

High-order 2D mesh curving methods with a piecewise linear target and application to Helmholtz problems

V. S. Ziel^{a,*}, H. Bériot^b, O. Atak^b, G. Gabard^c

^a*Institute of Sound and Vibration Research, University of Southampton, SO17 1BJ
Southampton, United Kingdom*

^b*Siemens Industry Software NV, Simulation and Test Solutions, Interleuvenlaan 68,
B-3001, Leuven, Belgium*

^c*Laboratoire d'Acoustique de l'Université du Maine, CNRS, LAUM UMR 6613, Av. O.
Messiaen, 72085 Le Mans, France*

Abstract

High-order simulation techniques typically require high-quality curvilinear meshes. In most cases, mesh curving methods assume that the exact geometry is known. However, in some situations only a fine linear FEM mesh is available and the connection to the CAD geometry is lost. In other applications, the geometry may be represented as a set of scanned points. In this paper, two curving methods are described that take a piecewise fine linear mesh as input: a least squares approach and a continuous optimisation in the H^1 -seminorm. Hierarchic, modal shape functions are used as basis for the geometric approximation. This approach allows to create very high-order curvilinear meshes efficiently ($q > 4$) without having to optimize the location of non-vertex nodes. The methods are compared on two test geometries and then used to solve a Helmholtz problem at various input frequencies. Finally, the main steps for the extension to 3D are outlined.

Keywords: High-Order Methods, Curvilinear Meshes, Integrated Legendre Polynomials, Linear Target Mesh, Helmholtz problems, Geometry Error

1. Introduction

The impact of the geometric accuracy on the convergence of high-order discretization methods has been the subject of several studies in the literature [1, 2, 3, 4, 5]. In [2], Bassi *et al.* examine the influence of the geometry representation on the performance of a high-order DGM simulation of flow around

*Corresponding author, vss1g14@soton.ac.uk

obstacles. They conclude that small, even imperceptible inaccuracies in the normal description can lead to large errors in the final flow results. The impact of the normal accuracy on transient acoustics simulations is investigated in [3]. Both, linear and quadratic boundaries are compared, as well as a mixed approach, where the normal description in the boundary conditions is quadratic while the geometric description remains linear. It is found that this mixed approach does not significantly reduce the error in the simulation. Full quadratic boundary treatment on the other hand leads to an improved accuracy, while the computational time increases only marginally. Dey et al. performed a theoretical analysis for second-order elliptic boundary value problems [4], and obtained that a geometric interpolation order of $q \geq p - 1$ has to be applied in order to preserve the exponential p -convergence of the numerical error. Numerical tests for tension simulations on plates with a circular or elliptic hole are described in [5], showing that the numerical error rapidly stagnates as the geometric inaccuracies become dominant.

This paper focuses on the Helmholtz equation, which is used for a wide variety of wave propagation applications. For this class of problems, the classical low-order FEM is known to suffer from pollution errors at high-frequency, and recent studies indicate that this can be effectively mitigated by resorting to higher-order polynomials in the solution basis (p -FEM) [6, 7]. However, p -FEM typically requires coarser input meshes as compared to classical FEM, and this may lead to inaccurate boundary representations if a low-order (i.e. linear or quadratic) approximation is kept for the geometry. This becomes particularly stringent at high values of kh , when the wavenumber k and the mesh size h are such that a single element carries a significant portion of the wavelength. However, it should be noted that the exact influence of the frequency on the geometric requirements for these type of problems is not well understood. Hence, an attempt is made in this paper to explore the relationship between the geometric discretisation error (GDE) and the geometry induced error (GIE) for Helmholtz scattering problems. The conclusions of this initial analysis provide further motivation for employing efficient mesh curving methodologies for these applications.

Typically, high-order polynomial interpolation is used to obtain curved mesh boundaries. The generation and optimisation of high-order curvilinear meshes is an active field of research, and most methods fall in the *a-posteriori* curving category (e.g. [8, 9]). In this approach, a standard linear mesh is generated first, using well-established meshing algorithms. In a second step, the boundary of the linear mesh is curved towards the target geometry. Finally, optimisation and untangling procedures can be applied to increase the element quality and to obtain a valid mesh for numerical simulations.

In general, the target geometry is assumed to be known exactly. However, in many practical situations, the original CAD-geometry may not be accessible, for instance when the simulations are outsourced to different companies. Instead, a fine linear description of the geometry is often provided. This might be the case for some multi-physics problems as well. In NVH applications for instance, acoustics engineers most often only have the structural mesh at their disposal.

In biomedical applications, geometries are often described by scanned data, that form a cloud of points [10].

This paper hence focuses on the comparison of two a-posteriori curving methods that take a fine linear mesh as target curve. This target curve is called the *model mesh* or *Mmesh*, and is denoted by \mathcal{M} . As its only function is to provide an improved description of the geometric boundary, it is only a mesh of the domain boundary. The boundary of the coarse mesh that is curved iteratively is referred to as *computational mesh* or *Cmesh*, and is denoted by \mathcal{C} . The number of Mmesh edges is denoted by $n_{\mathcal{M}}^e$, the number of Cmesh edges by $n_{\mathcal{C}}^e$. For the curving methods, it is assumed that the vertices of the Cmesh form a subset of the Mmesh vertices.

In a preprocessing step, the initial state of \mathcal{C} could be obtained by applying a coarsening algorithm to \mathcal{M} . The algorithm should ensure that areas of high curvature are less coarsened, whereas areas with low curvature allow for larger Cmesh elements. This could prevent tangling that otherwise has to be treated in a postprocessing step. In this paper, only the curving algorithms are examined. The input \mathcal{C} is assumed to be of appropriate mesh quality. Before the curving, Cmesh vertices are snapped onto nearby Mmesh vertices in order to fulfil the subset requirement. The principle of the Cmesh and Mmesh is illustrated in Figure 1.

In the extension to 3D, feature line treatment will be discussed as well. For this paper, the knowledge of the feature lines is assumed to be provided along with the mesh, e.g. in form of closed surfaces groups that are bounded by the feature lines. As for the Cmesh generation, feature lines could also be detected in a preprocessing step (see e.g. [11]). For the 2D curving, this additional data is not necessary, as feature points are expected to be covered by both, Mmesh and Cmesh vertices. The latter therefore won't be relocated during preprocessing and the curving process.




Figure 1: Schematic of (a) the fine linear Mmesh of the boundary, (b) the coarse linear Cmesh of the domain, and (c) the curved Cmesh.

Most a-posteriori curving methods proposed to date rely on high-order nodal (Lagrange) shape functions (e.g. [8, 12, 13, 14]). They consist in finding an appropriate nodal distribution in the curved elements to ensure (i) a good representation of the geometry and (ii) that the elements are untangled and exhibit acceptable aspect ratios. In this paper, we examine an alternative approach based on hierarchic Lobatto shape functions, which are normally used for the

field variables approximation in p -FEM [15]. As opposed to the usual *nodal* Lagrange-based curvilinear mesh curving methods, the proposed approach is *modal*, in the sense that the fitted geometric coefficients do not correspond to specific node locations anymore.

Two a-posteriori *modal* curving methods, initially introduced in [16], are further examined hereafter in 2D. The first method is a least squares approach, where only Mmesh vertices that lie on the exact geometry are used to describe the target curve. In the second method, the curved mesh is obtained from a global optimisation in the H^1 -seminorm. This optimisation takes full advantage of the δ -property of the Lobatto shape functions and hence does not include any matrix inversion. However, it relies on the linearly interpolated Mmesh elements that only form an approximation of the underlying exact geometry. For both modal methods, the curving is done element-wise, i.e. one Cmesh element at a time. It is assumed that all Cmesh vertices coincide with Mmesh vertices. This leads to a straightforward assignment of Mmesh vertices and edges to the Cmesh edges.

The two proposed modal curving methods are evaluated and benchmarked against the nodal approach on two test cases: a unit circle described by 12 Cmesh elements, and a distorted ellipse, described by 20 Cmesh elements. For both examples, the geometric discretisation error (GDE) and its convergence for an increasing geometric order q are reported. Afterwards, the curving algorithms are further tested by running a Helmholtz scattering problem on the curved meshes. The GDE-GIE relationship that was obtained with the nodally-curved meshes is checked and compared with the relationship obtained with the modally-curved meshes.

The structure of the paper is as follows. The influence of the geometry representation on a Helmholtz problem is first assessed in Section 2. In Section 3, the two modal curving methods are introduced. Their ability to represent a given Mmesh boundary is assessed in Section 4. The curving methods are then further benchmarked in Section 5 by running a Helmholtz scattering simulation on the curved meshes. The extension of the most promising method to 3D is outlined in Section 6. Finally, Section 7 summarizes the paper and discusses future developments.

2. Quantifying the geometric error influence

Prior to introducing the proposed mesh curving methodologies, this section examines the influence of the geometric representation on the accuracy of wave propagation problem, modelled by the Helmholtz equation.

2.1. Scattering test case

As benchmarking scenario, the scattering of a plane wave at a rigid cylinder of radius r_s is considered. An analytical solution \mathbf{p}_{ex} for the pressure \mathbf{p} can be provided [17], which is the solution of the following Helmholtz problem with




Figure 2: Sketch of the scattering test case.

angular frequency ω and wave number k on the bounded domain Ω :

$$\begin{aligned} \nabla^2 \mathbf{p} + k^2 \mathbf{p} &= 0, \text{ on } \Omega, \\ \frac{\partial \mathbf{p}}{\partial n} &= 0 \quad \text{on } \Gamma_s, \\ \frac{\partial \mathbf{p}}{\partial n} &= \frac{\partial \mathbf{p}_{\text{ex}}}{\partial n} + i\omega \mathbf{p}_{\text{ex}} - i\omega \mathbf{p} \quad \text{on } \Gamma_o, \end{aligned}$$

with Robin and Neumann boundary conditions on the outer boundary Γ_o and scattering boundary Γ_s respectively (see Figure 2). In all the remainder of the paper, parameters are made non-dimensional by setting the sound speed and density to unity. The corresponding Galerkin variational formulation consists in finding the solution $\mathbf{p} \in H^1(\Omega)$ such that

$$\begin{aligned} \int_{\Omega} \nabla \varphi \cdot \nabla \mathbf{p} \, d\Omega - k^2 \int_{\Omega} \varphi \mathbf{p} \, d\Omega \\ + i\omega \int_{\Gamma_o} \varphi \mathbf{p} \, d\Gamma_o = \int_{\Gamma_o} \varphi \left(\frac{\partial \mathbf{p}_{\text{ex}}}{\partial n} + i\omega \mathbf{p}_{\text{ex}} \right) d\Gamma_o, \quad \forall \varphi \in H^1(\Omega). \end{aligned} \quad (1)$$

To generate the high-order nodally curved grids, the open-source software *Gmsh* [12] is used. The meshes are generated as follows: (i) the computational domain Ω is first approximated by a set of non-overlapping straight sided triangular elements $\hat{\Omega} = \bigcup_{e=1}^{N_e} T_e$ obtained from a classical Delaunay triangulation (ii) this initial linear mesh is then curved a-posteriori, using a *nodal-curving* approach of order q and (iii) the resulting high-order Lagrange elements are untangled and smoothed applying the algorithms described in [18, 19].

