



A hybrid reduced-order and high-fidelity Discontinuous Galerkin Spectral Element framework for large-scale PMUT array simulations

Michelangelo G. Garroni¹ · Paola F. Antonietti¹ · Ilario Mazzieri¹ · Nicola Parolini¹ · Omer M. O. Abdalla²

Received: 18 March 2026 / Revised: 22 June 2026 / Accepted: 7 July 2026
© The Author(s) 2026

Abstract

Piezoelectric Micromachined Ultrasonic Transducers (PMUTs) are essential for next-generation ultrasonic sensing and imaging due to their bidirectional electromechanical behavior, compact design, and compatibility with low-voltage electronics. As PMUT arrays grow in size and complexity, efficiently modeling their coupled electromechanical-acoustic behavior becomes increasingly challenging. This work presents a novel computational framework that combines model order reduction with a Discontinuous Galerkin Spectral Element Method (DG-SEM) paradigm to simulate large PMUT arrays. Each PMUT's mechanical behavior is represented using a reduced set of vibration modes, which are coupled to an acoustic domain model to describe the full array. To further improve efficiency, a secondary acoustic domain is connected via DG interfaces, enabling non-conforming mesh refinement, with variable approximation order, and accurate wave propagation. The framework is implemented within the SPECTral Elements in Elastodynamics with Discontinuous Galerkin (SPEED) software, an open-source library. A key contribution of this work is the introduction of a novel parallel computational framework based on a domain decomposition procedure, high-order polynomial discretizations, METIS graph partitioning, and MPI communications. This new feature enables efficient large-scale wave propagation simulations involving very large PMUT arrays on distributed-memory architectures. The proposed methodology addresses key challenges in meshing, supporting high-fidelity simulations for both PMUT transmission and reception phases. Numerical results demonstrate the framework's accuracy, scalability, and efficiency for large PMUT array simulations.

Keywords PMUT · Model order reduction · Wave propagation · Parallel computing

Paola F. Antonietti, Ilario Mazzieri, Nicola Parolini, and Omer M.O. Abdalla contributed equally to this work.

✉ Michelangelo G. Garroni
michelangelogabriele.garroni@polimi.it

Paola F. Antonietti
paola.antonietti@polimi.it

Ilario Mazzieri
ilario.mazzieri@polimi.it

Nicola Parolini
nicola.parolini@polimi.it

Omer M. O. Abdalla
omermohamed.abdalla@polimi.it

¹ MOX-Laboratory for Modeling and Scientific Computing, Department of Mathematics, Politecnico di Milano, Piazza Leondardo da Vinci 32, 20133 Milan, Italy

² Department of Civil and Environmental Engineering, Politecnico di Milano, Piazza Leondardo da Vinci 32, 20133 Milan, Italy

1 Introduction

Piezoelectricity is the property of certain materials to generate an electric potential when subjected to mechanical stress and, conversely, to deform elastically when an electrical voltage is applied. These phenomena are referred to as the direct and inverse piezoelectric effects, respectively. In Piezoelectric Micromachined Ultrasonic Transducers (PMUTs), this coupled mechanical behavior enables both ultrasonic wave emission and reception: an applied alternating voltage induces diaphragm vibrations that radiate acoustic waves, while incident pressure waves deform the membrane and generate electrical signals. Thin piezoelectric films such as Lead Zirconate Titanate (PZT), Aluminum Nitride (AlN), and Scandium-doped Aluminum Nitride (ScAlN) are commonly integrated into multilayer diaphragms to achieve efficient electromechanical transduction [17].

With technological progress, systems increasingly demand higher miniaturization, complexity, and automation. Ultra-




Fig. 1 Cross-section of a representative PMUT stack, illustrating the multilayer structure in which each layer has a defined thickness and distinct material properties

sonic devices have followed this trend, particularly through the development of Micromachined Ultrasonic Transducers (MUTs) using Micro-Electro-Mechanical Systems (MEMS) technology, which allow compact, high-performance systems [7, 14]. Among them, PMUTs have gained significant attention due to their advantageous electrical properties and compatibility with low-voltage electronics [17]. Their Complementary Metal-Oxide Semiconductor compatibility enables integration of sensing and control electronics on a single substrate, supporting the development of miniaturized and efficient ultrasonic systems [7]. Figure 1 provides a representation of a multilayered PMUT stack cross-section.

Compared to bulk piezoelectric transducers, PMUTs benefit from batch microfabrication, resulting in compact arrays with improved yield and reduced cost. In contrast to capacitive MUTs, they do not require high bias voltages and exhibit lower impedance, simplifying electronic interfacing and improving energy efficiency [7, 14].

These advantages have enabled PMUT arrays to find applications in medical imaging, fingerprint sensing, sonar, non-destructive testing, flow measurement, and underwater communication [7, 14, 17].

As PMUT arrays grow in geometric complexity and face increasingly demanding operational constraints, the numerical models required to simulate their behavior accurately become significantly more computationally intensive. Consequently, there is a pressing need to develop innovative modeling strategies that deliver fast, accurate, and reliable tools capable of keeping up with the rapid evolution of PMUT technology [1].

The multiphysics nature of PMUT operation—combining structural deformation, piezoelectric coupling, electrical excitation, and acoustic radiation—makes modeling particularly challenging. Early analytical formulations based on classical plate and membrane theories provided closed-form estimates of resonance frequencies and modal behavior, but these simplified approaches cannot capture multilayer stacks, non-uniform electrode layouts, residual stresses, or fluid loading effects encountered in practical devices [7, 24].

To overcome these limitations, high-fidelity Finite Element Method (FEM) simulations have become the reference tool for detailed PMUT analysis. Three-dimensional multiphysics FEM models simultaneously solve structural mechanics and piezoelectric equations, enabling accurate prediction of mode shapes, displacement amplitudes, and electromechanical coupling coefficients across geometric and material parameters [13, 25]. Symmetry-based reductions and parametric sweeps are frequently employed to optimize diaphragm geometry and piezoelectric layer thickness [10, 28]. For realistic performance evaluation, structural models are often coupled with acoustic domains through pressure acoustics or fluid–structure interaction formulations. These coupled simulations allow estimation of radiated pressure fields, radiation impedance, directivity patterns, and bandwidth in air- or liquid-loaded conditions [11, 22, 24, 30]. Such analyses are crucial for air-coupled sensing and imaging arrays, where acoustic loading significantly affects resonance frequency and damping characteristics. Beyond linear frequency-domain analysis, nonlinear and transient multiphysics simulations have been explored to investigate complex dynamic phenomena. Time-domain coupling of piezoelectric, structural, and fluid equations has been used to model acoustic streaming and micropump actuation in ScAlN-based devices, revealing nonlinear mode interactions and even chaotic behavior under specific excitation regimes [16, 29]. These studies highlight the limitations of purely linear modal approximations when operating at large amplitudes.

While full 3D multiphysics FEM provides high predictive accuracy, it becomes prohibitively computationally expensive for large parameter sweeps and extended PMUT arrays. To address this issue, reduced-order modeling strategies have been developed. Modal superposition techniques project the electromechanical response onto a truncated basis of dominant eigenmodes extracted from high-fidelity simulations, significantly reducing computational cost while preserving accuracy [10, 17]. Equivalent circuit models based on Mason or Butterworth–Van Dyke formalisms [9] offer further simplification by translating electromechanical behavior into electrical analogs suitable for system-level analysis [8, 28]. At the array level, modeling complexity increases due to mutual acoustic coupling, mechanical crosstalk, and phase synchronization effects. Hybrid analytical–numerical formulations and reduced-dimensional acoustic models have been proposed to approximate inter-element interactions without explicitly resolving all structural degrees of freedom [8, 20, 31, 33]. Large-domain acoustic simulations typically incorporate absorbing boundary treatments such as perfectly matched layers (PMLs) to prevent artificial reflections [6, 12, 15, 32].

Despite substantial progress, key challenges remain: balancing computational efficiency and predictive accuracy for

large arrays, accounting for fabrication variability and residual stresses, capturing nonlinear behavior at high excitation levels, and ensuring stable acoustic–structure coupling in time-domain simulations. These limitations motivate the development of scalable numerical frameworks that combine model reduction with high-performance wave propagation solvers.

Unlike existing approaches, which typically rely on either fully resolved multiphysics models or simplified acoustic approximations [21], this work introduces a novel computational framework that combines reduced-order PMUT models, Discontinuous Galerkin (DG)-based acoustic coupling, and large-scale parallel computing. The resulting methodology enables accurate and scalable simulations of PMUT arrays with a complexity that is beyond the reach of current state-of-the-art techniques, see [2]. We consider an efficient model order reduction approach derived from [17] to reduce the computational complexity associated with simulating large PMUT arrays. The proposed novel framework is based on first solving the three-dimensional eigenvalue piezoelectric problem on a reference PMUT [2], from which a reduced set of significant vibration modes is extracted. The mechanical response of each transducer is then expressed as a linear combination of these modal shapes. The modal displacements are interpolated within a purely acoustic domain that models the PMUT arrays as bi-dimensional membranes. The proposed framework is implemented and validated in SPEED [4], an open-source software developed for large-scale wave propagation phenomena in complex three-dimensional media. Our framework integrates the Discontinuous Galerkin Spectral Element Method (DG-SEM) with domain decomposition and high-order polynomial approximation to ensure geometric flexibility and high-order accuracy. A key contribution of this work is the development of a distributed-memory parallel implementation, presented here for the first time, specifically designed to enable large-scale wave propagation simulations involving very large PMUT arrays. The framework leverages METIS [19] for mesh partitioning and MPI for inter-process communication, while explicit time integration provides excellent scalability and computational efficiency on High-Performance Computing (HPC) architectures. The primary challenges addressed include an optimized mesh generation and partitioning strategy, as well as the optimization of algorithmic operations to ensure both high-performance and accuracy in the simulations. By improving these aspects, the proposed approach leads to scalable, high-fidelity modeling of complex PMUT arrays in both transmission (TX) and reception (RX) phases.

The remainder of this paper is organized as follows. Section 2 formulates the acoustic problem coupled with the modal equations of the PMUT, addressing both the trans-

mission and reception scenarios. This section also includes a discussion on Perfectly Matched Layer (PML) [6, 12], highlighting its role in eliminating spurious reflections at the boundaries of the computational domain.

Section 3 presents the numerical discretization of the coupled problem in both space (using the DG-SEM) and time with Newmark scheme. Section 4 presents the domain decomposition strategy employed for efficient mesh partitioning, which is critical for handling large-scale PMUT arrays. Section 5 focuses on the implementation optimization designed to accelerate computations involving non-conforming meshes. Section 6 presents representative numerical results that validate the methodology and demonstrate its performance for large PMUT array simulations, and includes the evaluation of the scalability of the proposed approach on parallel computing architectures. Finally, in Sect. 7 we draw some conclusions and discuss future perspectives.


2 Multi-physics problem formulation for large PMUT array simulations

In this section, we introduce the general formulation of the coupled problem. Let $\Omega = \Omega_{\text{in}} \cup \Omega_{\text{out}}$ be a three-dimensional open Lipschitz domain, and $T \in \mathbb{R}^+$ be the final time. The boundary of the acoustic domain Ω is decomposed as $\partial\Omega = \Gamma_{\text{PMUT}} \cup \Gamma_{\text{N}} \cup \Gamma_{\text{ABC}}$, consisting of disjoint Γ_{PMUT} , Γ_{N} and Γ_{ABC} parts. Γ_{PMUT} is the portion of the boundary where the action of PMUT membranes is accounted for, Γ_{N} is where homogeneous Neumann conditions are imposed, and Γ_{ABC} is the remaining part of the boundary where absorbing condition are considered. Additionally, Γ_{I} represents an internal interface between the sub-domain Ω_{in} and Ω_{out} , cf. Fig. 2. The splitting of Ω_{in} and Ω_{out} enables the treatment of the two subdomains as potentially non-conforming regions while simultaneously allowing the assignment of distinct material properties within each subdomain. Moreover, the possible inclusion of the subdomain Ω_{PML} , as it will be described in the following section, can ensure the proper implementation of a Perfectly Matched Layer, thereby enabling the effective absorption of outgoing waves and the mitigation of artificial reflections at the truncated computational boundary.

