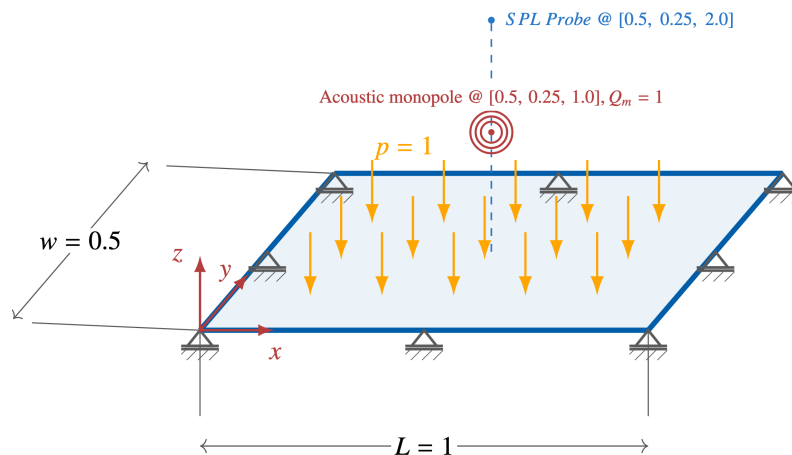


Figure 25: Rectangular plate simply supported along all four edges and subjected to a uniform surface load and an acoustic monopole excitation

The plate is subjected to a harmonic uniform surface load of unit amplitude $p = 1$, acting normal to the plate surface. In addition, the acoustic domain is excited by a harmonic monopole source

located at

$$\mathbf{x}_m = [0.5, 0.25, 1.0],$$

i.e., aligned with the plate center and positioned above the structure. The monopole generates an incident acoustic pressure field that loads the plate and induces structural vibrations. The resulting structural motion is transferred back to the acoustic solver through the normal-velocity boundary condition, thereby modifying the acoustic field and leading to a fully coupled vibroacoustic interaction.

As in the first example, the acoustic response is quantified in terms of the Sound Pressure Level (SPL), evaluated at the fixed observation point

$$\mathbf{x}_{\text{obs}} = [0.5, 0.25, 2.0].$$

A frequency sweep is performed to assess the combined effect of direct acoustic excitation and structurally radiated sound, and to identify the influence of structural resonances under bidirectional coupling.

Figure 26 shows the SPL frequency response obtained with weak (explicit) and strong (iterative) partitioned coupling strategies. At low frequencies, the response is dominated by the direct contribution of the acoustic monopole. Superimposed on this baseline, pronounced SPL peaks appear at frequencies corresponding to the structural resonance modes of the plate, indicating a strong interaction between the structural and acoustic subsystems.

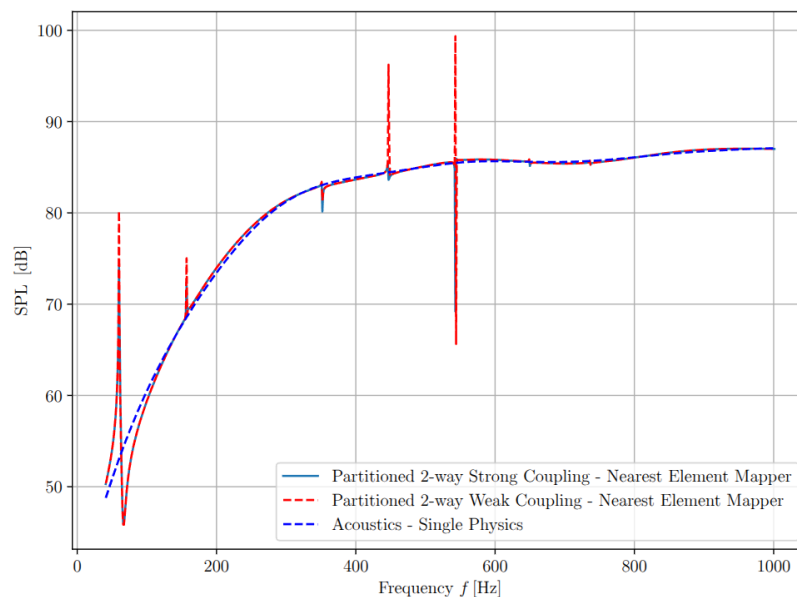


Figure 26: SPL Frequency Sweep - comparison between partitioned two-way strong coupling, partitioned two-way weak coupling, and single-physics acoustic solution.

Differences between weak and strong coupling are most evident in the vicinity of resonance frequencies. The weakly coupled solution tends to overpredict the SPL peaks, whereas the strongly coupled approach yields a smoother and more stable response due to the iterative enforcement of interface equilibrium.