The variational formulation (1) is discretized using p -FEM [15]. The polynomial order p representing the unknown pressure field is taken in the range $p = 1, 2, \dots, 10$. High-order Gauss-Legendre quadrature is used for the evaluation of the different integrals.

For Helmholtz FE formulations, the integrand in the mass bilinear form is typically polynomial. It can therefore be evaluated exactly using $2p + 2(q - 1)$. The stiffness term on the other hand involves a non-polynomial integrand (rational function), due to the presence of the Jacobian inverse in the field

derivatives, which needs to be approximated by higher-order quadrature rules.
140 From numerical experiments, it was however found that a quadrature rule of
 $2p + 2(q - 1)$ was sufficient to minimize the quadrature error, as setting a higher
quadrature rule did not allow to obtain more accurate solutions.

It is interesting to further examine how the proposed rule compares with
results from previous studies. In particular, in Ref. [4], Dey et al. examine
145 the influence of the geometrical and quadrature errors on second order elliptic
boundary value problems, including Helmholtz. They conclude that, for the ex-
ponential convergence to be preserved with curved geometries, the integration
order should be increased by $p - 1$. There, it is also stated that the geometric
order verifies $q \geq p - 1$ to reach optimal convergence of the global error. It is
150 interesting to note that the proposed pragmatically-derived rule follows these
requirements: in the exponential convergence regime ($q \geq p - 1$), a quadrature
order elevation of $2(q - 1)$ is always more conservative than what the theory
suggests. Otherwise, in the pre-convergence regime ($q < p - 1$), the geometri-
cal error becomes so dominant that the quadrature elevation only marginally
155 influences the global result.

It is worth emphasizing that the geometric order $q = 1, 2, \dots, 4$ and the field
approximation order $p = 1, 2, \dots, 10$ are completely independent of each other
and can be chosen separately.

2.2. Evaluation of the numerical error

The standard error evaluation procedure consists in integrating the differ-
ence between the analytical and the numerical solution over the computational
domain $\hat{\Omega}$ using some norm (typically L^2 or H^1). However, $\hat{\Omega}$ is not necessar-
ily a subset of Ω , and as a result, this may lead to evaluating the analytical
solution at quadrature points located inside the scatterer, i.e. outside of the
165 computational domain (see Figure 3(a)). To avoid this issue, a simple alterna-
tive approach is proposed to evaluate the numerical error. The relative L^2 error,
denoted $E_{2,R}$ is evaluated along a ring located in the middle of the domain (see
Figure 3(b)). Since the mesh elements are not aligned with this interior contour,
a trapezoidal rule is used to evaluate the integrals. The number of evaluation
170 points for the integration is chosen in an adaptive way, to ensure appropriate
convergence for each configuration. As the theoretical error estimate is inde-
pendent of any dimensions, it is expected that exponential convergence will be
observed.

2.3. Geometric Discretisation Error (GDE) measure

Many different approaches have been proposed in the literature to measure
175 the quality of the geometric approximation. They can roughly be classified
into three categories, namely normal-based, distance-based and area-based error
measures.

The two standard distance-based measures are the Hausdorff- and the Fréchet-
180 distance. Both are defined as two nested optimisations over the pointwise dis-

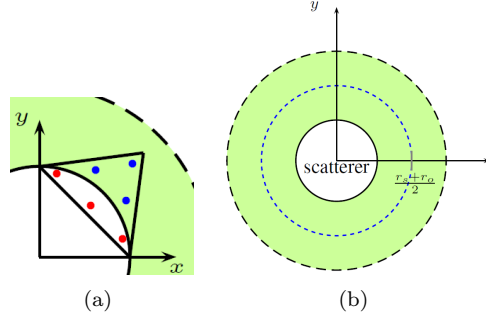


Figure 3: Schematic figures showing (a) the problem in the standard numerical error evaluation for the scattering test case and (b) the alternative approach. The red dots in (a) are the quadrature points that fell out of Ω .

tance between the two considered curves [20]. Therefore, their calculation is computationally costly [9]. In [14], a variation of the Fréchet distance was introduced. It results in a continuously differentiable measure that can be used for optimisation based mesh curving.

185 Another, less computationally intensive alternative, is to resort to area-based error measures. In Ref. [9], the relation between the two measure types is explored on several numerical examples. It is observed that an optimisation with respect to the area-based measure leads also to a reduction of the Hausdorff distance.

190 Normal- or curvature-based measures are less commonly used in the literature. Most of them rely on pointwise associations of the two curves, and therefore depend on their corresponding parametrisations. They either consider the maximal angle deviation of the normals [21] or the mean and standard deviation [22]. In the latter paper, also the mean and standard deviation of the differences in the Gaussian and the mean curvatures are considered.

195 Finding an optimal GDE measure that will fit every practical application on a given physical problem therefore remains a challenge. As pointed out by one of the reviewers, in many occasions, one or several pathological cases can be identified where the different GDEs may be proven defective and/or inconsistent.

200 Let us consider a concrete example, illustrated in Figure 4. Suppose the geometry is a circle of radius R and suppose the boundary mesh fits $x = R' \cos(\theta)$, $y = R' \sin(\theta)$ where $R' = R + a \cos(n\theta)$. Even if the radius perturbation a is small compared to R , by increasing n , the normal of the mesh will deviate more and more from the normal of the circle, leading to higher values in the normal-based GDE. However, by contrast, this will only loosely influence the area-based GDE, as the area between the approximation and the exact circle will not vary significantly for increasing n values. It is also worth noting that the normal-based approach proves defective at $n = 0$, see Figure 4(a). There, the circle normal is exactly represented at each point, although solving a scattering

205

210 problem with this different circle will obviously deteriorate the global solution.

It is hence difficult to infer which indicator will be most representative of the actual geometrical error influence for a given Helmholtz problem. Other secondary factors may also be of influence, like the regularity of the approximation and the type or positioning of the acoustic sources.




Figure 4: Artificially generated geometric approximations of the unit circle Γ_s (dotted line) with similar area-based GDE, but very different normal-based GDE.

215 In this study, an area-based measure was chosen to monitor the geometric accuracy. In practice, it is calculated as the surface area A between the approximated and the exact boundary Γ_s , divided by the length of Γ_s . The motivation behind this choice is that (by contrast to most normal-based approaches) the area-based measure is independent from the curve parametrisation.

220 For the circle, the polar coordinate transformation provides an inverse of the exact boundary's parametrisation, which can be used to calculate the area directly. Hence an approximation of both curves by a polygon, as proposed for instance in [9] is not necessary. Instead, parametrisations $c_i(s)$ and $\hat{c}_i(s)$ of the two curves are constructed for each element i over the same reference interval
225 $s \in [s_0, s_1]$.

It can be shown that the following relation holds

$$A_i(s + ds) - A_i(s) = \frac{1}{2} |(c_i(s + ds) - \hat{c}_i(s)) \times (c_i(s) - \hat{c}_i(s + ds))|$$

when travelling along the parametrisation of the curves from s to $s + ds$. Applying linear Taylor series expansions of c_i and \hat{c}_i and rearranging the formula, the following ordinary differential equation is obtained:

$$\frac{dA_i}{ds}(s) = \frac{1}{2} |(\hat{c}_i(s) - c_i(s)) \times (\hat{c}'_i(s) + c'_i(s))|, \quad (2)$$

where c' is the derivative of c . To get the area between the two curves over one element, Eq. (2) is solved for the initial value $A_i(s_0) = 0$. The elementwise areas A_i are then summed up to obtain the total area $A = \sum_i A_i$. Finally, the relative geometric discretisation error (referred to as the GDE) is obtained by
230 dividing A by the circumference of the circle $GDE_A = A/2\pi r_s$, as suggested in [9].

2.4. Numerical results and interpretation

2.4.1. Simulation results

To assess the influence of the geometric representation on the numerical error, p -convergence curves are first obtained at $\omega = 1$ for a varying geometric order q , see Figure 5. Expectedly, as p is increased, the numerical model first




Figure 5: Stagnating p -convergence curves of the numerical error evaluated along the ring $r = 1.5$ for a fixed $\omega = 1$ and varying mesh size (a) $n_C^e = 8$, corresponding to $h \approx 0.79$ and (b) $n_C^e = 32$, corresponding to $h \approx 0.20$.

converges exponentially. The loss of exponential convergence occurs at $p^* = q + 1$, in accordance with the theoretical analysis performed in [4]. The decrease in the numerical error slows down and, in most cases, the stagnation appears at $p = q + 2$. In the absence of other possible sources of error, this stagnation level is identified as the geometric induced error (GIE).

The GIE decreases for increasing geometric order q , which is in accordance with results previously described in the literature [5, 23]. Furthermore, as expected, applying an h -refinement leads to a rapid decrease of the GIE and this h -convergence rate is much larger for higher order meshes. It is, however, striking to note that the GIE is very marginally improved when going from a quadratic ($q = 2$) to a cubic ($q = 3$) geometric representation. This odd-even trend is also visible in the corresponding GDE_A plots, which are not reported here for conciseness. This will be commented further in Section 4.1.

The input frequency also plays an important role. In Figure 6, results for simulations with two different angular frequencies ω are compared. The GIE increases substantially with increasing frequencies, such that the stagnation level is sometimes not even obtained within the plotted p values. This confirms the intuitive statement that high-frequency simulations are more sensitive to geometric inaccuracies.

2.4.2. GDE-GIE relationship

The stagnation levels or GIEs are now evaluated for an extensive number of configurations, with varying frequencies $\omega \in \{1, 2, 4, 6, \dots, 30, 32, 64\}$ and




Figure 6: Stagnating p -convergence curves of the numerical error for the scattering test case, with numerical error evaluation along a ring, comparing different simulations for (a) $\omega = 1$ and (b) $\omega = 16$.

260 mesh size $n_{\mathcal{C}}^e \in \{8, 12, 16, \dots, 72, 76, 80\}$. The data sets obtained from these calculations are reported in Figure 7(a). Simulation results that did not reach the stagnation stage and for which the GIE cannot be firmly evaluated are discarded from the evaluations.