In $\Omega = \Omega_{\text{in}} \cup \Omega_{\text{out}}$, with $\Omega_{\text{in}} \cap \Omega_{\text{out}} = \emptyset$, we consider the following problem

$$\begin{cases} p_{\text{tt}} = \nabla \cdot (c^2 \nabla p), & \text{in } \Omega \times (0, T], \\ p_{\text{t}} = p = 0, & \text{in } \Omega \times \{0\}, \\ \nabla p \cdot \mathbf{n} = \frac{1}{c} p_{\text{t}}, & \text{on } \Gamma_{\text{ABC}} \times [0, T], \\ \nabla p \cdot \mathbf{n} = 0, & \text{on } \Gamma_{\text{N}} \times [0, T], \\ p|_{\Omega_{\text{in}}} = p|_{\Omega_{\text{out}}}, & \text{on } \Gamma_{\text{I}} \times [0, T], \\ (\nabla p|_{\Omega_{\text{in}}} - \nabla p|_{\Omega_{\text{out}}}) \cdot \mathbf{n}_{\text{I}} = 0, & \text{on } \Gamma_{\text{I}} \times [0, T], \\ \nabla p \cdot \mathbf{n} = f_{\text{PMUT}}, & \text{on } \Gamma_{\text{PMUT}} \times [0, T], \end{cases} \quad (1)$$

Fig. 2 A three-dimensional view of the domain decomposition (left) and the corresponding two-dimensional slice with boundary labels (right) for problem (1). For the sake of completeness, in these pictures, we consider the presence of Ω_{PML} . When $\Omega_{\text{PML}} = \emptyset$, then Γ_{PML} reduces to Γ_{ABC}



where $p(\mathbf{x}, t)$ denotes the acoustic pressure, $\kappa = \rho c^2$ is the bulk modulus with $\rho = \rho(\mathbf{x})$ the acoustic density and $c = c(\mathbf{x})$ is the sound speed in the medium. A first-order absorbing boundary condition is imposed on Γ_{ABC} , while a hard wall boundary condition is set on Γ_N . Continuity of the pressure and of its normal derivative is imposed on Γ_I . The terms $\mathbf{n}$ and $\mathbf{n}_I$ represent the unit outer normal to the domain Ω and to the domain Ω_{in} , respectively. The acoustic problem is thus coupled with the piezo-electro-mechanical problem, describing the behavior of PMUTs, by the term f_{PMUT} . Due to the large number of membranes typically employed in realistic scenarios, solving the full physics of the PMUTs may be computationally unaffordable. Thus, it is preferable to consider model order reduction techniques, where the behavior of bi-dimensional PMUT membranes is described by ordinary differential equations. Here, we consider the model discussed in [17], where the following steps are considered to obtain the reduced formulation. For a three-dimensional PMUT membrane, the constitutive law of piezoelectricity, written in terms of elastic displacement and electric potential, is discretized by means of finite elements. Next, a static condensation procedure is applied to the resulting ODE system, cf. [17, Eq. (8)], to get a differential problem of the form $M\ddot{\mathbf{U}} + K\mathbf{U} = \mathbf{0}$, where M is the mass matrix and K is the stiffness matrix of the coupled piezoelectric problem, and $\mathbf{U}$ the unknown membrane displacement. Let us denote with $\mathbf{U}_m$, with $m \in \{1, \dots, N_m\}$, the array of nodal values representing the m th piezoelectric displacement eigenmode of a single, isolated PMUT membrane. These eigenvectors represent a mass-normalized orthogonal basis, namely:

$$\begin{aligned} \mathbf{U}_m^T M \mathbf{U}_m &= 1, & \mathbf{U}_m^T K \mathbf{U}_m &= \omega_m^2, & \forall m, \\ \mathbf{U}_n^T M \mathbf{U}_m &= 0, & \mathbf{U}_n^T K \mathbf{U}_m &= 0, & \forall n \neq m, \end{aligned}$$

where ω_m is the corresponding natural frequency of mode m .

When considering an array of N_p PMUTs, we denote the displacement field of the k th PMUT as $\mathbf{U}_k(\mathbf{x}, t)$, with $k \in \{1, \dots, N_p\}$. It is possible to represent this displacement on

a reduced basis of N_m modal vectors as:

$$\mathbf{U}_k(\mathbf{x}, t) = \sum_{m=1}^{N_m} q_{k,m}(t) \mathbf{U}_{k,m}(\mathbf{x}),$$

where $\mathbf{U}_{k,m}$ denotes the m th mode shape assigned to the k th PMUT, and $q_{k,m}(t)$ is the corresponding generalized modal coordinate. Projecting the mechanical balance equations onto each retained mode, one obtains the ordinary differential equation for $q_{k,m}$ introduced in [17]:

$$\begin{cases} \ddot{q}_{k,m} = -\omega_{k,m}^2 q_{k,m} + \int_{\gamma_k} p \mathbf{U}_{k,m} \cdot \mathbf{n} \, d\sigma + \eta_{k,m} \phi_k, & \text{in } (0, T], \\ \phi_k = \phi_{k,0}(t) \mathbb{1}_{\{0 \leq t \leq \bar{t}\}}(t) + \frac{1}{C} \sum_m \eta_{k,m} q_{k,m} \mathbb{1}_{\{\bar{t} < t \leq T\}}(t), & \text{in } [0, T], \\ \nabla p \cdot \mathbf{n} = -\kappa \sum_{k,m} \ddot{q}_{k,m} \mathbf{U}_{k,m} \cdot \mathbf{n}, & \text{on } \gamma_k \times [0, T], \\ q_{k,m} = \dot{q}_{k,m} = 0, & \text{at } t = 0. \end{cases} \quad (2)$$

The total boundary portion occupied by all PMUTs is given by $\Gamma_{\text{PMUT}} = \bigcup_{1 \leq k \leq N_p} \gamma_k$, with γ_k denoting the boundary portion occupied specifically by the k th PMUT. In this coupled formulation, $\phi_k(t)$ is a low-voltage activation function, C is the piezoelectric capacitance of the PMUTs, $\omega_{k,m}$ is the natural frequency of mode m for the k th PMUT, while $\eta_{k,m}$ is the electro-mechanical coupling coefficient relating the applied potential ϕ_k to mode m . Due to their nature, PMUTs are able to capture both the transmission and reception phases. Setting $\bar{t} \in (0, T)$, for $t \leq \bar{t}$ the voltage $\phi_{k,0}$ is applied to produce an acoustic pressure, while for $t > \bar{t}$ it is directly produced by the pressure on the surface γ_k .

Remark We stress that the reference eigenmodes $\mathbf{U}_m$ and eigenvalues ω_m are computed for a single reference PMUT membrane in a pre-processing phase, as reported in [17]. These are then mapped to the array elements as $\mathbf{U}_{k,m}$ and $\omega_{k,m}$ to serve as the input for the acoustic simulation, as explained in the sequel.

2.1 Boundary treatment by Perfectly Matched Layer conditions

To correctly model wave propagation in unbounded domains, we consider an additional layer Ω_{PML} attached to Ω_{out} where PML conditions are imposed, cf. Fig. 2. Then, in agreement with [12, 15], we modify the first equation of (1) describing the acoustic wave propagation as follows

$$\begin{cases} p_{tt} + \alpha p_t + \beta p = \nabla \cdot (c^2 \nabla p) + \nabla \cdot \Phi - \gamma \psi, & \text{in } \Omega_{\text{PML}} \times (0, T], \\ \Phi_t = Z_1 \Phi + c^2 Z_2 \nabla p + c^2 Z_3 \nabla \psi, & \text{in } \Omega_{\text{PML}} \times (0, T], \\ \psi_t = p, & \text{in } \Omega_{\text{PML}} \times (0, T], \\ \psi = 0, & \text{in } \Omega_{\text{PML}} \times \{0\}, \\ \Phi = 0, & \text{in } \Omega_{\text{PML}} \times \{0\}, \end{cases} \quad (3)$$

where $\Phi(\mathbf{x}, t)$ represents the (vector) variable that introduces additional diffusion to the problem, and $\psi(\mathbf{x}, t)$ is an auxiliary variable enabling the formulation of the PML damping in a purely local differential form in time. The scalars $\alpha = (\zeta_1 + \zeta_2 + \zeta_3)$, $\beta = (\zeta_1 \zeta_2 + \zeta_2 \zeta_3 + \zeta_3 \zeta_1)$, $\gamma = \zeta_1 \zeta_2 \zeta_3$, and the matrices

$$Z_1 = \text{diag}(-\zeta_1, -\zeta_2, -\zeta_3),$$

$$Z_2 = \text{diag}(\alpha - 2\zeta_1, \alpha - 2\zeta_2, \alpha - 2\zeta_3),$$

$$Z_3 = \text{diag}(\zeta_2 \zeta_3, \zeta_3 \zeta_1, \zeta_1 \zeta_2),$$

are related to the specific damping profile $\zeta_i(x_i)$, $i \in \{1, 2, 3\}$ selected in the three spatial directions. The idea of the PML equations is to provide additional damping and diffusivity to the acoustic equation to shrink the waves' amplitude before they touch the external boundary of the PML, i.e., Γ_{ABC} . In this paper, in agreement with [12], we choose the following damping profile

$$\zeta_i(x_i) = \tilde{\zeta}_i \left(\frac{(|x_i| - h_i)}{\ell_i} - \frac{\sin\left(\frac{2\pi(|x_i| - h_i)}{\ell_i}\right)}{2\pi} \right) \mathbb{1}_{\{h_i \leq |x_i| \leq h_i + \ell_i\}}(x_i),$$

where for each i th direction, $i \in \{1, 2, 3\}$, h_i denotes half of the thickness of the propagation domain and ℓ_i the thickness of the PML. The function $\zeta_i(x_i)$ vanishes with continuity at the interface Γ_{PML} with the purely acoustic domain, and the quantity $\tilde{\zeta}_i$ is related to the sound speed c in the acoustic medium, the thickness of the PML, and the relative reflection coefficient R of the acoustic wave at the PML interface, as follows

$$\tilde{\zeta}_i = \frac{c}{\ell_i} \log\left(\frac{1}{R}\right), \quad i = 1, 2, 3.$$

The inverse of the reflection coefficient R has to be selected large enough to avoid the wave being reflected at the interface, but it has to be kept small enough to have a less

steep function $\zeta_i(x_i)$ in the transition region between the interface Γ_{PML} and the end of the PML. In the following, $R \in [10^{-5}, 10^{-4}]$.

The full set of unknowns and parameters defined in Eqs. (1), (2), and (3) is reported in Table 1.

3 Space discretization: Discontinuous Galerkin Spectral Element Method

The numerical discretization of (1) is based on a DG-SEM, which allows both for high-order polynomials of SEM and grid flexibility of the DG method [4] (Fig. 3).

