

Efficient implementation of high-order finite elements for Helmholtz problems

Hadrien Bériot^{1,*}, Albert Prinn² and Gwénaél Gabard²

¹*Siemens Industry Software, Simulation and Test Solutions, Interleuvenlaan 68, B-3001, Heverlee, Belgium*

²*Institute of Sound and Vibration Research, University of Southampton, SO17 1BJ, UK*

SUMMARY

Computational modeling remains key to the acoustic design of various applications, but it is constrained by the cost of solving large Helmholtz problems at high frequencies. This paper presents an efficient implementation of the high-order finite element method (FEM) for tackling large-scale engineering problems arising in acoustics. A key feature of the proposed method is the ability to select automatically the order of interpolation in each element so as to obtain a target accuracy while minimizing the cost. This is achieved using a simple local *a priori* error indicator. For simulations involving several frequencies, the use of hierarchical shape functions leads to an efficient strategy to accelerate the assembly of the finite element model. The intrinsic performance of the high-order FEM for 3D Helmholtz problem is assessed, and an error indicator is devised to select the polynomial order in each element. A realistic 3D application is presented in detail to demonstrate the reduction in computational costs and the robustness of the *a priori* error indicator. For this test case, the proposed method accelerates the simulation by an order of magnitude and requires less than a quarter of the memory needed by the standard FEM. © 2016 The Authors. *International Journal for Numerical Methods in Engineering* published by John Wiley & Sons Ltd.

Received 18 June 2015; Revised 28 September 2015; Accepted 6 November 2015

KEY WORDS: acoustics; high-order FEM; pFEM; Helmholtz problems; frequency sweeps

1. INTRODUCTION

In many engineering applications, solutions of the Helmholtz equation are required over a wide frequency range. For acoustics applications, one is often interested in modeling the sound field over the full audible frequency range, i.e. from 20 Hz to 20 kHz. The nature of the solution and the requirements in terms of mesh resolution vary drastically over this range. Ideally, a separate mesh should be prepared for each frequency of interest so as to obtain an accurate solution while using the coarsest mesh possible. In practice, this is rarely performed as this is more demanding for the end user. Instead, it is common to use a single mesh, designed for the highest frequency of interest, to cover a large frequency range, despite the large increase in computational cost incurred with this approach. These practical considerations are compounded by more fundamental limitations of the conventional finite element methods at mid to high frequencies. In particular, the pollution effect (which can be described as the cumulative build-up of dispersion error over the computational domain) leads to a rapid increase in numerical error at high frequencies [1]. This enhances even further the differences in mesh resolutions required for low and high frequencies.

A range of techniques have been proposed to reduce the pollution error inherent to the finite element method. Stabilized finite elements [2, 3] have been successfully applied to a number of

*Correspondence to: Hadrien Bériot, Siemens Industry Software, Simulation and Test Solutions, Interleuvenlaan 68, B-3001, Heverlee, Belgium.

†E-mail: hadrien.beriot@siemens.com

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

problems, but their performance decreases significantly on unstructured grids. Another alternative consists in embedding the oscillatory nature of the solution in the approximating functions, like in the partition of unity method [4], the discontinuous enrichment method [5], and the ultraweak variational formulation [6]. For a review of these methods and a comparison of their performance, see [7] and [8]. Although these methods have been shown to reduce the dispersion error, they still suffer from a number of drawbacks. Firstly, they are inherently less adapted to problems with inhomogeneous media, likely to be encountered in many applications. Secondly, the design of error estimators in view of their usage in a p -adaptive or hp -adaptive context remains an open question.