For a fixed ω , the plots show a linear dependency between $\log_{10}(\tilde{E}_2)$ and $\log_{10}(\text{GDE})$, with almost no dependency on the geometric order q . With increasing frequency, an offset is added to the relationship which changes the intercept. This offset is shown to logarithmically depend on the frequency ω . Combining these observations, a linear model of the form

$$\log_{10}(\tilde{E}_2) \approx \beta_{\text{const}} + \beta_{\text{geo}} \log_{10}(\text{GDE}) + \beta_{\omega} \log_{10}(\omega) \quad (3)$$

is fitted to the data using the statistical software *R* [24]. This model can be rearranged as a non-linear model of the form

$$\tilde{E}_2 \approx 10^{\beta_{\text{const}}} \text{GDE}^{\beta_{\text{geo}}} \omega^{\beta_{\omega}}. \quad (4)$$

The parameters β_i and the values of the adjusted \mathbf{R}^2 of the model are given in Table 1.

β_{const}	β_{geo}	β_{ω}	$\tilde{\mathbf{R}}^2$
1.3846	1.0106	1.2314	0.9975

Table 1: Parameter values obtained for the fitted GDE-GIE model (3), and the corresponding value of the adjusted \mathbf{R}^2 .

As an indicator of the quality of the regression model, the statistical parameter \mathbf{R}^2 measures how much of the variance in the data is explained by the model. It is defined as

$$\mathbf{R}^2 := \frac{\sum_i (\hat{d}_i - \tilde{d})^2}{\sum_i (d_i - \tilde{d})^2}, \quad (5)$$

with the i -th entry \mathbf{d}_i of the data vector \mathbf{d} , the mean $\tilde{\mathbf{d}}$ of the vector \mathbf{d} , and the i -th fitted value $\hat{\mathbf{d}}_i$ of the model (see [25]). In the given case, \mathbf{d}_i corresponds to the error GIE \tilde{E}_2 of the i^{th} simulation. Note that the definition in (5) is only valid in the case of uncorrelated residuals. This assumption is verified by examining the residuals against fitted plot. It is found to be well behaved, which confirms the reliability of the R^2 measure (it is not reported here for conciseness). We provide here the adjusted R^2 :

$$\tilde{R}^2 := 1 - \frac{m-1}{m-n_\beta-1} (1-R^2),$$

265 which takes into account the sample size m , and the number of parameters n_β corresponding to non-constant explanatory variables. A value of $\tilde{R}^2 = 0.9975$ is obtained, which indicates that the proposed GDE-GIE regression model (3) provides an excellent fit to the set of simulation data (see Figure 7(b)).

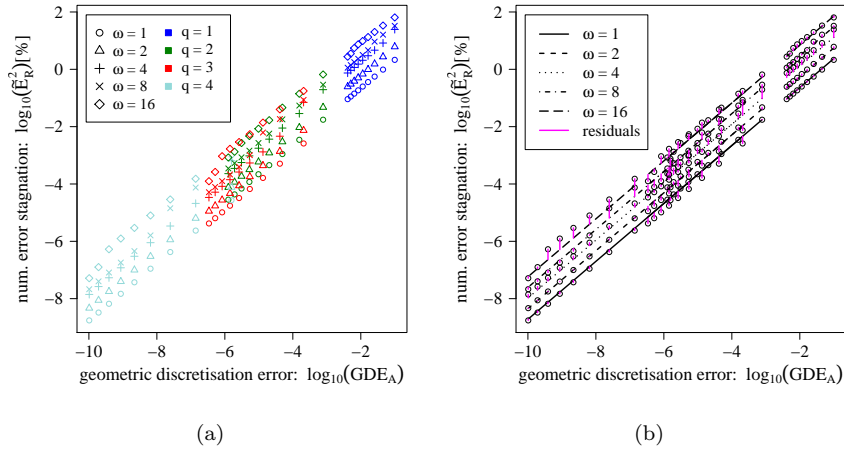


Figure 7: Plots of (a) a subset of the data set used for fitting the GIE \tilde{E}_2 against the GDE and (b) the fitting of the GIE \tilde{E}_2 against the GDE with model (3). In (a), the different colors represent the geometric order of the mesh, while the different shapes represent the ω value of the simulations.

270 This model infers that the GIE, i.e. the accuracy that can be expected from a given acoustics simulation increases linearly with the GDE. Furthermore, for a fixed GDE (i.e. for a given geometric approximation), the GIE is found to vary sup-linearly with the frequency (power of 1.25 approximately).

275 Simple physical considerations may be introduced to justify the latter observation. Consider the reflection of a plane wave impinging a solid wall at normal incidence. Consider now that the wall is artificially shifted to shorten or elongate the domain. The inexact location of the scatterer leads to a phase shift in the solution, which is proportional to the frequency. Intuitively, on this one-dimensional problem, the expected behaviour is hence that, for a fixed GDE, the resulting numerical error will grow linearly with the frequency. However, in

280 the provided 2D configuration, the additional complexity in the frequency dependence is thought to be brought by the presence of more complex scattering patterns, from e.g. normal discontinuities at inter-element junction or creeping waves along the scatterer. In practice, the super-linear dependence implies that the GDE should be divided by approximately 2.5 when doubling the frequency,
 285 in order to ensure an equivalent convergence of the global error.

This analysis provides a strong motivation for the introduction of high-order geometric representations. The next section presents two mesh curving methods that can be applied in order to ensure that the GDE error is controlled effectively.

290 3. Modal Curving methods

In this Section, two modal mesh curving algorithms are introduced, followed by a presentation of the blending method.

3.1. Motivations

As described in the introduction, the objective pursued here is to propose
 295 some strategies to efficiently curve a coarse linear Cmesh to approach a fine linear mesh, represented by the Mmesh (see Figure 1).

Each curved \mathcal{C} element is described as a linear combination

$$\hat{\mathbf{f}} = \sum_{i=1}^{n_q} \alpha_i b_i \quad (6)$$

of curving order q , with a polynomial basis $\{b_i\}_{i=1}^{n_q}$ and coefficients α_i . These coefficients will be determined by the curving processes described in this paper. As base functions $\{b_i\}_{i=1}^{n_q}$, integrated Legendre polynomials are used

$$\begin{aligned} l_0(\xi) &= \frac{1-\xi}{2}, \\ l_1(\xi) &= \frac{\xi+1}{2}, \text{ and} \\ l_k(\xi) &= \frac{1}{\|L_{k-1}\|_2} \int_{-1}^{\xi} L_{k-1}(x) dx, \text{ for } 2 \leq k, \xi \in [-1, 1], \end{aligned}$$

where $L_i(\xi)$ is the i^{th} Legendre polynomial. These polynomials are also called *Lobatto polynomials* and have several interesting properties which justify their consideration [15]:

- In 2D, the basis functions fulfill the equation

$$\int_{-1}^1 l'_i(\xi) l'_j(\xi) d\xi = \delta_{ij}, \text{ with } \delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases}, \quad (7)$$

300 which is referred to as the δ -property. It is based on the orthogonality of the Legendre shape functions in the L^2 -norm. This leads to explicit

formulas for the interpolation coefficients α_i , if they are obtained by optimisation in the H^1 -seminorm. In 3D, this property is not conserved for triangular face shape functions, but the construction of orthonormal bubble shape functions is described in [26]. This paper only considers 2D cases.

- The shape functions form hierarchical bases. Together with the δ -property, this results in hierarchical coefficients α_i , meaning that for every q increment, only one additional α_i needs to be evaluated.
- Furthermore, the Lobatto bases are modal bases, in the sense that they do not rely on specific nodal positions. This may simplify the embedding of the curved boundary for full simulation meshes. For nodal approaches, the location of element interior nodes has to be obtained from an optimization procedure during or after the curving (see for instance [27]). With a modal basis, such an “interior” optimization (controlled by the coefficients for the face shape functions or the mesh internal edges) is not always required, provided the element with curved boundary edges is valid and of adequate quality. In combination with the hierarchical property, the modal property also allows for local variations of the mesh curving order (i.e. local q -refinements).

The geometric polynomial approximation $\hat{f}(\xi)$ in Equ. (6) can be split up into a known linear interpolation and the interpolation with higher-order shape functions. The linear interpolation is fully described by the Cmesh vertex positions $\mathbf{p}_1 = (p_{1,x}, p_{1,y})$ and $\mathbf{p}_2 = (p_{2,x}, p_{2,y})$ that lie on the exact geometry. Therefore, the α_0 and α_1 are known and only α_j for $j = 2, \dots, q$ for the higher-order shape functions have to be determined. The interpolation with the non-linear shape functions will be called the *polynomial deflection*

$$\hat{f}(\xi) = \underbrace{\frac{1}{2}(-\xi + 1)\mathbf{p}_1 + \frac{1}{2}(\xi + 1)\mathbf{p}_2}_{\text{linear interpolation}} + \underbrace{\sum_{i=2}^q \alpha_i l_i(\xi)}_{\text{polynomial deflection}}. \quad (8)$$

Two different strategies are proposed to obtain the unknown interpolation coefficients α_i . These two strategies are referred to as the least-squares modal curving and the H^1 modal curving. They are introduced hereafter.

3.2. Least-squares modal curving

The first curving method aims at approaching the Mmesh vertices by minimizing their distance to the corresponding points on the Cmesh using a least-squares technique. With this approach, only the points on the exact geometry are taken as the target (i.e. the \mathcal{M} cloud of points), as the Mmesh vertices lie on the exact geometry. A similar strategy using Bézier curves is described in [10].

For the curving of a single Cmesh element, the n corresponding Mmesh vertices (x_i, y_i) are approximated by the polynomial $\tilde{f} = \sum_{j=0}^q \alpha_j l_j$. The co-

efficients α_j are determined such that they fulfill the following minimisation problem:

$$\min_{\alpha} \sum_{i=1}^n \|(x_i, y_i) - \tilde{\mathbf{f}}(\xi_i)\|_2^2,$$

where

$$\xi_i = -1 + 2 \frac{\sum_{j=1}^{i-1} h_j}{\sum_{j=1}^{n-1} h_j} \in [-1, 1]$$

is the parameter corresponding to (x_i, y_i) for a parametrisation based on the length h_j of the j -th Mmesh element on the Mmesh arc.