We next discuss the derivation of the DG-SEM formulation by assuming, without loss of generality, that $\Omega_{\text{PML}} = \emptyset$. Let $\Omega = \Omega_{\text{in}} \cup \Omega_{\text{out}}$ be a polygonal domain and consider its partition $\mathcal{T}_h = \mathcal{T}_{\text{in}} \cup \mathcal{T}_{\text{out}}$, s.t. $\mathcal{T}_{\text{in}} \cap \mathcal{T}_{\text{out}} = \emptyset$ with $\mathcal{T}_{\text{in}}$ and $\mathcal{T}_{\text{out}}$ containing hexahedral elements K_{in} and K_{out} of different granularity h_{in} and h_{out} , respectively. By defining N_{in} , and N_{out} , the number of elements in each partitions, one has $\mathcal{T}_{\text{in}} = \{K_{\text{in}}^j\}_{j=1}^{N_{\text{in}}}$ and $\mathcal{T}_{\text{out}} = \{K_{\text{out}}^j\}_{j=1}^{N_{\text{out}}}$, respectively. Let $\hat{K} = [-1, 1]^3$ be the reference element, and let $X_{K_a} : \hat{K} \rightarrow K_a$ be a trilinear invertible map with positive Jacobian J_{K_a} such that $\forall K_a \in \mathcal{T}_a$, $K_a = X_{K_a}(\hat{K})$, for $a = \{\text{in}, \text{out}\}$, with $X_{K_a}^{-1} : K_a \rightarrow \hat{K}$. Next, we define the internal faces as $F = \bar{K}_{\text{in}}^j \cap \bar{K}_{\text{out}}^k$ for certain $j \in \{1, \dots, N_{\text{in}}\}$ and $k \in \{1, \dots, N_{\text{out}}\}$ and collect them in $\mathcal{F}^I$. Similarly, we define the boundary faces $F = \bar{K}_a^{j_a} \cap \partial\Omega$ for certain $j_a \in \{1, \dots, N_a\}$ and collect them in $\mathcal{F}^B$. Moreover, we split $\mathcal{F}^B$ as $\mathcal{F}^B = \mathcal{F}_N^B \cup \mathcal{F}_{\text{ABC}}^B \cup \mathcal{F}_{\text{PMUT}}^B$, where Neumann, absorbing, and coupling conditions (with PMUT membranes) are imposed, respectively. Note that boundary faces are compliant with this subdivision.

Consider now two adjacent elements $K^- \in \mathcal{T}_{\text{out}}$ and $K^+ \in \mathcal{T}_{\text{in}}$, sharing a face $F \in \mathcal{F}^I$ and define $\mathbf{n}^+$ and $\mathbf{n}^-$ to be the outward unit normal vectors to the face F from K^+ and K^- , respectively. Figure 4 provides a visual representation of two adjacent elements.

For each face $F \in \mathcal{F}^I$ and for any sufficiently regular scalar function u and vector $\mathbf{v}$ we define the jump operator $\llbracket \cdot \rrbracket$ and the average operator $\{ \cdot \}$ by

$$\{u\} = \frac{1}{2}(u^+ + u^-), \quad \llbracket u \rrbracket = u^+ \mathbf{n}^+ + u^- \mathbf{n}^-,$$

$$\{\mathbf{v}\} = \frac{1}{2}(\mathbf{v}^+ + \mathbf{v}^-), \quad \llbracket \mathbf{v} \rrbracket = \mathbf{v}^+ \cdot \mathbf{n}^+ + \mathbf{v}^- \cdot \mathbf{n}^-,$$

where $\pm$ denotes the trace on u and $\mathbf{v}$ taken within the interior of $K^\pm$, respectively. Similarly, for each face $F \in \mathcal{F}^B$ we set $\{u\} = u$, $\llbracket u \rrbracket = u \mathbf{n}$ and $\{\mathbf{v}\} = \mathbf{v}$, $\llbracket \mathbf{v} \rrbracket = \mathbf{v} \cdot \mathbf{n}$.

Table 1 The three tables report the unknowns and the parameters of Eqs. (1), (2), and (3) with the corresponding units of measurement

Unknown	Symbol	Unit
Pressure	p	[Pa]
m th modal coordinate	q_m	[m]
Voltage	ϕ	[V]
PML auxiliary function	ψ	[Pa s]
PML auxiliary function	Φ	[Pa m s ⁻²]
Acoustic parameter	Symbol	Unit
PML damping function	α, Z_1, Z_2	[s ⁻¹]
PML damping function	β, Z_3	[s ⁻²]
PML damping function	γ	[s ⁻³]
Acoustic speed	c	[m s ⁻¹]
PMUTs parameter	Symbol	Unit
Natural frequency of mode m	ω_m	[rad s ⁻¹]
Electro-mechanical coupling coefficient	η_m	[C m ⁻¹]
Bulk modulus	κ	[Pa]
Piezoelectric capacitance	C	[F]
Applied potential	ϕ_0	[V]




Fig. 3 Example of mesh employed in the DG-SE approach. Different mesh sizes and polynomial degrees are employed for Ω_{in} and Ω_{out} , resulting in a non-conforming mesh




Fig. 4 Representation of two adjacent hexahedra K^+ and K^- with the shared face F and the corresponding normal vectors $\mathbf{n}^+$ and $\mathbf{n}^-$

For $a = \{\text{in}, \text{out}\}$ we consider the space

$$V_a^{r_a} = \left\{ v \in \mathcal{C}^0(\overline{\Omega_a}) : v|_{K_a} \circ \mathbf{X}_{K_a}^{-1} \in \mathbb{Q}_{r_a}(\hat{K}), \quad \forall K_a \in \mathcal{T}_a \right\},$$

where $\mathbb{Q}_{r_a}(\hat{K})$ is the space of polynomials of degree less than or equal to $r_a \geq 1$ in each coordinate direction. Next, we introduce

$$V^{\text{DG}} = \left\{ v \in L^2(\Omega) : v|_{\Omega_a} \in V_a^{r_a} \right\},$$

and provide the semi-discrete formulation: $\forall t \in (0, T]$ find $p_h = p_h(t) \in V^{\text{DG}}$ s.t.

$$\begin{aligned} & m(\ddot{p}_h, v_h) + a(p_h, v_h) + d(\dot{p}_h, v_h) \\ &= -\kappa \sum_{m=1}^{N_m} \int_{\Gamma_{\text{PMUT}}} \ddot{q}_m \mathbf{U}_m \cdot \mathbf{n} v_h \, d\sigma, \quad \forall v_h \in V^{\text{DG}}, \end{aligned} \quad (4)$$

with $p_h(0) = \dot{p}_h(0) = 0$ and where, for any $1 \leq m \leq N_m$, it holds

$$\begin{aligned} \ddot{q}_m(t) + \omega_m^2 q_m(t) &= \int_{\Gamma_{\text{PMUT}}} p_h(t) \mathbf{U}_m \cdot \mathbf{n} \, d\sigma + \eta_m \phi(t) \\ \forall t &\in (0, T]. \end{aligned} \quad (5)$$

For any $\omega, v \in V^{\text{DG}}$, the bilinear forms in (4) are defined as:

$$m(\omega, v) = \sum_{K_s \in \mathcal{T}_{\text{in}} \cup \mathcal{T}_{\text{out}}} \int_{K_s} \omega v \, d\Omega, \quad (6)$$

$$\begin{aligned} a(\omega, v) = & \sum_{K_s \in \mathcal{T}_{\text{in}} \cup \mathcal{T}_{\text{out}}} \int_{K_s} c^2 \nabla \omega \cdot \nabla v \, d\Omega \\ & - \sum_{F \in \mathcal{F}^I} \int_F \{c^2 \nabla_h \omega\} \cdot \llbracket v \rrbracket \, d\sigma \\ & - \sum_{F \in \mathcal{F}^I} \int_F \{c^2 \nabla_h v\} \cdot \llbracket \omega \rrbracket \, d\sigma \\ & + \sum_{F \in \mathcal{F}^I} \int_F \gamma_F \llbracket \omega \rrbracket \cdot \llbracket v \rrbracket \, d\sigma, \end{aligned} \quad (7)$$

$$d(\omega, v) = \sum_{F \in \mathcal{F}_{\text{ABC}}^B} \int_F c \omega v \, d\sigma, \quad (8)$$

where ∇_h represents the broken gradient and the parameter $\gamma_F = \alpha \bar{c}^2 \frac{\max(r_{\text{in}}, r_{\text{out}})^2}{\min(h_{\text{in}}, h_{\text{out}})}$, with $\bar{c}$ being the harmonic average of c on F and with α a sufficiently large positive parameter at our disposal, see, e.g., [5].

3.1 Algebraic formulation and time integration

Let $\{\psi_i\}_{i=1}^{N_h}$ be the Lagrangian basis functions for V^{DG} , cf. [3], and set $v_h(\mathbf{x}) = \psi_i(\mathbf{x})$, and $p_h(\mathbf{x}, t) = \sum_{j=1}^{N_h} p_j(t) \psi_j(\mathbf{x})$ in (6)–(8). Rearranging the terms in matrix form, the semi-discrete algebraic formulation of problem (4)–(5) reads: for any $t \in (0, T]$ find $\mathbf{p}(t) = (p_1(t), \dots, p_{N_h}(t))^T \in \mathbb{R}^{N_h}$ and $\mathbf{q}(t) = (q_1(t), \dots, q_m(t))^T \in \mathbb{R}^{N_m}$, s.t.

$$M \ddot{\mathbf{p}}(t) + K \mathbf{p}(t) + C \dot{\mathbf{p}}(t) = \mathbf{f}(\ddot{\mathbf{q}}(t)), \quad (9)$$

$$\ddot{\mathbf{q}}(t) + W \mathbf{q}(t) = \mathbf{g}(\mathbf{p}(t), \phi(t)), \quad (10)$$

with $\mathbf{p}(0) = \dot{\mathbf{p}}(0) = \mathbf{0}$, $\mathbf{q}(0) = \dot{\mathbf{q}}(0) = \mathbf{0}$. In (9), $M \in \mathbb{R}^{N_h \times N_h}$ is the diagonal mass matrix associated with $m(\cdot, \cdot)$, i.e., the matrix representation of the L^2 inner product, evaluated with reduced integration, $K \in \mathbb{R}^{N_h \times N_h}$ is the stiffness matrix related to $a(\cdot, \cdot)$, $C \in \mathbb{R}^{N_h \times N_h}$ is the matrix encoding the absorbing conditions $d(\cdot, \cdot)$, while $\mathbf{f} \in \mathbb{R}^{N_h}$ is the vector describing the forcing term acting on Γ_{PMUT} , i.e.,

$$\begin{aligned} f_i(t) = & -\kappa \sum_{m=1}^{N_m} \ddot{q}_m(t) \int_{\Gamma_{\text{PMUT}}} \mathbf{U}_m(t, \mathbf{x}) \cdot \mathbf{n} \psi_i(\mathbf{x}) \, d\sigma, \\ & 1 \leq i \leq N_h. \end{aligned} \quad (11)$$

In (10), $W \in \mathbb{R}^{N_m \times N_m}$ is the matrix $W = \text{diag}(\omega_1^2, \dots, \omega_{N_m}^2)$ and $\mathbf{g} \in \mathbb{R}^{N_m}$ is such that

$$\begin{aligned} g_m(t) = & \sum_{j=1}^{N_h} p_j(t) \int_{\Gamma_{\text{PMUT}}} \psi_j(\mathbf{x}) \mathbf{U}_m(\mathbf{x}) \cdot \mathbf{n} \, d\sigma \\ & + \eta_m \phi(t) \quad 1 \leq m \leq N_m. \end{aligned} \quad (12)$$

Next, we let $S \in \mathbb{N}^+$ be the number of sub-intervals of length Δt in which the time interval $[0, T]$ is divided in such a way that $t^n = n\Delta t$, $n \in \{0, \dots, S\}$. Equations (9)–(10) can be discretized with the following second-order predictor-corrector Newmark scheme, shown in Algorithm 1. Note that a similar scheme is employed when a PML boundary is introduced in the computational model.