Another alternative to reduce the pollution error is to increase the order of interpolation of the polynomial basis. Higher-order approximations lead to significantly reduced resolution requirements, which in turn result in smaller numbers of degrees of freedom for a specific Helmholtz problem. In [9], several polynomial bases commonly used in spectral and FE methods (Lagrange, integrated Legendre, and Bernstein shape functions) are compared. High-order polynomial approximations are shown to efficiently control the pollution effect and lead to more efficient iterative calculations. However, the link between the number of degrees of freedom and the computational cost for large-scale problems is not so direct as p -FEM models typically result in larger matrix bandwidth. Defining the optimal element order p that minimizes the computational effort required to achieve a given accuracy on a given problem is therefore not trivial. This issue is examined by Vos *et al.* for different equations including the Helmholtz operator [10]. For two-dimensional problems, results indicate that the optimal performance for a given error tolerance is typically obtained with an order $p = 6$. In [11], three-dimensional Helmholtz problems are considered, but only the evaluation of the system matrix is monitored and not the factorization of the global system matrix, which remains the most computationally intensive operation. Their study indicates that the benefits of high-order approximations are even greater in three dimensions. In [12], Giorgioni *et al.* compare the performance of high-order continuous and discontinuous Galerkin methods for various 2D smooth and non-smooth scattering problems. The computational cost is assessed in terms of CPU time for solving the linear system, and it is shown that high-order methods exhibit an efficiency orders of magnitude above that of classical linear approximations. In [13], Huerta *et al.* perform a theoretical complexity analysis of the operations involved in solving continuous and discontinuous Galerkin models. This study is carried out for a constant accuracy, based on a simple one-dimensional error estimator. They conclude that for an engineering accuracy, high-order methods are superior for all major operations, that is, for the creation of the element matrices as well as for the cost of solving the global linear system. The superiority of continuous approximations over discontinuous approximations is also clearly established.

High-order finite elements form the basis for the development of hp -adaptive methods where both the mesh and the polynomial order are concurrently adjusted automatically to achieve a prescribed level of error. This approach is motivated by the fact that a suitable combination of spatial refinement (h -adaptivity) and order refinement (p -adaptivity) can provide exponential convergence on a class of problems where h -FEM or p -FEM methods achieve only algebraic convergence [14]. Automatic hp -adaptive computations can be viewed as an iterative procedure that generates a sequence of converging solutions rather than a unique approximation. The performance of these methods is largely dependent on the design of accurate *a posteriori* error estimates, which are used to guide the mesh refinements and the selection of polynomial orders. The most common error estimates are residual-based or recovery-based [15], and the design of accurate *a posteriori* error estimators for Helmholtz problems remains challenging [16]. In particular, these local estimators do not incorporate the pollution effect, and as a result, their accuracy deteriorates as the frequency increases [17]. hp -adaptivity also relies on robust, automated mesh generation. While this is a subject of active ongoing research [18–20], the complete integration of the mesh generator with a finite-element solver would require a complete overhaul of the architecture of existing industrial codes [21].

The objective of this paper is to propose an efficient implementation of the p -FEM approach for large-scale three-dimensional acoustic calculations required over a range of frequencies. A key aspect in this approach is the design of a robust *a priori* error indicator to select the polynomial order in each element. This allows to use a single mesh for multiple calculations over a wide range of frequencies. It was also shown how the use of hierarchical shape functions is important to accelerate

the assembly process by reducing the cost of evaluating large element matrices associated with high-order elements. This combination leads to an efficient numerical method able to provide accurate solutions with significantly reduced memory and time. The aim of this study is to not only to assess the intrinsic computational performance of the p -FEM but also to discuss how this method can be used efficiently in an industrial context for large-scale acoustic simulations.

The present paper is structured as follows. Section 2 introduces the Helmholtz finite element formulation as well as the high-order FEM approximation basis. The proposed implementation of the high-order FEM for tackling multi-frequency Helmholtz problems is then presented. In Section 3, the computational performance of p -FEM for three-dimensional Helmholtz problems is assessed in terms of polynomial order and element size. In Section 4, the *a priori* error indicator is presented, and its performance is assessed with uniform or strongly non-uniform meshes. Finally, in Section 5, the benefits of the proposed approach are demonstrated with an industrial test case, while Section 6 summarizes the main findings in this paper.