The least-squares approach then leads to a linear system of equations $\mathbf{A}\alpha = \mathbf{b}$. Both the x - and y -coordinate are described independently, which leads to two independent linear systems, one for each coordinate, with $\alpha_i = (\alpha_{i,x}, \alpha_{i,y})$. The system for the x -coordinate is shown below (subscript x is dropped for readability):

$$\mathbf{A} = \begin{pmatrix} a_{2,2} & \cdots & a_{2,q} \\ \vdots & \ddots & \vdots \\ a_{q,2} & \cdots & a_{q,q} \end{pmatrix}, \quad \alpha = \begin{pmatrix} \alpha_2 \\ \vdots \\ \alpha_q \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} b_2 \\ \vdots \\ b_q \end{pmatrix}, \quad \text{with}$$

$$a_{k,j} = \sum_{i=1}^n l_k(\xi_i) l_j(\xi_i),$$

$$b_k = \sum_{i=1}^n x_i l_k(\xi_i) - \sum_{j=0}^1 \alpha_j \sum_{i=1}^n l_k(\xi_i) l_j(\xi_i).$$

The system solved for the y -coordinate follows analogously.

3.3. H^1 modal curving

335 A second modal-based curving method is introduced, where a projection-based interpolation [15] is applied to obtain the α_i coefficients in the polynomial description of the curved edges. Thereby, the δ -property (7) of the Lobatto shape functions is used, by considering an optimisation in the H^1 -seminorm. By contrast to the previously introduced method, which only considers the \mathcal{M}
340 cloud of points, this approach requires a *continuous* curve as target. In this study, a simple linear interpolation of the Mmesh vertices is considered.

The difference between the target curve and the approximating curve is minimised in each coordinate independently. Only the formula for the x -coordinate is presented hereafter, showing the approximation of the target f_x by \hat{f}_x . The
345 polynomial approximation of the y -coordinate follows analogously and will result in an additional set of α_i -values. Similarly to $\tilde{\mathbf{f}}(\xi)$ in Eq. (8), the parametrisation of the target curve can be split up into a linear part and a deflection d . Here, d is generally non-polynomial and depends on the Mmesh vertices of the target segment.

The polynomial approximation is done in the H^1 -seminorm by finding α_i such that

$$\min_{\alpha} |f_x - \hat{f}_x|_{H^1} = \min_{\alpha} \int_{-1}^1 (f'_x - \hat{f}'_x)^2 d\xi.$$

The modal coefficients α_k that solve the minimisation problem have to satisfy

$$\frac{\partial}{\partial \alpha_k} \int_{-1}^1 (d' - \sum_{i=2}^q \alpha_i l'_i(\xi))^2 d\xi = 0,$$

which can be reformulated as

$$\alpha_k = \int_{-1}^1 d' l'_k d\xi, \quad k = 2, 3, \dots, q.$$

The calculation of the α_k coefficients depends on the parametrisation of the Mmesh arc that forms the target curve. For this, a physical length based parametrisation will be considered. The lengths of the subintervals in the reference space are adapted to the lengths of the Mmesh elements. The resulting formula for the α_k is

$$\begin{aligned} \alpha_k = & \sum_{i=1}^{n-1} \left(\int_{\xi_i}^{\xi_{i+1}} \frac{\sum_{j=1}^{n-1} h_j}{2h_i} (x_{i+1} - x_i) l'_k(\xi) d\xi \right) \\ & + \left(\frac{1}{2} p_{1,x} - \frac{1}{2} p_{2,x}, \right) \int_{-1}^1 l'_k(\xi) d\xi, \end{aligned}$$

350 with n the number of Mmesh vertices on the Mmesh arc, h_j the lengths of the j -th Mmesh element on the Mmesh arc, and ξ_i the parameter corresponding to the i -th Mmesh vertex on the Mmesh arc as defined above.

3.4. Blending method

355 The two curving methods previously introduced provide a parametrization of the curved \mathcal{C} elements along the boundary. Nevertheless, they do not provide any description of the geometric transformation inside the domain. This is required for evaluating the Jacobians of the neighbouring curved triangles (see Figure 1 (c)). This additional step is referred to as *blending* and follows the idea of the Lobatto shape functions construction [15].

Elements with a curved edge are represented by a mapping from the reference triangle to the physical space of the form

$$\mathbf{f}_i(\xi, \eta) = \sum_{i=1}^3 \mathbf{x}_i \varphi_T^{v_i}(\xi, \eta) + \sum_{i=1}^3 \boldsymbol{\kappa}^{e_i}(\xi, \eta),$$

where the first term is the linear mapping of the straight sided edges defined by the vertices \mathbf{x}_i with $\varphi_T^{v_i}$ the usual vertex Lobatto shape functions in the triangle (see [15]), and $\boldsymbol{\kappa}^{e_i}$ maps the edge e_i and has the property

$$\boldsymbol{\kappa}^{e_i}(\xi, \eta) = \begin{cases} \mathbf{0} & \text{on } e_j, i \neq j \\ \mathbf{d}^{e_i}(\tilde{\xi}^{e_i}(\xi, \eta)) & \text{on } e_i \end{cases}, \quad (9)$$

with $\mathbf{d}^{e_i}(\tilde{\xi}^{e_i}) = \mathbf{f}_c^{e_i}(\tilde{\xi}^{e_i}) - \mathbf{f}_s^{e_i}(\tilde{\xi}^{e_i})$ the difference between the actual, possibly curved edge $\mathbf{f}_c^{e_i}$ and the straight sided edge $\mathbf{f}_s^{e_i}$. The auxiliary variable $\tilde{\xi}^{e_i} \in$

$[-1, 1]$ is defined as $\tilde{\xi}^{e_i}(\xi, \eta) = \lambda_{i+1}(\xi, \eta) - \lambda_i(\xi, \eta)$, with $\lambda_i(\xi, \eta)$ the standard Barycentric coordinates of the triangle and $\lambda_4 = \lambda_1$. It is applied to define the 2D edge shape functions along edge e_i

$$\psi_j^{e_i}(\xi, \eta) = \lambda_i(\xi, \eta) \lambda_{i+1}(\xi, \eta) \frac{l_j(\tilde{\xi}^{e_i}(\xi, \eta))}{l_1(\tilde{\xi}^{e_i}(\xi, \eta)) l_0(\tilde{\xi}^{e_i}(\xi, \eta))}.$$

They vanish on all edges other than e_i and their trace along e_i corresponds to l_j (see [15]). From this construction it follows for the x -component of κ^{e_i}

$$\kappa_x^{e_i}(\xi, \eta) = \sum_{j=2}^q \psi_j^{e_i}(\xi, \eta) \alpha_{x,j}^{e_i},$$

360 with $\alpha_{x,j}^{e_i}$, the α -coefficient of l_j for the x -component of edge e_i . The function for the y -component is constructed analogously.

The resulting formula of the reference mapping is

$$\mathbf{f}_i(\xi, \eta) = \sum_{i=1}^3 \mathbf{x}_i \varphi_T^{v_i}(\xi, \eta) + \sum_{i=1}^3 \sum_{j=2}^q \alpha_j^{e_i} \psi_j^{e_i}(\xi, \eta).$$

It forms the geometric representation of the simulation domain. The geometric description of the element enters the simulation equations in form of its Jacobians which can be easily calculated from the formula of \mathbf{f}_i . If more than
 365 one triangle edge is curved, as it is e.g. necessary to prevent tangled elements, the formula remains the same. The coefficients of each edge enter the equation independently.

Note again that no nodes have to be placed for the modal blending. This is an advantage compared to mesh curving methods based on Lagrange polynomials,
 370 for which the positioning of interior nodes typically requires a careful treatment (see e.g. [27]).

4. Comparison of the methods

The two modal-based curving methods are now benchmarked on two geometries, namely the unit circle and a distorted ellipse. The quality of the resulting
 375 curved boundary \mathcal{C} is assessed by measuring the geometric discretization error GDE for each configuration.

4.1. Circle test case

The performance of the two methods is first examined on the unit circle geometry. As a reference, high-order nodally curved meshes are also generated,
 380 using *Gmsh*, following the approach described in Section 2.1. In order to ensure that the Cmesh vertices form a subset of the Mmesh vertices, a preprocessing

step is applied. For each Cmesh vertex the closest Mmesh vertex in Euclidean distance is determined and the position of the Cmesh vertex is adjusted to fall together with the Mmesh vertex. For thin geometries the application of the
 385 Euclidean distance instead of the geodesic distance could generally lead to an invalid Cmesh. Here it is assumed that the Mmesh is refined enough to prevent this. In practice, a refinement of the same size as the geometry thickness should be sufficient, as the maximal distance from a Cmesh vertex to its closest Mmesh vertex is one half of the Mmesh element size. The accuracy of the resulting
 390 curved meshes is assessed by measuring the distance to the exact geometry (and not to the target mesh), using the parametric area-based GDE measure Eq. (2) introduced in Section 2.3.

For the interpretation of the results, it is useful to introduce the so-called MpC ratio, defined as

$$\text{MpC} = \frac{n_{\mathcal{M}}^e}{n_{\mathcal{C}}^e},$$

and representing the element size ratio between the refined target Mmesh and the coarse linear Cmesh.




Figure 8: q -convergence of the relative area to the exact circle with the H^1 (a) and the least squares (b) modal curving methods for different MpC ratios.

Figure 8 shows q -convergence results of GDE_A for the two proposed curving
 395 methods with 12 Cmesh elements. The results for the high-order Lagrange elements, obtained using the exact geometry as target curve, are included as a reference.

Stagnation is observed for the q -convergence curves for the H^1 modal curving
 400 approach. With increasing MpC, the stagnation onset shifts to higher q and the stagnation level decreases. From this behaviour it is deduced that the stagnation is introduced and controlled by the geometric accuracy of \mathcal{M} . With increasing q , the approximation of \mathcal{M} becomes more accurate and the difference between \mathcal{M} and the exact circle becomes the dominating error. As this error is fixed
 405 with the Mmesh size, the overall error stagnates.

Finally, the results with the least squares approach are shown in Figure 8(b). As the method considers only the geometrically exact Mmesh vertices, the linear interpolation in between vertices is not taken into account. Therefore, the stagnation induced by the inaccuracy of the Mmesh is overcome. However, it should be emphasized that the method demands Mmesh refinement of q elements per Cmesh element in order to be valid.

It is worth emphasizing that the q -convergence of all methods shows an odd-even stepwise behaviour. For $q = 2k$ and $q = 2k + 1$, nearly the same GDE_A values are observed.

This behaviour is found to be specific to the circle, as it is not observed on other geometries (see Section 4.2). An intuitive explanation is that it is linked to the symmetry of the circle arcs over each element, but this could not yet be confirmed by simple numerical experiments.

It is now worth examining the computational costs of the modal curving method. Even with the considered implementation in Matlab and a high Mmesh refinement, the CPU time for the calculation of the α values is a fraction of a second ($\sim 0.1s$ for $MpC = 500$, $\sim 0.15s$ for $MpC = 1000$) for both methods. It is therefore considered negligible compared to the CPU time for the acoustic simulations.