The Newmark time-integration method is employed with $\gamma = 1/2$ and $\beta = 0$. For this choice, the scheme is second-order accurate, and the corresponding update relations reduce to the explicit leapfrog formulation. Since the method is explicit, it is conditionally stable under the Courant-Friedrichs-Lewy (CFL) condition

$$\Delta t \leq K_{\text{CFL}} \frac{h_{\min}}{\max(c(\mathbf{x}))}, \quad \mathbf{x} \in \Omega, \quad (13)$$

being $h_{\min}$ the minimum Gauss–Legendre–Lobatto node spacing across all mesh elements [23], and K_{CFL} the CFL constant to be determined (cf. Sect. 6.1).

Algorithm 1 Newmark time integration scheme

- 1: Set $\mathbf{q}^0 = \dot{\mathbf{q}}^0 = \ddot{\mathbf{q}}^0 = \mathbf{0}$ and $\mathbf{p}^0 = \dot{\mathbf{p}}^0 = \ddot{\mathbf{p}}^0 = \mathbf{0}$ and compute the voltage ϕ^0
 - 2: **for** $n = 0, \dots, S$ **do**
 - 3: Predictor (PMUT): set $\mathbf{q}^{n+1} \leftarrow \mathbf{q}^n + \Delta t \dot{\mathbf{q}}^n + \frac{1}{2} \Delta t^2 \ddot{\mathbf{q}}^n$
 - 4: Predictor (PMUT): set $\dot{\mathbf{q}}^{n+1} \leftarrow \dot{\mathbf{q}}^n + \frac{1}{2} \Delta t \ddot{\mathbf{q}}^n$
 - 5: Solve (PMUT): $\ddot{\mathbf{q}}^{n+1} \leftarrow -W \mathbf{q}^{n+1} + \mathbf{g}(\mathbf{p}^n, \phi^n)$
 - 6: Corrector (PMUT): $\dot{\mathbf{q}}^{n+1} \leftarrow \dot{\mathbf{q}}^{n+1} + \frac{1}{2} \Delta t \ddot{\mathbf{q}}^{n+1}$
 - 7: Predictor (acoustic): set $\mathbf{p}^{n+1} \leftarrow \mathbf{p}^n + \Delta t \dot{\mathbf{p}}^n + \frac{1}{2} \Delta t^2 \ddot{\mathbf{p}}^n$
 - 8: Predictor (acoustic): set $\dot{\mathbf{p}}^{n+1} \leftarrow \dot{\mathbf{p}}^n + \frac{1}{2} \Delta t \ddot{\mathbf{p}}^n$
 - 9: Solve (acoustic): $\ddot{\mathbf{p}}^{n+1} \leftarrow M^{-1}(\mathbf{f}(\ddot{\mathbf{q}}^{n+1}) - K \mathbf{p}^{n+1} - C \dot{\mathbf{p}}^{n+1})$
 - 10: Corrector (acoustic): $\dot{\mathbf{p}}^{n+1} \leftarrow \dot{\mathbf{p}}^{n+1} + \frac{1}{2} \Delta t \ddot{\mathbf{p}}^{n+1}$
 - 11: Update: $\mathbf{q}^n \leftarrow \mathbf{q}^{n+1}$, $\dot{\mathbf{q}}^n \leftarrow \dot{\mathbf{q}}^{n+1}$, $\ddot{\mathbf{q}}^n \leftarrow \ddot{\mathbf{q}}^{n+1}$
 - 12: Update: $\mathbf{p}^n \leftarrow \mathbf{p}^{n+1}$, $\dot{\mathbf{p}}^n \leftarrow \dot{\mathbf{p}}^{n+1}$, $\ddot{\mathbf{p}}^n \leftarrow \ddot{\mathbf{p}}^{n+1}$
 - 13: **end for**
-

4 Parallel implementation strategies for PMUT simulations

In the context of HPC, achieving competitive runtimes for large-scale simulations is of paramount importance. This

requires fully leveraging the computational capabilities of parallel codes on multi-core clusters. A key challenge in this context is the efficient distribution of computational workload across cores, which often involves adopting domain-partitioning strategies tailored to the specific requirements of the simulation. In this work, we address a particularly challenging load-balancing problem arising from the need to account for multiple constraints simultaneously. Specifically, our goal is to balance the computational cost associated with the

1. boundary conditions for the PMUTs-array;
2. acoustic propagation in $\Omega_{\text{out}} \cup \Omega_{\text{PML}}$;
3. connectivity setup for non matching grids on Γ_I .

Additionally, we aim to ensure that the membranes representing the PMUTs, namely the boundary Γ_{PMUT} , are not split across multiple cores, as such fragmentation can degrade both performance and efficiency due to increased inter-process communication and loss of structural coherence. To this end, we propose a tailored partitioning strategy that respects the structural integrity of the PMUTs while distributing the simulation load evenly across processing units. This approach enhances the scalability of the simulation framework and ensures efficient utilization of HPC resources without compromising the accuracy of the underlying discretization.

4.1 Setup of the computational domain structure

The computational domain considered in the present case study consists of three main regions. The first one, which is referred to as *inner domain*, Ω_{in} , contains the PMUT array and constitutes the region of primary geometrical and physical interest. The second one, which is referred to as *outer domain*, Ω_{out} , represents the acoustic wave propagation region. The third one, Ω_{PML} , describes the external diffusive layer where Perfectly Matched Layer conditions are applied in order to mimic an unbounded medium.

For geometric fidelity and accuracy requirements associated with the vibrating membranes, the grid in Ω_{in} must be highly refined. Imposing a mesh of comparable resolution throughout Ω_{out} would lead to a prohibitive number of elements-potentially on the order of billions-thus rendering the computation intractable with standard computational resources. To reduce the overall number of degrees of freedom while preserving the required accuracy near the membranes, a coarser discretization is therefore adopted in Ω_{out} .

Accordingly, the inner and the outer subdomains are meshed independently, allowing for local refinement in Ω_{in} and a coarser mesh in the propagation region Ω_{out} . Conformity is instead imposed between Ω_{out} and Ω_{PML} to ensure a

consistent treatment of the absorbing layer. The use of distinct discretizations in Ω_{in} and Ω_{out} results in a globally nonconforming mesh. To properly couple the two regions, an internal interface Γ_I is introduced between Ω_{in} and Ω_{out} , enabling the matching of the finer and coarser meshes across the subdomain boundary.

4.1.1 Load balancing strategy for $\Omega_{\text{out}} \cup \Omega_{\text{PML}}$

The two domains Ω_{out} and Ω_{PML} are partitioned directly using Metis 5.1.0 [19]. The idea is to exploit multiple times a Metis routine, namely `METIS_PartMeshDual`, that receives as input the set of hexahedra composing the two regions and returns the ownership of the mesh elements to the computational cores. The routine can be directly applied to the full set of mesh elements composing Ω_{out} and Ω_{PML} if the primary interest is to balance the number of hexahedra among the unit processes. However, in most cases, we will obtain an unbalanced workload for the processors, since not all processors will compute the integral terms on $\mathcal{F}^I$ in (7), cf. Fig. 5 (left).

To overcome this bottleneck, we first split Ω_{out} into two subdomains: Ω_{DG} containing the hexahedral elements that share a face with Γ_I and the remaining ones. Next, to equally distribute elements over processors, we call three times `METIS_PartMeshDual`: one for Ω_{PML} , one for Ω_{DG} and one for the remaining part $\Omega_{\text{out}} \setminus \Omega_{\text{DG}}$. Figure 5 shows the Ω_{DG} layer connected to the outer domain Ω_{out} (middle) and a schematic representation of the Metis partition (right).

4.1.2 Mesh generation and load balancing strategy for Ω_{in}

The concept behind the mesh generation for the inner domain is to define a set of fundamental unit blocks that can be combined to form the complex membrane pattern of the PMUT array. In general, the membranes are not arranged in a Cartesian x - y grid but are only aligned along the x -direction, as shown in Fig. 6. If we define *row* the set of membranes aligned along x and *column* the set of PMUTs that winds along y , then each row in the array is offset relative to the previous one by a distance determined by the pitch of the PMUT, d_p , that is, the distance of the centers of two neighboring membranes, see Fig. 6.

While element partitioning for $\Omega_{\text{out}} \cup \Omega_{\text{PML}}$ is performed by employing the Metis library, for the inner domain Ω_{in} , we consider a different strategy. Indeed, we need a sufficiently refined mesh for Ω_{in} to accurately capture the geometry of the PMUT membranes and to properly resolve the applied voltage input signal on them. On the other hand we need to balance among processors: (i) the hexahedral elements that share a face with Γ_I , (ii) the number of PMUT membranes, and (iii) the remaining elements. Moreover, an additional issue to face in simulations involving a large computational




Fig. 5 Schematic illustration of the outer domain decomposition: non-optimized partitioning (left), DG layer Ω_{DG} (middle), optimal partitioning (right). The set of processors is $\{P_1, P_2, P_3, P_4\} = \{\text{blue}, \text{yellow}, \text{red}, \text{green}\}$. (color figure online)




Fig. 6 Top view of a minimal example of inner domain Ω_{in} . The elementary structures originating Ω_{in} are depicted with different colors. The circular regions represent the membranes γ_k belonging to Γ_{PMUT} , while the remaining part is Γ_N . (color figure online)

domain is generating and extracting the mesh (to be partitioned) in a single step due to RAM limitations. To account for all these constraints, the main approach proposed in this study is to cut the mesh into smaller blocks and provide a tool to merge them into a single mesh directly within the code, since the number of membranes involved is extremely large, reaching more than ten thousand PMUTs. Consequently, generating the complete mesh in a single step is often impractical due to memory limitations and time resources.

The proposed strategy, which simultaneously considers the constraints imposed by PMUT arrays and the requirements of load balancing, is to generate a set of elementary files, $\{f_s^{el}\}$,

with $s \in \{1, \dots, N_{sub}^{el}\}$, each containing the mesh of a subdomain of the inner region Ω_{in} . These subdomains are then assigned to individual computational processes, which collectively reconstruct the complete array structure. Because PMUT arrays are designed according to a predefined pattern, multiple processes can adopt the same elementary file, applying an appropriate coordinate shift to give rise to the global arrangement. Figure 7 (left) provides an example of how those blocks appear.

In this way, four main constraints are directly satisfied: the PMUTs are distributed among different cores, keeping them entirely on the same process unit, the load balance of the inner domain is guaranteed, the DG faces on the finer side are well balanced, and the RAM limitation is overcome. Figure 6 illustrates a minimal example of the inner domain, along with the elementary structures used in its construction.

Next, we detailed the procedure developed to address the mesh generation and partitioning in Ω_{in} , when the number of PMUTs N_p is sufficiently large. Let $N_{p,row}$ denote the number of PMUTs in a row, and let $N_{block} = N_{p,row} + 1$ represent the number of elementary blocks required to form a complete row in the inner domain. The procedure begins by verifying that the number of computational cores N_{proc} is an integer multiple of N_{block} , i.e. $N_{proc} = MN_{block}$ with $M \in \mathbb{N}^+$. Upon failure, the system raises an exception in order to avoid unfavorable workload imbalances, particularly within the finer regions of the mesh.