6.4.3 Example 3: Simply Supported Hemispherical Shell Subjected to an Incident Plane Wave and Surface Pressure Load

The final numerical example considers a fully coupled vibroacoustic problem involving a simply supported hemispherical shell subjected to both a harmonic acoustic plane wave and a uniform structural pressure load. This case represents the most challenging benchmark in the present study, combining geometric curvature, distributed acoustic excitation, and strong structure–acoustic interaction.

The hemispherical shell has radius $R = 0.25 \text{ m}$ and is simply supported along its circular base, as illustrated in Fig. 27. The structure is modeled using IBRA Kirchhoff–Love shell elements, enabling a smooth CAD-integrated representation of the curved geometry. The shell has a uniform thickness $t = 0.01 \text{ m}$ and is assumed to be made of a linear elastic material with Young’s modulus $E = 1 \times 10^9 \text{ Pa}$, Poisson’s ratio $\nu = 0.45$, and mass density $\rho = 1200 \text{ kg/m}^3$.

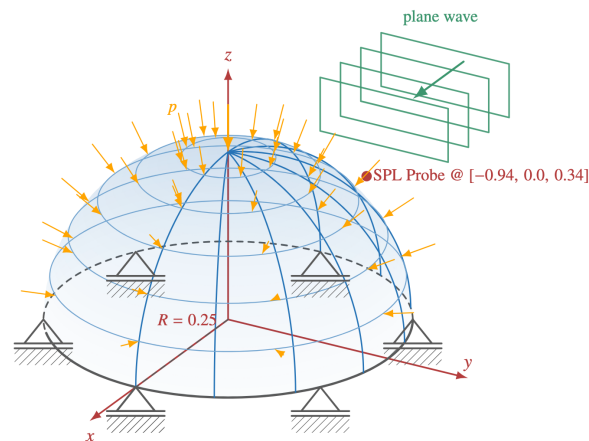


Figure 27: Simply supported hemisphere subjected to a pressure load and an acoustic plane wave excitation

In addition to the acoustic excitation, a pressure load of prescribed amplitude $p = 100 \text{ Pa}$ is applied normal to the shell surface, providing a direct structural excitation mechanism. The acoustic domain is excited by a harmonic plane wave propagating in the positive x -direction, which induces a spatially varying acoustic pressure field on the shell. Both the acoustic pressure and the structural surface load contribute to the shell vibration.

The resulting structural motion is transferred back to the acoustic solver through the normal-velocity boundary condition, leading to a fully bidirectional vibroacoustic coupling.

The acoustic response is evaluated in terms of the sound pressure level (SPL) at a fixed observation point located at

$$\mathbf{x}_{\text{obs}} = [-0.94, 0.0, 0.34].$$

Figure 28 presents the SPL frequency response obtained from the partitioned two-way coupling approach, comparing weak (explicit) and strong (iterative) coupling strategies, together with a purely acoustic reference solution. At low frequencies, the response is dominated by the incident plane wave and the applied surface load, resulting in similar SPL levels across all approaches. At higher frequencies, pronounced SPL peaks emerge due to structural resonance modes of the hemispherical shell, indicating strong vibroacoustic interaction.

Differences between weak and strong coupling become more pronounced in the vicinity of these resonances. The weakly coupled solution exhibits sharper and higher SPL peaks, whereas the strongly coupled solution provides a smoother and more stable response due to the iterative enforcement of interface consistency. This highlights the importance of strong coupling strategies for curved structures subjected to combined acoustic and structural excitation.

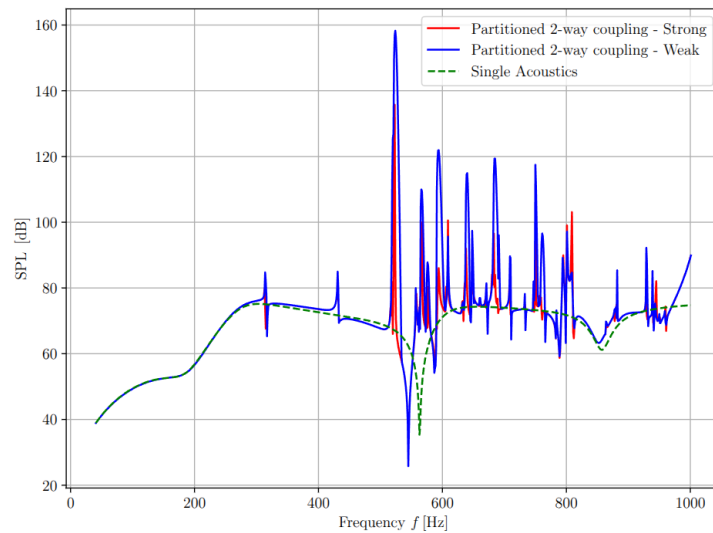


Figure 28: SPL Frequency Sweep - comparison between partitioned two-way strong coupling, partitioned two-way weak coupling, and single-physics acoustic solution.