4.2. Distorted ellipse

4.2.1. Geometry and meshes

As a second, more realistic test case, a distorted ellipse is considered. The same geometry was also tested in [28]. It is obtained by adding cosine perturbations (bumps) along the curve of a high aspect ratio ellipse

$$\begin{pmatrix} x(\theta) \\ y(\theta) \end{pmatrix} = \begin{pmatrix} (r_x + h_b \cos(\theta n_b)) \cos(\theta) \\ (r_y + h_b \cos(\theta n_b)) \sin(\theta) \end{pmatrix} \quad \text{for } \theta \in [0, 2\pi], \quad (10)$$

with radii $r_x = 0.5$, $r_y = 5$. The number of bumps is set to $n_b = 10$ and their height to $h_b = 0.1$. The resulting geometry is shown in Figure 9(a).

In real-life applications, the meshes are not expected to be regularly sized. Therefore, the linear meshes on the distorted ellipse are randomly perturbed using n_v uniformly distributed pseudo random values γ_i chosen in the interval $[-1/(3n_v), 1/(3n_v)]$, where n_v is the number of equidistant nodes. The final reference nodes are obtained as $\theta_i = \tilde{\theta}_i + \gamma_i$, with $\tilde{\theta}_i = (i-1)\pi/(n_v-1)$ equally spaced in $[0, \pi]$. The mesh vertices are then calculated with the parametrisation (10). This constructs the upper half of the geometry. For the other half, $\theta_i + \pi$ are mapped onto the physical space. The resulting curve is closed using $\gamma_1 = \gamma_{n_v} = 0$. As in the circle test case, the Cmesh vertices are finally snapped onto the closest Mmesh vertices before the actual curving.

4.2.2. GDE measure for the comparison

In the initial circle test case, the GDE error was evaluated by measuring the area between the Cmesh and the exact geometry. Here, we consider that the exact geometry is not available and directly compute the area between the Cmesh

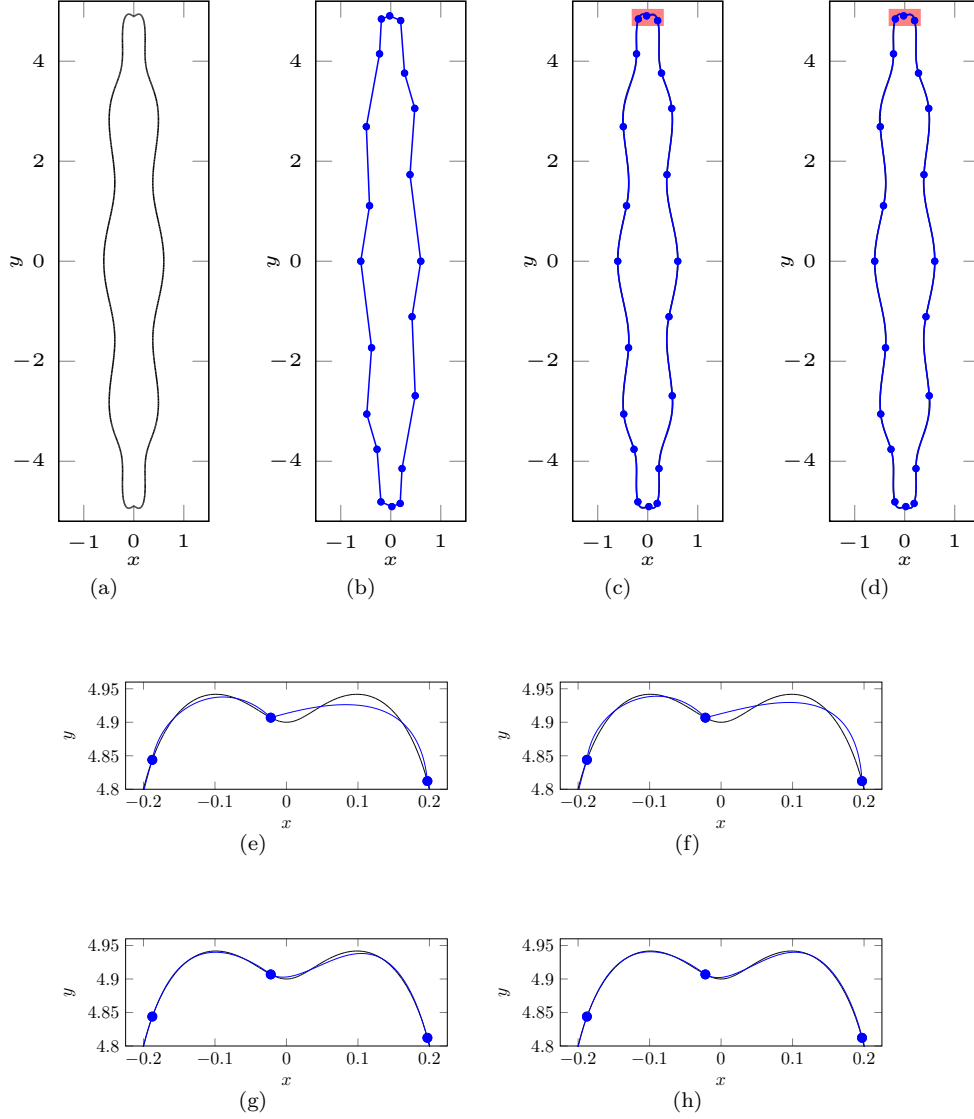


Figure 9: The distorted ellipse test case. (a) shows the target geometry, (b) the linear Cmesh with 20 elements. Third order curves after the curving are shown for (c) the H^1 modal approach and (d) the least squares approach. The red rectangle indicates the zoom-in area which is shown in figures (e-f) for $q = 3$ and (g-h) for $q = 4$. Figures (e) and (g) to the H^1 modal approach and (f) and (h) to the least squares.

and the target curve \mathcal{M} . The latter is measured by dividing the space between M and C into sufficiently small convex quadrangles $\mathbf{Q}_i = (\mathbf{q}_{i,1}, \mathbf{q}_{i,2}, \mathbf{q}_{i,3}, \mathbf{q}_{i,4})$

with respective area

$$A_{\mathbf{q}_i} = \frac{1}{2} |(\mathbf{q}_{i,2} - \mathbf{q}_{i,4}) \times (\mathbf{q}_{i,1} - \mathbf{q}_{i,3})|.$$

The quadrangles' size is determined adaptively for each configuration to reach convergence for each configuration. Finally, a relative global GDE measure is obtained by summing over all quadrangles' areas and dividing by the length $l_{\mathcal{M}}$ of \mathcal{M} :

$$\text{GDE}_{A_{\mathbf{q}}} := \frac{\sum_i A_{\mathbf{q}_i}}{l_{\mathcal{M}}}.$$

The change in the target curve from the exact geometry to the Mmesh is motivated by the fact that there is no explicit inverse of the parametrisation for the distorted ellipse. In the circle case, the inverse parametrisation was used in the implementation of the GDE calculation. Furthermore, on realistic test cases, a parametrisation of the geometry may not be available.

4.2.3. Numerical results

The number of Cmesh elements is set to $n_{\mathcal{C}}^e = 20$. The number of Mmesh elements $n_{\mathcal{M}}^e$ is varied from 198 to 7998. Figure 9(b) shows the linear Cmesh. Curving results with the H^1 modal and the least squares approach are shown in Figures 9(c) and 9(d), respectively. In Figures 9(e-j), the area around the tip of the ellipse is examined more closely for $q = 3$ (e-f) and for $q = 4$ (g-h), where Figures (e) and (g) show curves obtained with the H^1 modal and (f) and (h) with the least squares approach. The q -convergence curves of the two modal curving methods with these parameters are shown in Figure 10.




Figure 10: q -convergence curves of the GDE (distance to the target Mmesh) on the distorted ellipse for the H^1 (a) and least squares (b) curving methods with fixed $n_{\mathcal{C}}^e = 20$.

As observed on the circle test case, the curves for different Mmesh sizes coincide for small q values. Also, the stagnation levels are obtained for larger q values as the target Mmesh is refined. As expected for this more general geometry, odd and even order shape functions contribute to a similar extent.

The two modal curving methods converge as expected and show similar q -convergence rates. For the coarsest tested Mmeshes, the order restriction of the least squares model can be observed. In these cases, the number of Mmesh vertices is not high enough to robustly determine the α_i . Note that for
465 the coarser Mmeshes, the least squares seems to obtain less accurate results. However, it should be emphasized that the GDE is measured here with respect to the Mmesh.

5. Helmholtz scattering with modally curved meshes

The boundary meshes that resulted from the two proposed curving methods
470 are extended to 2D domain meshes by using the blending technique described in Section 3.4 and benchmarked on a Helmholtz scattering problem. The GDE-GIE relationship obtained in Section 2 is then tested against scattering data obtained on the modally curved circle and the distorted ellipse.

5.1. Scattering at the rigid cylinder (2D)

To benchmark the methods, the diffraction of a plane wave by a rigid cylinder,
475 presented in Section 2, is used.

Results

The numerical error follows the definition in Section 2.2. The results are depicted in Figure 11, along with reference results obtained on meshes that
480 were nodally curved using *Gmsh*, up to $q = 4$. The Figures show p -convergence curves for varying curving orders $q = 1, 2, \dots, 10$.

The results show the expected behaviour: (i) with increasing field order p the simulation error decreases until reaching the GIE and (ii) the GIE decays for increasing q values, in an irregular odd-even fashion (owing to the problem
485 symmetry). Results for both, $\text{MpC} = 100$ and $\text{MpC} = 5000$ are presented to examine the influence of the Mmesh accuracy. For the H^1 modal approach and $\text{MpC} = 100$, the stagnation levels of higher q orders fall together. As the refinement is increased to $\text{MpC} = 5000$, the curves separate and lower stagnation levels are achieved for higher q values. This matches the results of Section 4.1,
490 where it is shown that for coarse MpC -refinements, the inaccuracy of the target \mathcal{M} dominates the accuracy of the curving method and leads to higher GDE values. The corresponding plots with the least squares approach show that the MpC -refinement has no significant effect on the simulation accuracy. This also matches the GDE results of Section 4.1.

Figure 12 shows the numerical error for $p = 10$ for increasing frequencies for
495 each q -order. A highly refined target mesh is chosen with $\text{MpC} = 5000$. The reference nodally-curved *Gmsh* results are also shown (black marks).

The modal least squares and the reference behave more or less similarly. However, the H^1 modal approach leads to generally better simulation results.