Subsequently, a connectivity structure A_i , for $i = 0, \dots, N_{proc} - 1$, is constructed for each process i to define adjacency relationships. This ensures that each process i has precise knowledge of its neighboring domains labeled in A_i , and thereby enabling efficient reconstruction of the global mesh connectivity. As an example, we can retrieve the mesh depicted in Fig. 6 by setting $N_{block} = 3$, and $M = 3$ in Fig. 7. The corresponding blocks $B_{n,m}$ and the core-ownership $L_{n,m}$ of each block, for $n = 1, \dots, N_{block}$, $m = 1, \dots, M$, can be

grouped in the following matrices, respectively:

$$B = \begin{bmatrix} B_{3,1} & B_{3,2} & B_{3,3} \\ B_{2,1} & B_{2,2} & B_{2,3} \\ B_{1,1} & B_{1,2} & B_{1,3} \end{bmatrix}, \quad L = \begin{bmatrix} 2 & 5 & 8 \\ 1 & 4 & 7 \\ 0 & 3 & 6 \end{bmatrix},$$

from which the adjacency structure A_i , $\forall i = 0, \dots, 8$, is derived

$$\begin{aligned} A_0 &= [1 \ 3 \ 4], & A_1 &= [0 \ 2 \ 3 \ 4 \ 5], & A_2 &= [1 \ 4 \ 5], \\ A_3 &= [0 \ 1 \ 4 \ 6 \ 7], & A_4 &= [0 \ 1 \ 2 \ 3 \ 5 \ 6 \ 7 \ 8], & A_5 &= [1 \ 2 \ 4 \ 7 \ 8], \\ A_6 &= [3 \ 4 \ 7], & A_7 &= [3 \ 4 \ 5 \ 6 \ 8], & A_8 &= [4 \ 5 \ 7]. \end{aligned}$$

So far, all the elementary files f_i^{el} are independent of each other, due to the fact that nodes, faces, and hexahedra have a local numeration starting from 1. Then, we need an algorithm to merge all elementary blocks and define a global numbering for the mesh elements. We start by writing all nodes lying on shared interfaces $\Gamma_{n,m}^s$, for $n = 1, \dots, N_{\text{block}}$, $m = 1, \dots, M$, cf. Fig. 7 (right), in specific files S_i , $i = 0, \dots, N_{\text{proc}} - 1$. A quicksort algorithm orders the nodes according to the coordinates x , y , and z ; in this way, the procedure for matching the nodes' IDs of different interfaces $\Gamma_{n,m}^s$ that represent the same physical surface is fast and affordable. Notice that this step must be performed sequentially, and the connectivity structures A_i provide a way to speed up the procedure. Rank zero keeps its original identifiers, while all other processes receive the identifiers of shared nodes from neighboring ranks with smaller IDs, ensuring that each common node is assigned a unique identifier and duplicates are removed. This procedure leads to non-continuous local numbering due to missing identifiers. A final global renumbering step is then performed across all mesh files to restore continuity and establish a consistent global numbering of nodes, faces, and elements. The full procedure is synthetically shown in Algorithm 2.

Algorithm 2 Inner domain partition and merging algorithm

Require: N_{proc} multiple of N_{block}

- 1: **for** $i \in \{0, \dots, N_{\text{proc}} - 1\}$ **do**
 - 2: Construct core-connectivity structure A_i
 - 3: Assign to core i its mesh file f_i^{el}
 - 4: Renumber nodes, faces, and elements of f_i^{el}
 - 5: Sort and store blocks' interface nodes in file S_i
 - 6: **for** j in A_i s.t. $j < i$, $i \neq 0$ **do**
 - 7: Match nodes in file S_i and S_j
 - 8: Retain node IDs of file S_j and discard IDs in S_i
 - 9: **end for**
 - 10: **end for**
 - 11: Renumber nodes according to global IDs
-

5 Fast assembly of interface terms for non-matching grids

The nonconformity between the inner and outer discretizations entails the construction of an interface mesh on $\Gamma_I = \partial\Omega_{\text{in}} \cap \partial\Omega_{\text{out}}$, defined as the geometric intersection of two hexahedral meshes with different resolutions. The number of resulting DG interfaces increases with the number of membranes under consideration. Efficiently identifying and matching interface elements becomes a critical task before the numerical solution of the governing equations can be carried out. While this strategy provides clear advantages in terms of geometric flexibility and reduction of the number of degrees of freedom, it also introduces significant challenges. In particular, a fast neighbor-matching algorithm is required to determine pairs of elements sharing DG interfaces, with the dual objective of reducing preprocessing overhead and improving performance during the time-stepping loop. In this work, we address these issues and propose an optimized methodology for interface construction and matching on Γ_I , significantly reducing the computational burden associated with the intersection process. Note that the intersection faces $F \in \mathcal{F}^I$ between the coarse and the fine boundaries are defined starting from the quadrature nodes located on the surface with the finer resolution. These quadrature nodes therefore determine the connectivity between elements (without explicitly computing the geometric intersection of the faces), which is a crucial aspect for the evaluation of the integrals in (7), cf. [18]. In particular, we focus our attention to the terms in (7) of the form


$$\int_F g^-(\mathbf{x}) f^+(\mathbf{x}) d\sigma, \quad (14)$$

being $F = \partial K^+ \cap \partial K^-$ and where superscript $\pm$ denotes the restriction of polynomial functions g and f to K^+ or K^- , respectively. To compute the integral in (14), we recall that every element K of the mesh is the image through a trilinear map X_K of the reference element $\hat{K} = (-1, 1)^3$. Therefore, using the change of variables $\mathbf{x} = X_{K^\pm}(\hat{\mathbf{x}})$ it holds

$$\begin{aligned} & \int_F g^-(\mathbf{x}) f^+(\mathbf{x}) d\sigma \\ &= \int_{\hat{F}} g^-(X_{K^-}(\hat{\mathbf{x}}^-)) f^+(X_{K^+}(\hat{\mathbf{x}}^+)) |J_{X_{K^+}}(\hat{\mathbf{x}}^+)| d\hat{\sigma}, \end{aligned}$$

being $\hat{F}$ (a portion of) a face of $\partial\hat{K}$ mapped through X_{K^-} or X_{K^+} to F , see Fig. 8. The above integral is thus evaluated using a suitable quadrature rule with N_q nodes and $\{\omega_q\}_{q=1}^{N_q}$ weights as follows

Fig. 7 Schematic representation of the mesh partitioning into blocks $B_{n,m}$ (left) and the corresponding interfaces $\Gamma_{n,m}^s$, with $n = 1, \dots, N_{\text{block}}$ and $m = 1, \dots, M$, used for block merging (right). The blocks illustrate the content of the elementary files f_s^{el} owned by the computational processes. Repetition of these blocks along the x - and y -directions allows the computational domain to be extended to an arbitrary size



$$\begin{aligned} & \int_{\hat{F}} g^-(X_{K^-}(\hat{\mathbf{x}}^-)) f^+(X_{K^+}(\hat{\mathbf{x}}^+)) |J_{X_{K^+}}(\hat{\mathbf{x}}^+)| d\hat{\sigma} \\ & \approx \sum_{q=1}^{N_q} g^-(X_{K^-}(\hat{\mathbf{x}}_q^-)) f^+(X_{K^+}(\hat{\mathbf{x}}_q^+)) |J_{X_{K^+}}(\hat{\mathbf{x}}_q^+)| w_q \\ & = \sum_{q=1}^{N_q} \hat{g}^-(\hat{\mathbf{x}}_q^-) \hat{f}^+(\hat{\mathbf{x}}_q^+) |J_{X_{K^+}}(\hat{\mathbf{x}}_q^+)| w_q, \end{aligned}$$

where w_q , for $q = 1, \dots, N_q$ are suitable weights and since $g^-(\mathbf{x}_q) = \hat{g}^-(\hat{\mathbf{x}}_q^-)$ and $f^+(\mathbf{x}_q) = \hat{f}^+(\hat{\mathbf{x}}_q^+)$.

In our process, we first define $\mathbf{x}_q \in F$ as the image of quadrature nodes $\hat{\mathbf{x}}_q \in \hat{F}$ through X_{K^+} , i.e., $\mathbf{x}_q = X_{K^+}(\hat{\mathbf{x}}_q^+)$, and then, we map them back by using X_{K^-} to find $\hat{\mathbf{x}}_q^- \in \hat{F}$, i.e., $\hat{\mathbf{x}}_q^- = X_{K^-}^{-1}(\mathbf{x}_q)$, cf. Fig. 8. So, in practice, we implicitly define F and $\hat{F}$ through quadrature points.

To accelerate the aforementioned computation, information regarding element faces ∂K^+ and ∂K^- in $\mathcal{F}^I$ is stored in two separate files. A third file contains the coordinates of a single vertex of each face ∂K^- belonging to $\mathcal{T}_{\text{out}}$. The algorithm for face-matching proceeds as follows: the finer faces ∂K^+ loop over their quadrature nodes $\{\mathbf{x}_q\}_{q=1}^{N_q}$, and look for the coarser face containing the current node $\mathbf{x}_j$.

In this scenario, where hundreds of thousands of faces appears, to narrow the number of faces among the preselected ones, we first consider all the coarse faces ∂K^- whose normal $\mathbf{n}^-$ is opposite, up to a tolerance, to that of the current fine face $\mathbf{n}^+$. Hence, the search for the correct neighboring element is reduced to a smaller set of faces. Then, among the latter, we consider the ones whose vertex V_{K^-} lies within a circular neighborhood $C_r(\mathbf{x}_j)$ of the quadrature node of the finer face, with radius $r = \max\{\text{diam}(\partial K^-), \text{diam}(\partial K^+)\} + \varepsilon$, $\varepsilon > 0$. If those conditions are satisfied, the face is selected among the possible ones intersecting ∂K^+ .

Finally, to properly find the face ∂K^- sharing the node $\mathbf{x}_j$ with ∂K^+ , a Newton–Raphson routine checks whether the quadrature point lies within the element K^- , computing its reference position in $[-1, 1]^3$. This procedure is shown in Algorithm 3, and Fig. 9 provides a visual explanation.

Algorithm 3 Face matching procedure


```

1: Fix tol > 0
2: for  $\partial K^+$  s.t.  $\partial K^+ \cap \Gamma_I \neq \emptyset$  do
3:   for  $\mathbf{x}_q$  with  $q = 1, \dots, N_q$  do
4:     for  $\partial K^-$  s.t.  $\partial K^- \cap \Gamma_I \neq \emptyset$  do
5:       if  $|\mathbf{n}_+ + \mathbf{n}_-| < \text{tol}$  and  $V_{K^-} \in C_r(\mathbf{x}_j)$  then
6:         Compute reference position  $\hat{\mathbf{x}}_j$  with the Newton-
           Raphson method
7:         if  $\hat{\mathbf{x}}_j \in [-1, 1]^3$  then
8:            $K^-$  is a neighbour of  $K^+$ 
9:         end if
10:      end if
11:    end for
12:  end for
13: end for
    
```

6 Numerical results

This section presents the numerical results obtained from the PMUT array simulations. We first consider the case of a single PMUT configuration, which serves as a reference case for the proposed computational model. Once the model's consistency is established, we investigate more complex scenarios involving multi-element PMUT arrays. The simulations are performed under two distinct operational conditions: a pure transmission (TX) framework, where the focus is on the emitted acoustic field, and a transmission–reception (TX–RX) framework, which additionally assesses the PMUTs sensing and reception capabilities.