8. Conclusions

This report presented the implementation and application of partitioned co-simulation workflows for aeroelastic and vibroacoustic problems within the Kratos Multiphysics framework. The work focused on the practical realization of solver coupling, data exchange, and interface mapping, with particular attention to reproducibility and modularity of the co-simulation setup.

A central aspect of the presented developments is the use of CAD-integrated structural analysis through the IBRA framework, employing Kirchhoff–Love shell formulations for both aeroelastic and vibroacoustic simulations. The use of isogeometric B-Rep representations enabled smooth structural kinematics and consistent geometric descriptions across the coupled solvers, while also introducing specific challenges related to multi-patch coupling.

In the aeroelastic part, IBRA- and FEM-based structural models were coupled with a low-fidelity Vortex Lattice Method aerodynamic solver using an RBF-based mapping strategy. Particular attention was given to the transfer of rotational information and to the treatment of multi-patch structural models, where penalty-based coupling was employed. While this approach allowed the simulations to be carried out successfully, the results highlighted the sensitivity of penalty formulations with respect to parameter tuning, motivating the investigation of Nitsche-type coupling methods as a more robust alternative in future work.

For the vibroacoustic applications, a frequency-domain partitioned framework was implemented by coupling harmonic structural response analysis in Kratos with an external Boundary Element Method acoustic solver implemented in MATLAB. Nearest-neighbor and nearest-element mapping techniques were used to exchange interface quantities between the structural and acoustic domains. The numerical examples demonstrated that the co-simulation framework captures the expected resonance-driven acoustic behavior for increasing levels of problem complexity.

Overall, the developments presented in this report contribute to the establishment of structured and reproducible partitioned co-simulation workflows, combining CAD-integrated structural models with external aerodynamic and acoustic solvers within Kratos Multiphysics. These workflows provide a consistent basis for further investigations of coupled aeroelastic and vibroacoustic problems using heterogeneous solvers and discretizations.

9. REFERENCES

- [1] Bucher, P. "CoSimulation and Mapping for Large Scale Problems". Dissertation. Technical University of Munich, 2023.
- [2] T. Wang. "Development of Co-Simulation Environment and Mapping Algorithms". Dissertation. Technical University of Munich, 2016.
- [3] Wilson, P., Teschemacher, T., Bucher, P., & Wüchner, R. (2021). Non-conforming FEM-FEM coupling approaches and their application to dynamic structural analysis. *Engineering Structures*, 241, 112342.
- [4] Apostolatos, A., Emiroğlu, A., Shayegan, S., Péan, F., Bletzinger, K. U., & Wüchner, R. (2021). An isogeometric b-rep mortar-based mapping method for non-matching grids in fluid-structure interaction. *Advanced Modeling and Simulation in Engineering Sciences*, 8(1), 9.
- [5] Bauer, A. M. CAD-integrated isogeometric analysis and design of lightweight structures. Doctoral dissertation. Technische Universität München, 2020.
- [6] Dadvand, P., Rossi, R., & Oñate, E. (2010). An object-oriented environment for developing finite element codes for multi-disciplinary applications. *Archives of computational methods in engineering*, 17, 253-297.
- [7] Teschemacher, T., Bauer, A. M., Oberbichler, T., Breitenberger, M., Rossi, R., Wüchner, R., & Bletzinger, K. U. (2018). Realization of CAD-integrated shell simulation based on isogeometric B-Rep analysis. *Advanced Modeling and Simulation in Engineering Sciences*, 5, 1-54.
- [8] Mok, D. P. (2001). *Partitionierte Lösungsansätze in der Strukturdynamik und der Fluid-Struktur-Interaktion*.
- [9] Turek, S., Hron, J., Razzaq, M., Wobker, H., & Schäfer, M. (2010). Numerical benchmarking of fluid-structure interaction: A comparison of different discretization and solution approaches (pp. 413-424). Springer Berlin Heidelberg.
- [10] Valdés Vázquez, J. G. (2007). *Nonlinear Analysis of Orthotropic Membrane and Shell Structures Including Fluid-Structure Interaction*.