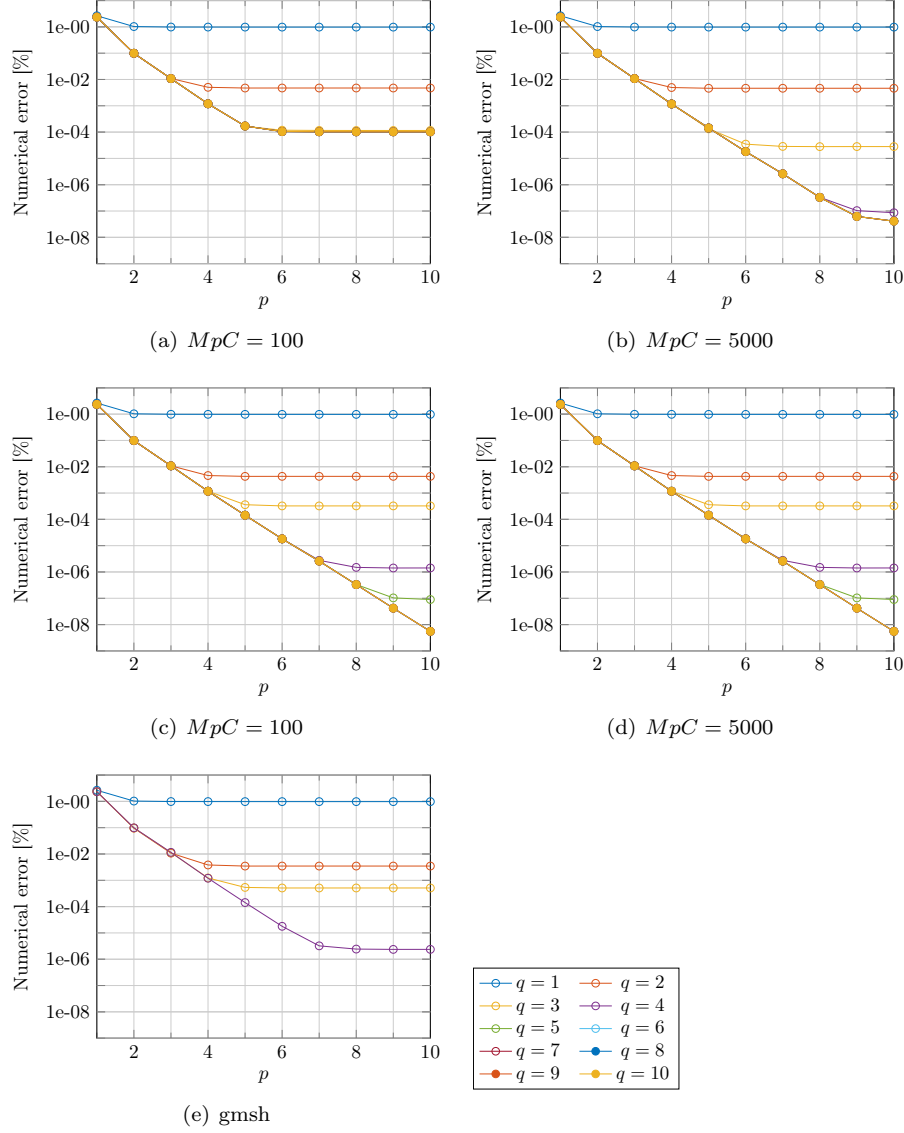


Figure 11: Comparing p -convergence curves for different MpC -values of the circle scattering test case with 12 Cmesh elements and $\omega = 1$ on meshes curved with (a)-(b) the H^1 modal approach, (c)-(d) the least squares curving, and (e) the nodal Gmsh results.

500 This is visible especially on the lower frequency range on cubic ($q = 3$) and quartic ($q = 4$) meshes.

We recall that the H^1 modal curving is based on the optimisation of the H^1 -seminorm which corresponds to the L^2 -norm of the derivatives. In the

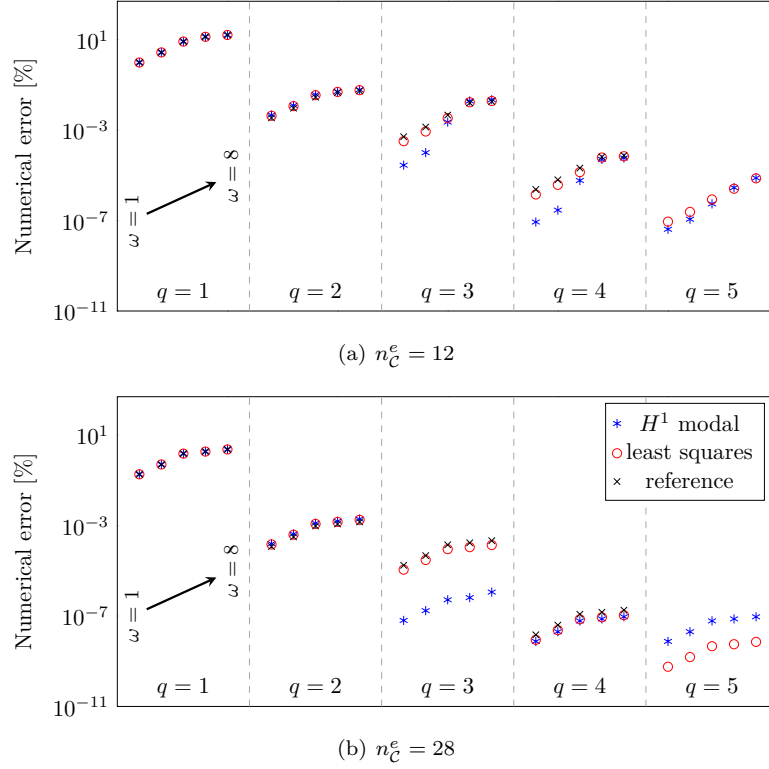


Figure 12: Numerical error at $p = 10$ with $q = 1, 2, \dots, 5$ and $\omega = 1, 2, 4, 6, 8$ for (a) $n_C^e = 12$ and (b) $n_C^e = 28$. H^1 modal (blue); least squares (red) and *Gmsh* reference (black).

simulations, the scattering direction is affected by the normal direction of the geometry. It is assumed that the reduced numerical error are caused by an optimised normal representation obtained with the continuous H^1 -optimisation.

This superiority is not observed anymore for $q \geq 5$, when the inaccuracy of the target Mmesh becomes the dominating factor. Further MpC-refinement, or better continuous reconstruction of the \mathcal{M} cloud of points is expected to overcome this stagnation and to allow even more improvement in the simulation results.

For high ω values, another effect prevails on the field error. Especially for $n_C^e = 12$, the p -convergence curves did not yet stagnate for the plotted $p = 10$. The geometric accuracy is not dominating the error and no comparison of the influence of the curving methods on the field error can be drawn from these data points.

GDE-GIE relationship

In Section 2.4.2, the GDE-GIE-relationship was only considered for simulations on nodally curved meshes obtained with *Gmsh*. Now the identified model

520 is tested for simulations on meshes curved with the H^1 and the least squares modal approaches described in Section 3. As before, the GIE is defined as the stagnation value in the q -convergence plots (see e.g. Figure 11).




Figure 13: Plot of the data set of the GIE \tilde{E}_R^2 against the GDE for the H^1 modal curving. The different colors represent the geometric order of the mesh, while the different shapes represent the ω value of the simulations.

Figure 13 shows the full data set for the H^1 modal curving. Most of the data points follow the expected linear behaviour. The data points for order $q = 3$ as well as some for $q = 4$ are standing out which correspond to the cases of improved simulation results discussed in the last Section. This indicates that the area based GDE measure does not cover all the information needed to explain the GIE-GDE relationship. It is assumed that the missing GDE component is linked to the deviation of the normal direction between the exact geometry and the approximation. As part of a future work, a possible measure would be the maximal or mean angle between the exact and the approximated normal, as described in [21, 22]. A drawback of these measures is their dependency on the curve parametrisations, as the angles are compared pointwise.

The remaining data for orders $q \in \{1, 2, 5, 6, 7, 8, 9, 10\}$ as well as the full data set of the least squares curving are tested with the model described in Eq. (3). Both data sets and the corresponding fits are shown in Figure 14. The model parameters β_i of the fits as well as the adjusted R^2 are presented in Table 2. The two parameters β_{geo} and β_{ω} that model the influence of the GDE and of the frequency ω are close to those obtained with the results on the nodally curved meshes. Compared to the nodal results, significantly smaller values are obtained for the the offset parameter β_{const} . This means that, neglecting the effects for $q = 3$ and $q = 4$ for the H^1 modal curving, the modally curved meshes lead to better simulation results for the same area-based GDE measure.

5.2. Scattering on the distorted ellipse

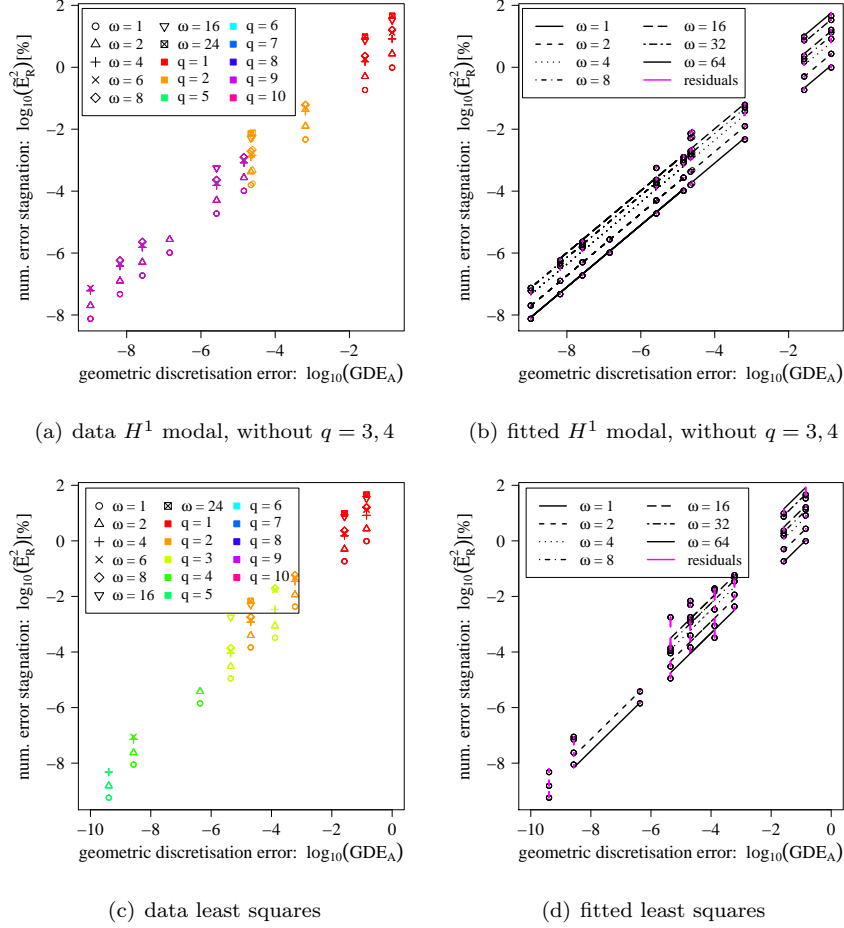


Figure 14: Plots of the data and fitting of the GIE \tilde{E}_R^2 against the GDE for (a)-(b) the H^1 modal curving, without data for $q = 3, 4$ and (c)-(d) the least squares meshes for $q = 1, \dots, 10$. In (a) and (c), the color decodes the geometric order of the mesh, while the shapes decode the ω value of the simulations.