Fig. 8 The figure shows the physical hexahedra K^- and K^+ , sharing the common face F , together with their quadrature nodes $\mathbf{x}_q$, and the corresponding reference hexahedra $\hat{K}^\pm$ and their reference nodes $\hat{\mathbf{x}}_q^\pm$



6.1 Single PMUT excitation in TX phase

The first analysis focuses on a single PMUT, i.e. $N_p = 1$ in (2), operating in transmission mode. An alternating low-voltage signal of the form $\phi(t) = A \sin(2\pi f t) g(t)$ is applied to the surface of the PMUT, which is free to vibrate over the upper boundary of the computational domain. The modulation function $g(t)$ is introduced to ensure a smooth onset and offset of the excitation signal, thereby minimizing spurious frequency components and numerical artifacts associated with abrupt signal imposition. The function is defined as follows:

$$g(t) = \exp(1 - \alpha(T_d - t)) \mathbb{1}_{\{0 \leq t < T_d/2\}} + \exp(1 - \alpha(t)) \mathbb{1}_{\{T_d/2 \leq t \leq T_d\}}, \quad (15)$$

where $\alpha(\tau) = \left(1 - \frac{2\beta(\tau)^{1+2n}}{1 + \beta(\tau)^{2n}}\right)^{-2}$, with $\beta(\tau) = (2\tau/T_d - 1)$ and $n = 10$. The excitation parameters are set to an amplitude of $A = 1.0$ V, a driving frequency of $f = 2.5$ MHz, and a total duration of $T_d = 5$ μ s. The imposed voltage is plotted in Fig. 11 (left). The computational domain comprises the PMUT surrounded by an acoustic medium for wave propagation, which extends for 2000 μ m both in the radial and z directions. The radius of the PMUT is 65 μ m. To avoid spurious reflections, first-order absorbing boundary conditions are imposed along the lateral and bottom boundaries of the acoustic domain, while a hard-wall constraint is imposed at the top surface supporting the PMUT as in (1), cf. Fig. 10. A time step $\Delta t = 2 \times 10^{-4}$ μ s and a second-order polynomial degree in space are selected. The final time is set to $T = 12$ μ s. Water is employed for both the inner and outer domains, hence the material density and




Fig. 9 Example of the matching algorithm of DG faces. The red and blue grids refer to the coarse and fine faces, respectively, and the grey dots are the Gauss-Legendre quadrature nodes of the finer faces. The blue-filled face is a representative fine face, and the center $\mathbf{x}_j$ of the circle is its quadrature node considered; the red-filled faces are those selected by the algorithm. Notice that considering only one vertex of the coarse face instead of all four increases the number of residual faces, but it avoids excessive RAM consumption. (color figure online)

the velocity of acoustic waves are $\rho = 1 \times 10^{-3}$ kg m $^{-3}$ and $c = 1481$ m s $^{-1}$, respectively. The first three axially symmetric modes, i.e., $N_m = 3$ in (2), are considered.

Figure 11 illustrates the temporal profile of the smooth input voltage applied to the PMUT surface, together with the resulting pressure excitation at the center of the membrane. The results are compared with a COMSOL Multiphysics V6.3 simulation of the PMUT array, where the device is modeled using a hybrid solid-shell approach: thin passive layers are represented using 2D first-order shear deforma-




Fig. 10 Test case of Sect. 6.1. Mesh configuration used for the single PMUT simulation. Panels a–c show the lateral, top, and zoomed top views, respectively. The PMUT, highlighted in panel (c), is excited by an input voltage signal applied to its upper surface




Fig. 11 Test case of Sect. 6.1. Time evolution of the input signal imposed on the PMUT surface (left), and output pressure computed at the PMUT center (right). SPEED and COMSOL signals are in blue and red, respectively. (color figure online)

tion theory shell elements, while the PZT layer and the silicon substrate are retained as 3D solid elements. Displacement continuity is enforced across shell–solid interfaces. The front encapsulation is modeled as an acoustic fluid with impedance-matched properties, and the surrounding acoustic domain solves the linearized acoustic wave equation [2]. This configuration provides an accurate approximation of a full 3D model while reducing computational cost, and the comparison demonstrates close correspondence between the two simulation frameworks. It is important to emphasize that this single-element comparison is intended solely as an accuracy validation benchmark. For isolated or small-scale configurations, COMSOL’s specialized local solver remains highly optimized and faster. However, the primary advantage of the proposed approach lies in its high scalability within parallelized computing environments. While large-scale PMUT array simulations quickly become computationally unfeasible in COMSOL due to memory limits and scaling bottlenecks, SPEED leverages parallel architectures to efficiently handle large-scale, multi-element configura-

tions that are otherwise out of reach (cf. Sects. 6.3 and 6.4). Figure 12, instead, provides a visualization of the acoustic wavefield propagation on an x – y plane extracted at a depth of $z = -1000\ \mu\text{m}$, together with the temporal evolution of the pressure at $P = (0\ \mu\text{m}, 0\ \mu\text{m}, -1000\ \mu\text{m})$. The comparison with COMSOL further confirms the consistency and agreement between the two simulations.

From a physical standpoint, the output pressure waveform at the center of the membrane firmly tracks the electrical excitation during the initial cycles, demonstrating efficient electromechanical coupling. However, as time progresses, structural inertia and the acoustic radiation impedance of the surrounding fluid medium introduce a minor phase lag and a smooth decaying tail once the electrical excitation ceases. The spatial wavefield slice in Fig. 12 clearly illustrates the transition from a highly localized near-field pressure fluctuation directly below the membrane into a classical spherical acoustic wave as it propagates into the bulk fluid, validating the model’s ability to capture diffraction limits.




Fig. 12 Test case of Sect. 6.1. Snapshot at $t = 4 \mu\text{s}$ of a slice in the x - y plane (left), and output pressure computed at the $P = (0 \mu\text{m}, 0 \mu\text{m}, -1000 \mu\text{m})$ for SPEED and COMSOL (right)




Fig. 13 Test case of Sect. 6.1. Three dimensional surface plots showing the time evolution of the PMUT membrane displacement field $U_1(\mathbf{x}, t)$ in μm at three time instants: $t_1 = 0.0 \mu\text{s}$, $t_2 = 1.0 \mu\text{s}$, $t_3 = 4.0 \mu\text{s}$. The

plots show neutral (flat), positive (upward), and negative (downward) deflections, respectively. The maximum amplitude of the displacement is $0.0176 \mu\text{m}$

To illustrate the dynamic behavior of the structural response, Fig. 13 presents the reconstructed displacement field $U_1(\mathbf{x}, t)$ of the PMUT membrane across three distinct time instants. At $t = 0.0 \mu\text{s}$, the membrane resides in its undeformed state, aligned entirely at the zero-plane. As time progresses, the spatial evolution driven by the modal coordinates $q_{1,m}(t)$ becomes distinct: at $t = 1.0 \mu\text{s}$, the membrane exhibits a uniform upward deflection ($U_1 \geq 0$), while at $t = 4.0 \mu\text{s}$, it transitions into a completely downward displacement phase ($U_1 \leq 0$). The visual transition between the positive and negative operational regimes is distinguished across the three subplots, highlighting the fundamental mode shape profiles scaling over time.

To validate the numerical stability of the proposed time-stepping scheme and determine the operational threshold of the CFL condition in Eq. (13), a parametric analysis was conducted. Figure 14 illustrates the temporal evolution of the acoustic pressure for different K_{CFL} values. The results demonstrate that the numerical simulation remains perfectly stable, maintaining a well-bounded oscillatory behavior without numerical dissipation, for values up to

$K_{\text{CFL}} = 0.20526$ (represented by the grey curve). However, a strict stability threshold is encountered immediately thereafter: at $K_{\text{CFL}} = 0.20527$ (indicated by the magenta curve), high-frequency numerical oscillations rapidly emerge and grow exponentially around $t \approx 6.5 \mu\text{s}$. This sharp transition into numerical instability clearly identifies the critical CFL limit for our explicit integration scheme, confirming that the time step must be strictly bounded below this threshold to ensure a physically meaningful and stable solution. In our simulations, we enforce a time step $\Delta t = 2 \times 10^{-4} \mu\text{s}$, i.e., $K_{\text{CFL}} = 0.13164$, to be conservative.

6.2 Single PMUT with an obstacle in TX-RX phase

In general, PMUTs are designed to be used in both the transmission and the receiving phases, i.e., to trigger the propagation of the acoustic signal up to a specific target and to capture the response that they receive back from the target itself. To this aim, the following simulation deals with both TX and RX phases.





Fig. 14 Test case of Sect. 6.1. Stability analysis under varying K_{CFL} . The simulation remains stable up to the optimal value $K_{CFL} = 0.20526$ (grey curve), while a marginal increase to $K_{CFL} = 0.20527$ (magenta curve) triggers exponential numerical instability near the end of the simulation timeline. Curves are vertically offset for visual clarity

An obstacle representing the target is put inside the acoustic propagation domain as shown in Fig. 15. The setup is the following: an AC low-voltage $\phi(t) = A \sin(2\pi ft)h(t)$ is used with

$$h(t) = \exp(1 - \alpha(T_d - t)) \mathbb{1}_{\{0 \leq t \leq T_d\}},$$

where $\alpha(\tau)$ is defined as in the test case of Sect. 6.1, with voltage parameters $A = 2.0$ V, $f = 2.5$ MHz, and a total duration $T_d = 1$ μ s. The acoustic water domain is a cylinder of radius $r = 2000$ μ m composed of three layers of height $h_1 = 3c$, $h_2 = 25$ μ m, and $h_3 = 1000$ μ m, respectively, where h_1 is exactly the distance traveled by the acoustic wave in a time equivalent to $3T_d$ and $c = 1481$ m s⁻¹ is the wave speed in water. In this way, the acoustic pressure wave fully reaches the obstacle with radius $r_{\text{obs}} = 1700$ μ m placed at a distance h_1 from the PMUT center, and the pressure is received back from the membrane after a time equivalent to $6T_d$ (see Fig. 16). From time T_d up to the final time $T = 12$ μ s, the PMUT is left in open-circuit condition with no externally applied voltage. The resulting mechanical vibrations and reflected pressure waves induce a voltage across

Fig. 15 Test case of Sect. 6.2. Geometry used for the single PMUT simulation with an obstacle in TX-RX phases. Panels a–c show the geometry employed, the thickness of the cylinder layers, and a zoom of a representative obstacle for graphical reasons, respectively



the electrodes via the direct piezoelectric effect. Absorbing boundary conditions are imposed on the lateral and bottom sides of the outer cylinder, while a homogeneous Neumann condition is enforced on the top side and on the obstacle's surface.

Physically, this dual-phase setup mimics an authentic pulse-echo ultrasonic sensing cycle, tracking both the voltage and the corresponding pressure computed at the PMUT center, as shown in Fig. 16. The open-circuit condition imposed after the pulse duration T_d alters the PMUT behavior, allowing the membrane to transition from an active actuator to a passive sensor in the receiving phase. From T_d up to $t = 6T_d$, the voltage amplitude decays progressively as the pressure on the PMUT surface reduces. The subsequent delay of $6T_d$ corresponds exactly to the bidirectional acoustic time-of-flight through the water layers to the obstacle and back, i.e., $2h_1/c$. Consequently, the minor increase in values captured after $t > 6T_d$ in both the voltage and pressure plots, cf. Fig. 16, provides evidence of the direct piezoelectric effect, where the reflected acoustic wavefront physically deforms the mechanical layers of the PMUT, transforming structural strain back into a readable electrical potential. Finally, the reduction in the reflected peak amplitude relative to the transmitted wave characterizes the geometric spreading losses and acoustic energy scattering caused by the finite radius of the cylinder obstacle.

Figure 17 presents two snapshots of the iso-contours of the acoustic pressure; the one on the left is a snapshot at time $t = 3T_d$, where the waves move towards the obstacle represented in green, while the picture on the right highlights the pressure waves that reach the PMUT surface at time $6T_d$ after the reflection.