	β_{const}	β_{geo}	β_{ω}	\tilde{R}^2
nodal	1.3846	1.0106	1.2314	0.9975
H^1 modal	0.9106	1.0012	1.2287	0.9994
least squares modal	0.9175	1.0594	1.3802	0.9965

Table 2: Parameter values of the fits of model (3) to the numerical error evaluation and GDE definitions for simulations on directly modally curved meshes of orders $q \in \{1, 2, 5, 6, 7, 8, 9, 10\}$, for simulations on meshes curved with the modal least squares approach of orders $q \in \{1, 2, 3, \dots, 10\}$, and the corresponding values of the adjusted \tilde{R}^2 . The parameter values obtained for the nodally curved meshes with $q \in \{1, 2, 3, 4\}$ are provided as reference.

545 The GDE-GIE relationship is now examined when solving a Helmholtz scattering problem on the distorted ellipse. This will allow to determine whether the previously drawn conclusions also apply to more complex cases.

Test case and reference solution

550 The domain Ω chosen for the simulation is shown in Figure 15, with $r_x = 2$ and $r_y = 7$. The incident field is defined as a plane wave $\mathbf{p}_i = \exp(-ikx)$ with wavenumber k , travelling in positive x -direction. The numerical set-up is very similar to the one proposed for the rigid cylinder scattering, except that, in the absence of analytical solution, the reference solution is obtained from an overkilled simulation. The latter is obtained as follows: a highly refined 555 quadratic mesh of size $h = 10^{-3}$ is first generated on the boundary to minimize the geometric error. The mesh size is then relaxed towards the domain outer surface, where a 10-element layered Perfectly Matched Layer (PML) [29] is generated, in order to absorb all outgoing waves with minimal spurious reflections. For optimal efficiency, the element order is adjusted locally using the a priori 560 error indicator described in [7], with a target L^2 -error of $10^{-8}\%$. The reference solution obtained at $\omega = 4$ and $\omega = 16$ is shown in Figures 15(b) and 15(c) respectively. The field error is evaluated as the relative L^2 -error with respect to this reference solution along an ellipse with $r_x = 1.25$ and $r_y = 6$.

GDE-GIE relationship

565 The data and the fitted model are shown in Figure 16. Table 3 provides the model parameters.

	β_{const}	β_{geo}	β_{ω}	$\tilde{\mathbf{R}}^2$
H^1 modal	0.85762	1.0916	1.3861	0.9783

Table 3: Parameter values obtained for the fitted GDE-GIE model (3) to the data of the scattering at the distorted ellipse, and the corresponding value of the adjusted \mathbf{R}^2 .

The model fit is not as good as it was for the circular geometry, however the general model trend is still observed, including the sup-linear dependency in ω . This further confirms that the evaluation of the area based GDE towards the 570 Mmesh is a reasonable choice, provided that the Mmesh is sufficiently refined.

Note that from $q = 3$ to $q = 4$, a drop of the field error is observed in the data. This is similar to the drop in the data for $q = 3, 4$ in Figure 13 and could therefore coincide with the already described area-based GDE model limitation. The missing aspect of the normal approximation becomes more 575 significant for this more intricate geometry with increased complexity in the scattering structure.

6. Extension to 3D

In this paper, the algorithms were described and tested only for 2D geometries. The main steps for the extension of the methods to 3D are explained in




Figure 15: (a) Sketch of the domain for the scattering test case and reference solutions for (b) $\omega = 4$ and (c) $\omega = 16$.

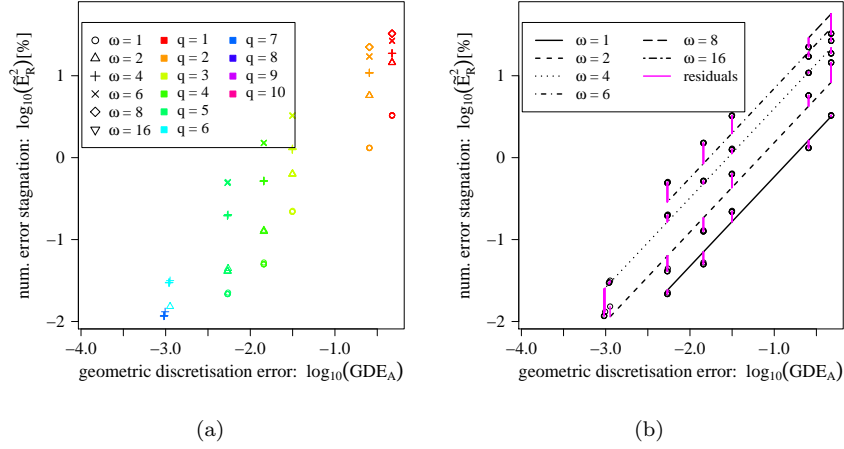


Figure 16: Plot of (a) the data set of the GIE \tilde{E}_R^2 against the GDE for the H^1 modal curving of the distorted ellipse test case. The different colors represent the geometric order of the mesh, while the different shapes represent the ω value of the simulations. In (b), the corresponding model fit and residuals are shown.




Figure 17: Cmesh, Mmesh, and curved mesh of a cylinder. In red, the feature lines are indicated.

580 this section. For a 3D geometry, the boundary forms a surface which we assume to be linearly triangulated. As in the 2D case we consider a coarse linear Cmesh and a fine linear Mmesh as curving target, where the Cmesh vertices form a subset of the Mmesh vertices (see Fig. 17).

585 Again, the advantageous property of the modal Lobatto shape functions can be exploited. In 3D their approximation basis is organized into vertex, edge, face and bubble shape functions [15]. Edge shape functions contribute only on the local edge and vanish on the neighbouring edges. Similarly, all face shape functions contribute only on the face interior and vanish on the element vertices and edges. Finally, bubble shape functions are only active inside the volume elements and vanish on the element boundaries. These properties of the shape functions can be exploited to determine the curving coefficients. The edges are curved first, then the face interior shape functions are added to the boundary

590

element deformation. Finally the bubble shape functions contribution is taken into account through the blending.

595 In 2D, the Mmesh arc over a single Cmesh element is bounded by the \mathcal{M} vertices coinciding with the \mathcal{C} vertices. In 3D, the edge target curves for the Cmesh edges do in general not fall together with Mmesh edges. Instead, they are defined as the intersection between the plane in Cmesh edge normal direction and the triangles in the Mmesh. The resulting target curve is still piecewise
600 linear and the edge curving itself follows the 2D algorithms.

At this step, the treatment of feature lines can be introduced. A feature line is a contour line on the geometry that may not be smoothed out during the curving step. In Figure 17, the feature lines are indicated in red. We assume that edges that fall on a feature line are marked as such on both \mathcal{M} and \mathcal{C} . Then,
605 the target of a \mathcal{C} edge on a feature line is simply defined by the corresponding Mmesh vertices on the feature line.

For the curving of the triangular face, all Mmesh vertices are considered that fall between the edge target curves of the edges bounding the face. The resulting target surface is a piecewise linear, triangulated surface. To find a
610 parametrisation of the triangular face, mean value coordinates (see [30]) are used. They are length ratio and angle ratio preserving and correspond to the idea of length-based parametrisation in 2D.

As previously mentioned, the δ -property does not hold for the standard triangle face shape functions as they are introduced in [15]. This means that
615 no explicit formula can be found for the corresponding face α values in the H^1 optimisation-based algorithms used in the H^1 modal approach. Instead, a linear system has to be solved for the face curving. Alternatively, orthogonalised shape functions may be used [26].

7. Summary, conclusions, and future work

620 In the first part of this paper, the influence of geometry representation on the accuracy of a simple Helmholtz exterior problem was examined. Based on intensive numerical calculations, a statistical model was provided which relates the geometry induced numerical errors (GIE) to the corresponding area-based geometric discretisation error (GDE). This model inferred that the GIE varies
625 (i) linearly with the GDE and (ii) sup-linearly with the frequency ω .

In a second part, two 2D curving modal methods were proposed: a least squares approach and a H^1 -seminorm optimisation (called H^1 modal method). Unlike the usual nodal high-order (Lagrange) methods, the proposed algorithms are based on the approximation of the curve using hierarchical, modal shape
630 functions. A blending method is also described which allows to propagate the curved boundaries towards the domain discretization seamlessly (i.e. without requiring to run any optimization)

The two methods were compared. The main drawback of the H^1 modal approach relates to the stagnation due to the Mmesh accuracy. A preprocessing
635 step running a mesh reconstruction method for the Mmesh before the curving

could improve this, similar to what is proposed in [10]. The advantage of least squares, using only exact points, is clearly observed for the circle test case. It becomes less distinct though on the more complicated geometry. Also, if the Mmesh becomes too coarse, the least squares approach may become unstable. Note that the here considered least squares method is only a basic approach. There exist techniques to potentially avoid instabilities, such as applying Legendre polynomials with constraints for the polynomial interpolation. Another one would be to use optimised node distributions (e.g. Gauss-Lobatto-Legendre nodes) by allowing inexact target points along the interpolated Mmesh. Since the H^1 curving approach bypasses the instabilities as well, further optimizations on the least squares approach were not considered and the focus was shifted to the H^1 -seminorm optimisation.

A general improved level of numerical error for a given geometric accuracy is observed for the modally curved meshes in comparison with the nodally curved meshes. Besides this, the GIE further drops for simulations on modally curved meshes. It is assumed that this is linked to the improved description of the normal direction that is caused by the H^1 -seminorm optimisation.

Under the assumption of a highly refined target mesh, the H^1 modal method is considered the most promising approach. It shows the expected stagnation of the GDE due to the inaccuracy in the target curve, but otherwise converges with the best obtained order. Furthermore, it does not show instabilities for coarse Mmeshes.

In this paper, the aspect of mesh validity and quality was not considered. When the element edge (in 2D) or edges and faces (in 3D) on the boundary are curved, tangling can occur, leading to invalid elements. Many mesh untangling and smoothing algorithms are available in the literature, see e.g. [9, 31, 32]. For the described modal approaches, an optimisation-based algorithm could be applied, as suggested in [33, 34], in order to find optimal coefficients to curve the internal edges (2D) or edges/faces (3D). For an enhanced control on the element quality, coefficients for element interior shape functions (i.e. bubble shape functions) should be included.

As future work, the evaluation of the GDE-GIE relationship should be enhanced with other types of the measuring the GDE. Especially the combination of the area-based measure with normal- or curvature based measures seems most promising from the current point of view. Future work further involves the extension of the H^1 modal method to 3D, including the treatment of feature lines and the implementation of efficient untangling strategies. Finally, as modal curving methods allow to adjust the order independently in each element, geometrical order (q -)adaptive refinement strategies will be examined.

Acknowledgements

This work was performed as part of the CRANE project (Community and Ramp Aircraft Noise, www.crane-eid.eu, GA: 606844) funded by the European Union under the Framework Programme 7. The authors acknowledge

the use of the IRIDIS High Performance Computing Facility, and associated
 support services at the University of Southampton, in the completion of this
 work The authors also gratefully acknowledge VLAIO (Flanders Innovation &
 Entrepreneurship) for their support of the EUREKA project HEAVENS.

References

- [1] R. Sevilla, S. Fernández-Méndez, A. Huerta, NURBS-enhanced finite ele-
 ment method (NEFEM), International Journal for Numerical Methods in
 Engineering 76 (1) (2008) 56–83.
- [2] F. Bassi, S. Rebay, High-order accurate discontinuous finite element solu-
 tion of the 2D Euler equations, Journal of Computational Physics 138 (2)
 (1997) 251 – 285.
- [3] T. Toulorge, W. Desmet, Curved boundary treatments for the discontin-
 uous Galerkin method applied to aeroacoustic propagation, AIAA journal
 48 (2) (2010) 479–489.
- [4] S. Dey, M. S. Shephard, J. E. Flaherty, Geometry representation issues
 associated with p-version finite element computations, Computer Methods
 in Applied Mechanics and Engineering 150 (1) (1997) 39–55.
- [5] X. Luo, M. S. Shephard, J.-F. Remacle, The influence of geometric ap-
 proximation on the accuracy of high order methods, Rensselaer SCOREC
 report 1.
- [6] A. Lieu, G. Gabard, H. Bériot, A comparison of high-order polynomial and
 wave-based methods for Helmholtz problems, Journal of Computational
 Physics 321 (2016) 105–125.
- [7] H. Bériot, A. Prinn, G. Gabard, Efficient implementation of high-order
 finite elements for Helmholtz problems, International Journal for Numerical
 Methods in Engineering 106 (3) (2016) 213–240.
- [8] Z. Xie, R. Sevilla, O. Hassan, K. Morgan, The generation of arbitrary order
 curved meshes for 3D finite element analysis, Computational Mechanics
 51 (3) (2013) 361–374. doi:10.1007/s00466-012-0736-4.
- [9] J.-F. Remacle, J. Lambrechts, C. Geuzaine, T. Toulorge, Optimizing the
 geometrical accuracy of 2D curvilinear meshes, Procedia Engineering 82
 (2014) 228–239.
- [10] K. Bock, J. Stiller, Generation of high-order polynomial patches from scat-
 tered data, in: Spectral and High Order Methods for Partial Differential
 Equations-ICOSAHOM 2012, Springer, 2014, pp. 157–167.
- [11] X. Jiao, N. R. Bayyana, Identification of c^1 and c^2 discontinuities for surface
 meshes in cad, Computer-Aided Design 40 (2) (2008) 160–175.

- [12] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities, *International Journal for Numerical Methods in Engineering* 79 (11) (2009) 1309–1331.
- 720 [13] D. P. Sanjaya, K. J. Fidkowski, Improving high-order finite element approximation through geometrical warping, in: *AIAA 22nd Aeroacoustics Conference*, Dallas, Texas, USA, 2015, aIAA Paper 2015–2605.
- [14] E. Ruiz-Gironés, J. Sarrate, X. Roca, Defining an 2-disparity measure to check and improve the geometric accuracy of non-interpolating curved high-order meshes, *Procedia Engineering* 124 (2015) 122–134.
- 725 [15] P. Solin, K. Segeth, I. Dolezel, *Higher-order finite element methods*, CRC Press, 2003.
- [16] V. Ziel, H. Bériot, O. Atak, G. Gabard, Comparison of 2D boundary curving methods with modal shape functions and a piecewise linear target mesh, *Procedia Engineering* 203.
- 730 [17] H. Bériot, E. Perrey-Debain, M. B. Tahar, C. Vayssade, Plane wave basis in galerkin bem for bidimensional wave scattering, *Engineering Analysis with Boundary Elements* 34 (2) (2010) 130–143.
- [18] C. Geuzaine, A. Johnen, J. Lambrechts, J.-F. Remacle, T. Toulorge, The generation of valid curvilinear meshes, in: *IDIHOM: industrialization of high-order methods-a top-down approach*, Springer, 2015, pp. 15–39.
- 735 [19] T. Toulorge, C. Geuzaine, J.-F. Remacle, J. Lambrechts, Robust untangling of curvilinear meshes, *Journal of Computational Physics* 254 (0) (2013) 8 – 26.
- [20] H. Alt, B. Behrends, J. Blömer, Approximate matching of polygonal shapes, *Annals of Mathematics and Artificial Intelligence* 13 (3-4) (1995) 251–265.
- 740 [21] M. Aliabadi, W. Hall, Nonisoparametric formulations for the three-dimensional boundary element method, *Engineering Analysis* 5 (4) (1988) 198–204.
- 745 [22] A. M. McIvor, R. J. Valkenburg, A comparison of local surface geometry estimation methods, *Machine Vision and Applications* 10 (1) (1997) 17–26.
- [23] X. Luo, M. S. Shephard, J.-F. Remacle, R. M. O’Bara, M. W. Beall, B. Szabó, R. Actis, p-version mesh generation issues., in: *Proceedings of the 11th International Meshing Roundtable*, 2002, pp. 343–354.
- 750 [24] R Core Team, *R / A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria (2013).
URL <http://www.R-project.org>

- [25] L. Fahrmeir, T. Kneib, S. Lang, B. Marx, Regression: models, methods and applications, Springer Science & Business Media, 2013.
- 755 [26] P. Šolín, T. Vejchodský, M. Zítka, F. Ávila, Imposing orthogonality to hierarchic higher-order finite elements, *Mathematics and Computers in Simulation* 76 (1) (2007) 211–217.
- [27] J. S. Hesthaven, T. Warburton, Nodal discontinuous Galerkin methods: algorithms, analysis, and applications, Springer Science & Business Media, 760 2007.
- [28] V. Schmid, H. Bériot, O. Atak, G. Gabard, High-order curved mesh generation by using a fine linear target mesh, in: *ECCOMAS Congress 2016 Proceedings*, Vol. 1, ECCOMAS, 2016, pp. 493–503.
- 765 [29] A. Bermúdez, L. Hervella-Nieto, A. Prieto, R. Rodri, et al., An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems, *Journal of Computational Physics* 223 (2) (2007) 469–488.
- [30] M. S. Floater, Mean value coordinates, *Computer aided geometric design* 20 (1) (2003) 19–27.
- 770 [31] D. Moxey, D. Ekelschot, Ü. Keskin, S. Sherwin, J. Peiró, A thermo-elastic analogy for high-order curvilinear meshing with control of mesh validity and quality, *Procedia Engineering* 82 (2014) 127–135.
- [32] X. Luo, M. Shephard, L.-Q. Lee, C. Ng, L. Ge, Tracking adaptive moving mesh refinements in 3D curved domains for large-scale higher order 775 finite element simulations, in: R. Garimella (Ed.), *Proceedings of the 17th International Meshing Roundtable*, Springer Berlin Heidelberg, 2008, pp. 585–601.
- [33] A. Gargallo-Peiró, X. Roca, J. Peraire, J. Sarrate, Defining quality measures for validation and generation of high-order tetrahedral meshes, in: 780 J. Sarrate, M. Staten (Eds.), *Proceedings of the 22nd International Meshing Roundtable*, Springer International Publishing, 2014, pp. 109–126.
- [34] A. Gargallo-Peiró, X. Roca, J. Peraire, J. Sarrate, Optimization of a regularized distortion measure to generate curved high-order unstructured tetrahedral meshes, *International Journal for Numerical Methods in Engineering* 785 103 (5) (2015) 342–363.

Vitae

Verena Ziel (formerly Schmid) studied Applied Mathematics at the Technische Universität Dresden, Germany. She received her Diploma in 2014 on the simulation of hexagonal structures on curved surfaces. From 2014 to 2017 she

790 was a fellow of the Marie Skłodowska-Curie project CRANE. In this scheme
she worked in the acoustics research group at Siemens Industry Software NV
in Leuven, Belgium and later at the Institute of Sound and Vibration Research
at the University of Southampton, UK. Since 2014 she is a PhD student at
the University of Southampton, working on mesh curving algorithms and their
795 application in Helmholtz simulations.

Hadrien Bériot received his PhD degree in Mechanical Engineering from
University of Compiègne, France in 2008, where he worked on High-order meth-
ods for aeroacoustics propagation. His current research interests include High-
order finite element methods, Perfectly Matched Layers, Domain Decomposition
800 Methods and flow induced noise generation and propagation. He is currently
research manager at Siemens Industry Software NV., where he is responsible for
the technology monitoring and prototyping supporting the development of the
leading commercial software Virtual.Lab and SimCenter Acoustics.

Onur Atak is a RTD Project Leader in the Acoustics research team at
805 Siemens Industry Software NV. He received his MSc degree in 2009 from the
Sound and Vibration Masters programme at Chalmers University of Technology.
He received his PhD degree from KU Leuven in 2014, where he worked on
Wave Based Modeling Methods for acoustic inclusion and multiple scattering
problems in the mid-frequency range. His current research interests are Fast
810 Boundary Element Methods, high order and adaptive methods, IsoGeometric
analysis, high order surface reconstruction, inverse material characterization,
model order reduction schemes etc.

Gwénaél Gabard is a senior researcher at the acoustics lab LAUM at
the University of Le Mans (France). He previously worked as an associate
815 professor at the Institute of Sound and Vibration Research at the University of
Southampton (UK). He is currently the holder of an Industrial Chair, funded by
the Safran Group, which is focused on the development of innovative acoustic
treatments for aerospace applications. More generally, his areas of expertise
include aero-acoustics, computational methods and fluid mechanics.