Figure 18 illustrates the reconstructed displacement field $U_1(\mathbf{x}, t)$ of the PMUT membrane in the TX-RX configuration at three operational stages. At $t_1 = 1.0$ μ s, the membrane reaches a prominent upward deflection corresponding to the minimum pressure value of the initial excitation. At $t_2 = 6.0$ μ s, the structural state transitions into a downward deflection phase precisely as the back-reflected acoustic signal is received. Subsequently, at $t_3 = 10.2$ μ s, the mem-





Fig. 16 Test case of Sect. 6.2. Time evolution of the input signal imposed on the PMUT surface (left) with a zoom on the region for $t \in [5, 10]$ μs , and output pressure computed at the PMUT center (right). The time instant $t = 6T_d$ is highlighted with a dashed line

Fig. 17 Test case of Sect. 6.2. Snapshots of the acoustic pressure waves contour plot: on the left, the one hitting the obstacle (in green), on the right, the one that travels back up to the PMUT



brane exhibits a secondary positive deflection representing the point where the returned signal undergoes attenuation ($\sim 1 \times 10^{-4} \mu\text{m}$). This comparison directly underscores the multi-scale amplitude evolution inherent to the complete transmission-reception cycle.

6.3 Multi-array PMUT simulation

In general, absorbing boundary conditions are well-suited for simulations where wave propagation is predominantly orthogonal to the boundary. This approach is particularly effective when the computational domain is spherical or cylindrical, namely, when the shape of the boundary matches the shape of the propagating waves (cf. Sects. 6.1 and 6.2). However, classical absorbing conditions can lead to undesirable internal reflections in domains with corners. Figure 19 illustrates the influence of the PML on a single-element array within a computational domain featuring corners. The two images correspond to the same simulation performed with and without Ω_{PML} . In the absence of this layer, wave reflections arising at the domain boundaries lead to noticeable

distortions in the wavefront, thereby compromising the accuracy of the numerical solution. Conversely, when the PML is applied in combination with standard absorbing boundary conditions, spurious reflections are effectively suppressed, allowing for a smoother and more physically consistent propagation of acoustic waves within the domain.

Multi-array simulations involve a large number of membranes. This section focuses on the simulation of the 64-element array, cf. Fig. 21. Each PMUT array consists of 184 membranes arranged in two adjacent columns of 92 each, distributed according to the configuration illustrated in Fig. 6. This configuration and the physical parameters used in the sequel have been taken from [26, 27].

The central idea behind this type of simulation, regardless of the number of arrays employed, is to define an acoustic propagation domain with a Cartesian geometry rather than a spherical or cylindrical one. Using Cartesian coordinates significantly reduces the total number of degrees of freedom required for the mesh compared to alternative geometries. Hard-wall conditions are imposed on the top




Fig. 18 Test case of Sect. 6.2. Three-dimensional surface plots showing the time evolution of the PMUT membrane displacement field $U_1(\mathbf{x}, t)$ in μm for the TX-RX configuration at three time instants: $t_1 = 1.0 \mu\text{s}$, $t_2 = 6.0 \mu\text{s}$, $t_3 = 10.2 \mu\text{s}$




Fig. 19 Test case of Sect. 6.3. Schematic representation of a slice of the acoustic domain $\Omega_{\text{in}} \cup \Omega_{\text{out}}$ and the surrounding perfectly matched layer region Ω_{PML} . The upper figure corresponds to the simulation without the PML, while the lower figure shows the solution obtained when the PML is included. The dashed line represents the interface Γ_{PML}

surface Γ_N . A Perfectly Matched Layer encloses the computational domain, with absorbing boundary conditions applied along the outer boundary Γ_{ABC} , see Eq. (1). Each array operates independently; PMUTs in different arrays may be

driven by distinct voltage excitations. In the following simulation, the applied voltage signal follows the expression in Eq. (15), with parameters $f = 2.5 \text{ MHz}$ and $A = 7.0 \text{ V}$, over a total activation time of $T_d = 4 \mu\text{s}$. All arrays simultaneously begin and end the transmission phase. The domain consists of two regions described in Sect. 4, i.e., an inner and an outer domain. In the inner region, the material properties are defined as $\rho_{\text{in}} = 1140 \text{ kg m}^{-3}$ and $c_{\text{in}} = 1130 \text{ m s}^{-1}$. The outer region corresponds to water, characterized by $\rho_{\text{out}} = 1000 \text{ kg m}^{-3}$ and $c_{\text{out}} = 1481 \text{ m s}^{-1}$. The inner domain has a thickness of $h_{\text{in}} = 500 \mu\text{m}$. The perfectly matched layer extends over a distance equal to one wavelength in water, $\lambda_w = 592.4 \mu\text{m}$, and the overall dimensions of the domain are $l_x = 22874.4 \mu\text{m}$, $l_y = 15570.64 \mu\text{m}$, and $l_z = 12940.4 \mu\text{m}$ along the x , y , and z directions, respectively. The radius of the PMUT is $r_p = 65 \mu\text{m}$, and the distance between two adjacent PMUT centers, known as pitch, is $d_p = 150 \mu\text{m}$. Figure 21 contains the geometrical information of the simulation. The time step of the simulation is $\Delta t = 2 \times 10^{-4} \mu\text{s}$, and a second-order polynomial degree in space is considered. The final time is set to $T = 15 \mu\text{s}$. The number of elements in the mesh is 33.8M, and the degrees of freedom are 262.6M. Figure 22 illustrates the membrane excitation after $0.2 \mu\text{s}$ and




Fig. 20 Test case of Sect. 6.3. Three-dimensional surface plots showing the spatial evolution of the array-centered PMUTs displacement field $U_k(\mathbf{x}, t)$ in μm at two time instants: $t_1 = 0.4 \mu\text{s}$, $t_2 = 10.4 \mu\text{s}$

the resulting acoustic wave propagation within the computational domain along the three spatial directions. The pressure field is accurately captured within the computational domain, as shown by the alternating pattern of positive and negative planar pressure waves. Ten probes placed at $P = (0, 0, z)$, with $z \in \{-500 \mu\text{m}, \dots, -9500 \mu\text{m}\}$, allow for evaluating the acoustic pressure signal in both the near and far fields (see Fig. 23).

The physical implications of scaling up to a 64-element array are demonstrated in the directional characteristics of the acoustic field. Unlike the isolated spherical wavefronts produced by a single PMUT, the closely spaced multi-element array triggers constructive spatial interference. This originates a highly directional planar acoustic beam traveling along the z -axis, which is essential for medical imaging applications requiring high lateral resolution. This phenomenon is further highlighted by the probe data in Fig. 23: the probes nearest to the array ($z = -500 \mu\text{m}$ to $-3500 \mu\text{m}$) capture the complex and fluctuating pressure variations characteristic of the acoustic near-field, where individual element wavefronts are still interfering. Conversely, at deeper probes ($z > -5500 \mu\text{m}$), the signals smooth out into a stable profile, marking the transition into the acoustic far-field dominated by steady beam directional propagation and geometric attenuation.

To capture the collective array dynamics, Fig. 20 displays the reconstructed multi-element displacement field $U_k(\mathbf{x}, t)$, with $k \in \{1, \dots, 16\}$, across 16 PMUTs in the central elements of the array at two different time instants: $t_1 = 0.4 \mu\text{s}$ and $t_2 = 10.4 \mu\text{s}$. The plots track the synchronized evolution of the elements in both an upward and a downward phase. Notably, while the fully excited elements display synchronized, uniform displacement at t_1 (left), the unexcited state at t_2 reveals minor acoustic crosstalk effects visible as residual, low-amplitude wave propagation across the array (right).

The simulation ran on a high-performance cluster consisting of 40 CPU compute nodes. Each CPU node is configured with two AMD EPYC 7413 processors, each featuring 24 cores, yielding a total of 96 threads per node. These nodes are provisioned with 512 GB of RAM each, providing substantial memory capacity to support large-scale simulations and data-intensive computations within a parallel computing framework. By utilizing 21 nodes, and 1935 processes, the simulation required approximately 12 h to complete.

6.4 Performance comparisons and scalability test

A comparative performance analysis was carried out to evaluate the effectiveness of the DG interface optimization described in Sect. 5 and to quantify the computational and communication overhead associated with the evaluation of the integrals over Γ_{PMUT} in Eqs. (11) and (12), as discussed in Sect. 4. The first test case employed a computational domain





Fig. 21 Test case of Sect. 6.3. Representation of the three-dimensional view of the geometry employed for the 64-elements simulation. The dots represent the probes at which the solution is computed. Notice that the figure is not scaled

consisting of a 20-element array distributed across 315 processes. Two configurations were investigated: one without DG interface optimization and one with the optimization enabled. In the absence of DG optimization, the setup of 318698 interfaces required approximately 160 000 s. By contrast, enabling DG interface optimization reduced the setup time to approximately 630 s. This represents a reduction of more than two orders of magnitude in interface setup time, underscoring the critical importance of DG optimization for improving scalability and significantly reducing preprocessing overhead. The second test case was performed on a computational domain comprising a single-element array partitioned across 96 processes. In the first configuration, an arbitrary mesh partitioning was employed, with no constraint enforcing the association of an entire PMUT with a single process. Under this configuration, the solution of Eq. (10) incurred an inter-process communication cost of approximately 0.167 s per timestep, arising solely from the exchange of local contributions to the PMUT's solution on each process with all other processes. This communication overhead substantially exceeded the pure computational cost, thereby indicating poor parallel efficiency.

In the second configuration, the tailored partitioning strategy described in Sect. 4 was adopted. This approach resulted in an average synchronization waiting time of approximately 0.0008 s per PMUT per timestep for each process, corresponding to a reduction of more than two orders of magnitude. It is worth noting that, in the former configuration, increasing either the number of membranes or the number of processes leads to a corresponding increase in communication time, whereas in the latter configuration, the communication cost remains, on average, constant. This

Fig. 22 Test case of Sect. 6.3. The figure shows the snapshot of the surface containing the PMUTs at time $t = 0.2 \mu\text{s}$ with a zoom of the middle region (top), and the solution at time $t = 4 \mu\text{s}$ depicted over 3 slices (bottom)



behavior is further illustrated by the scalability analysis presented in the following.

Overall, these results demonstrate that DG interface optimization is essential for reducing preprocessing time in large-scale simulations, while PMUTs performance is highly sensitive to communication patterns and domain decomposition strategies. Inappropriate partitioning can result in communication-dominated runtimes, thereby severely limiting parallel efficiency. These findings underscore the necessity of balanced mesh partitioning and efficient data exchange mechanisms to achieve scalable and efficient parallel performance.

Scalability refers to the ability of a computational system to maintain or improve performance as workload and resources increase. In high-performance and parallel computing, it depends on the efficient use of processors, memory, and communication, requiring well-designed parallel algorithms, effective memory management, and controlled communication overhead to ensure competitive performance. In the specific context of the problem under investigation, the main constraints that influence scalability are the ones dis-




Fig. 23 Test case of Sect. 6.3. The figure shows the acoustic pressure computed at ten equispaced points $P = (0, 0, z)$, with $z \in \{-500 \mu\text{m}, \dots, -9500 \mu\text{m}\}$

cussed in Sect. 4 and can be categorized into three main factors:

1. PMUTs blocks: the scalability of the system is directly affected by the configuration and number of PMUT array




Fig. 24 Test case of Sect. 6.4. The figure shows the theoretical trend (black) and the real one (blue) of the strong scalability test conducted across four different simulations (left). The value on the y-axis is nor-

malized with respect to the time t_0 of the first simulation. On the right, the weak scalability test performed on five different simulations (color figure online)

blocks. As the array size increases, the computational load associated with the membrane becomes dominant. Efficient data partitioning and optimized communication between array elements are therefore essential to maintain scalability in large configurations.

2. Acoustic propagation domain: the computational complexity of simulating the acoustic propagation domain scales with both the spatial resolution and the physical model. High-resolution domains require finer discretization, leading to increased memory usage and processing demands. Load balancing and domain decomposition strategies are of paramount importance to ensure that the computations remain scalable across multiple processes. Moreover, simulations with the PML enlarge the DOFs and hence require an additional complexity to be taken into account.
3. DG interfaces: they represent another critical scalability constraint. These interfaces govern the communication and numerical flux exchanges between subdomains. As the number of hexahedral elements increases, so does the volume of interfacial data transfer, potentially leading to communication bottlenecks. Maintaining scalability in this context requires efficient inter-process communication schemes and optimized data structures that ease rapid access and exchange of boundary information.

Overcoming these challenges requires a combination of algorithmic optimization, advanced parallelization strategies, and system design to ensure that performance scales effectively with increasing computational resources. A strong scalability test is performed to assess the validity of the approach discussed in the previous sections. Four simulations are conducted within the same computational domain, exploiting the cluster of the Mathematical Department of Politecnico

di Milano (DMAT), whose characteristics are reported in the previous section. The number of processes selected is 22, 44, 88 and 176, respectively. In this way, communication and load balancing are performed both among processes within a single node and between different nodes. Figure 24 (left) shows the comparison between the theoretical and normalized real-time resources required in the four cases. The plot in logarithmic scale shows a very good trend compared to the expected behavior. Moreover, a weak scalability test was assessed over five different simulations proportionally increasing both the number of degrees of freedom and the number of processes, thereby preserving a fixed ratio of approximately 140000 DOFs per process. Figure 24 (right) shows an almost ideal weak scalability, by holding the execution timestep constant.

7 Conclusions and future developments

This work presented a comprehensive methodology for the efficient numerical simulation of large-scale PMUTs arrays, based on a hybrid approach consisting of model order reduction techniques and a high-performance Discontinuous Galerkin Spectral Element framework implemented within the SPectral Elements in Elastodynamics with Discontinuous Galerkin (SPEED) environment. By integrating the mechanical modal behavior of individual PMUTs into an acoustic domain discretized through high-order elements, the proposed approach achieves a balance between accuracy and computational efficiency. The resulting model successfully captures both transmission and reception phenomena, enabling scalable and high-fidelity simulations of complex array configurations. The implemented domain decomposition and optimization strategies proved effective in enhancing computational performance, ensuring that the framework

remains suitable for massively parallel simulations involving thousands of transducers. The numerical results confirm the robustness and scalability of the proposed approach, making it a promising tool for the design and optimization of next-generation PMUT-based ultrasonic systems. Beyond the present developments, future work will focus on extending the current framework to account for thermoacoustic effects, which are increasingly relevant as PMUTs operate at higher power densities and within miniaturized environments. Incorporating thermo-mechanical coupling and heat generation due to dielectric and mechanical losses will enable the study of temperature-induced frequency shifts, performance degradation, and thermal noise in PMUT arrays. Moreover, the integration of nonlinear thermoelastic models could provide deeper insight into self-heating phenomena and their influence on both emission and reception behavior. Such extensions would pave the way toward multiphysics simulations capable of predicting the full electromechanical–acoustic–thermal response of PMUT systems under realistic operational conditions. Ultimately, these developments will support the design of more reliable, energy-efficient, and thermally stable ultrasonic microsystems, broadening their applicability in demanding environments such as biomedical diagnostics, industrial sensing, and harsh-condition monitoring.

Author contributions PFA: Supervision; Writing—editing. OMOA: Validation; COMSOL simulations. MGG: Conceptualization; Formal analysis; Investigation; Methodology; SPEED coding; Validation; Writing—original draft and editing. IM: Formal analysis; Methodology; Writing—original draft and editing; SPEED coding. NP: Supervision; Writing—editing.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement. This work is part of the *PROgramma per la realizzazione di dispositivi ad Ultrasuoni e loro promozione Democratica* (PROUD) project funded by the Italian Ministry of Enterprises and Made in Italy (MIMIT): Prog. n.: F/310211/01-05/X56 (Decree n.4187, 19-12-2023). The authors would like to thank STMicromicroelectronics for their valuable support and contributions within the framework of the PROUD project. This work also benefited from helpful discussions and technical support from all project partners. P.F. Antonietti, M.G. Garroni, I. Mazzieri, and N. Parolini are members of INdAM-GNCS group. The present research is part of the activities of “Dipartimento di Eccellenza 2023–2027”. The simulations discussed in this work were, in part, performed on the HPC Cluster of the Department of Mathematics of Politecnico di Milano, which was funded by MUR grant “Dipartimento di Eccellenza 2023–2027”.

Data availability The datasets generated during the current study are available from Michelangelo Gabriele Garroni upon reasonable request.

Declarations

Conflict of interest The author declares no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adap-

tation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abdalla OMO, Garroni MG, Moretti E, Massimino G, Corigliano A, Frangi A et al (2025) Efficient modeling and simulation of large PMUT arrays for biomedical applications. In: Proceedings of the 2025 26th international conference on thermal, mechanical and multi-physics simulation and experiments in microelectronics and microsystems (EuroSimE 2025). Institute of Electrical and Electronics Engineers Inc, pp 1–5
2. Abdalla OMO, Massimino G, Savoia AS, Quaglia F, Corigliano A (2022) Efficient modeling and simulation of PMUT arrays in various ambients. *Micromachines* 13(6):962
3. Antonietti PF, Mazzieri I, Quarteroni A, Rapetti F (2012) Non-conforming high order approximations of the elastodynamics equation. *Comput Methods Appl Mech Eng* 209–212:212–238
4. Antonietti PF, Bonaldi F, Mazzieri I (2020) Simulation of three-dimensional elastoacoustic wave propagation based on a Discontinuous Galerkin spectral element method. *Int J Numer Methods Eng* 121(10):2206–2226
5. Artoni A, Antonietti PF, Mazzieri I, Parolini N, Rocchi D (2024) A hybrid Finite Volume-Spectral Element Method for aeroacoustic problems. *Comput Math Appl* 155:193–208
6. Baffet D, Grote M, Imperiale S, Kachanovska M (2019) Energy decay and stability of a perfectly matched layer for the wave equation. *J Sci Comput* 81(3):2237–2270
7. Birjis Y, Swaminathan S, Nazemi H, Raj GCA, Munirathnam P, Abu-Libdeh A, Emadi A (2022) Piezoelectric micromachined ultrasonic transducers (PMUTs): performance metrics, advancements, and applications. *Sensors (Basel)* 22(23):9151
8. Bybi A, Khouili D, Granger C, Garoum M, Mzerd A, Hladky A (2020) Experimental characterization of a piezoelectric transducer array taking into account crosstalk phenomenon. *Int J Eng Technol Innov* 10(4):259–268
9. Chaudhary S, Singh J, Haque FZ (2022) BVD and Mason's modelling of piezoelectric bulk acoustic resonators for high frequency applications. *Mater Today Proc* 65:2556–2559
10. Feng G-H, Lin H-J (2019) Piezoelectric micromachined ultrasonic transducers with a cost-effective bottom-up fabrication scheme for millimeter-scale range finding. *Sensors* 19(21):4696
11. Gao Y, Wang X, Zhao L et al (2025) High frame rate ultrasonic flowmeter based on PMUT array with bidirectional acoustic beams. *Microsyst Nanoeng* 11:245
12. Grote MJ, Sim I (2010) Efficient PML for the wave equation. <https://arxiv.org/abs/1001.0319>
13. He L, Xu W, Wang Y, Zhou J, Ren J (2022) Sensitivity-bandwidth optimization of PMUT with acoustical matching using finite element method sensors 22(6):2307
14. He Y, Wan H, Jiang X, Peng C (2022) Piezoelectric micromachined ultrasound transducer technology: recent advances and applications. *Biosensors (Basel)* 13(1):55
15. Kaltenbacher B, Kaltenbacher M, Sim I (2013) A modified and stable version of a perfectly matched layer technique for the 3-d

- second order wave equation in time domain with an application to aeroacoustics. *J Comput Phys* 235:407–422
16. Luo Z, Cai J, Zhang S, Gu Y, Lou L, Wu T (2022) Nonlinearity of piezoelectric micromachined ultrasonic transducer using AlN thin film. *IEEE Open J Ultrason Ferroelectr Freq Control* 2:96–104
 17. Massimino G, Quaglia F, Corigliano A, Frangi A (2021) Model order reduction for the analysis of large arrays of piezoelectric micromachined ultrasonic transducers in water. *Appl Acoust* 182:108231
 18. Mazzieri I (2012) Non-conforming high order methods for the elastodynamics equation. PhD thesis, POLITesi digital repository, Politecnico di Milano
 19. METIS. <https://github.com/karypislab/metis>
 20. Narváez F, Hosseini S, Farkhani H, Moradi F (2021) Analysis and design of a PMUT-based transducer for powering brain implants. In: Proceedings of the 2021 IEEE 19th international new circuits and systems conference (NEWCAS), pp 1–4
 21. Nayak B, Gupta H, Roy K, Ashok A, Shastri V, Pratap R (2021) An experimental study of the acoustic field of a single-cell piezoelectric micromachined ultrasound transducer (PMUT). arXiv preprint. [arXiv:2108.04660](https://arxiv.org/abs/2108.04660)
 22. Pan J, Bai C, Zheng Q, Xie H (2023) Review of piezoelectric micromachined ultrasonic transducers for rangefinders. *Micromachines* 14(2):374
 23. Raviart P-A, Thomas JM (1983) Introduction à l'analyse numérique des équations aux dérivées partielles. Collection Mathématiques Appliquées pour la Maîtrise. Masson, Paris
 24. Roy K, Lee JE, Lee C (2023) Thin-film PMUTs: a review of over 40 years of research. *Microsyst Nanoeng* 9:95
 25. Sanvito S, Passoni M, Giusti D, Terenzi M, Prelini C, La Mura M, Savoia A (2024) PMUT-based system for continuous monitoring of bolted joints preload. *Sensors* 24:4150
 26. Savoia AS, Mazzanti A, Ottaviani S, Novaresi L, Malcovati P, Ghisu DU, Bonizzoni E, Terenzi M, Quaglia F (2022) A 4-channel fully integrated ultrasound imaging front-end transceiver for 1-D PMUT arrays. In: IEEE international ultrasonics symposium (IUS), pp 1–4
 27. Savoia AS, Casavola M, Boni E, Ferrera M, Prelini C, Tortoli P, Giusti D, Quaglia F (2021) Design, fabrication, characterization, and system integration of a 1-D PMUT array for medical ultrasound imaging. In: IEEE international ultrasonics symposium (IUS), pp 1–3
 28. Shelton S, Chan M, Park H, Horsley D, Boser B, Izyumin I, Przybyla R, Frey T, Judy M, Nunan K, Sammoura F, Yang K (2009) CMOS-compatible AlN piezoelectric micromachined ultrasonic transducers. In: 2009 IEEE International ultrasonics symposium, pp 402–405
 29. Smyth K, Kim S (2015) Experiment and simulation validated analytical equivalent circuit model for piezoelectric micromachined ultrasonic transducers. *IEEE Trans Ultrason Ferroelectr Freq Control* 62(4):744–765
 30. Su X, Ren X, Wan H, Jiang X, Liu X (2022) Simulation and optimization of piezoelectric micromachined ultrasonic transducer unit based on AlN. *Electronics* 11:2915
 31. Xu T, Zhang Z, Rocha R, Zeng L, Xu C (2024) Theoretical validation of a single-channel air-coupled PMUT with multi-frequency operation for compressed 3D spatial sensing. *IEEE Open J Ultrason Ferroelectr Freq Control* 4:37–51
 32. Zhang J, Zhao M, Zhang G, Zhang J, Du X (2024) 3D acoustic scaled boundary perfectly matched layer (SBPML) for acoustic-structure interaction problems. *Eng Anal Bound Elem* 164:105765
 33. Zhao L, Xia J, Yang C, Lu Y (2024) Modeling and experimental studies of mutual acoustic radiation effect on PMUT array. *Sens Actuators, A* 377:115677

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.