2. HIGH-ORDER FINITE ELEMENTS

2.1. Problem description

We consider the propagation of harmonic, linear sound waves in a uniform, quiescent medium within a bounded domain Ω . Throughout this paper, an implicit time dependence $e^{+i\omega t}$ is assumed, where ω is the angular frequency. The complex amplitude u of the acoustic pressure is governed by the Helmholtz equation:

$$\nabla^2 u + k^2 u = -s \quad \text{in } \Omega, \quad (1)$$

where $k = \omega/c$ is the acoustic wavenumber (with c the sound speed) and s is an arbitrary volume source term.

A generic Robin condition is applied on the boundary $\Gamma = \partial\Omega$ of the domain

$$\nabla u \cdot \mathbf{n} + \gamma u = g \quad \text{in } \Gamma, \quad (2)$$

where \mathbf{n} is the unit normal to Γ pointing out of the domain Ω . The coefficient $\gamma(\omega)$ can be used to represent an acoustic treatment and $g(\omega)$ an imposed normal velocity applied on the boundaries.

This algebraic system is typically solved at a large number N_f of individual frequencies, sorted in ascending order over a frequency range

$$\mathcal{F} = \{f_1, f_2, \dots, f_{N_f}\}, \quad (3)$$

and for a large number N_{rhs} of right-hand sides.

Using the Galerkin method, Equations (1) and (2) can be combined into a weak variational formulation. It consists in finding the solution $u \in H^1(\Omega)$ such that

$$K(v, u) - k^2 M(v, u) + \gamma(\omega) M_\Gamma(v, u) = S(v) + S_\Gamma(v), \quad \forall v \in H^1(\Omega), \quad (4)$$

with the sesquilinear forms

$$M(v, u) = \int_{\Omega} \bar{v} u \, d\Omega, \quad K(v, u) = \int_{\Omega} \nabla \bar{v} \cdot \nabla u \, d\Omega, \quad M_\Gamma(v, u) = \int_{\Gamma} \bar{v} u \, d\Gamma, \quad (5)$$

and the source terms

$$S(v) = \int_{\Omega} \bar{v} s \, d\Omega, \quad S_\Gamma(v) = \int_{\Gamma} \bar{v} g \, d\Gamma, \quad (6)$$

where $\bar{\cdot}$ denotes the complex conjugate.

To define a finite element approximation of the variational formulation (4), the domain Ω is partitioned into a set of non-overlapping finite elements Ω_e of typical size h . The approximate solution u^h and the associated test functions v^h are constructed using the classical H^1 -conforming p -FEM with hierarchical shape functions (see Šolín *et al.* [22] for details).

2.2. Lobatto shape functions

In one dimension, the reference element is defined as $E^L = \{\xi \in \mathbb{R}; -1 < \xi < 1\}$. The first-order $p = 1$ only involves the standard linear nodal shape functions:

$$l_0(\xi) = \frac{1-\xi}{2}, \quad l_1(\xi) = \frac{1+\xi}{2}. \quad (7)$$

For higher orders $p > 1$, the basis is complemented by $p - 1$ shape functions given by the Lobatto functions (also called integrated Legendre polynomials):

$$l_i(\xi) = \sqrt{i - \frac{1}{2}} \int_{-1}^{\xi} L_{i-1}(s) ds, \quad \text{with } i = 2, 3, \dots, p, \quad (8)$$

where $L_i(\xi)$ are the Legendre polynomials of order i [23].

An important property of the Lobatto functions is the following orthogonality property:

$$\int_{-1}^1 \frac{dl_i}{d\xi} \frac{dl_j}{d\xi} d\xi = \delta_{ij}, \quad \text{with } i, j = 0, 1, \dots, p, \quad \text{and } i \text{ or } j \geq 2, \quad (9)$$

which ensures an optimal conditioning of the stiffness matrix for Helmholtz problems [24].

Another crucial property of the functions $l_i(\xi)$ is that they form a hierarchical basis, meaning that the basis used at order p is fully retained in the basis of order $p + 1$. In other words, one can construct the basis at order $p + 1$ by adding more functions to the basis of order p rather than defining a completely different set of basis functions. An important consequence is that this basis naturally allows for local p -refinement because two adjacent elements using different interpolation orders can coexist (Section 2.4). Standard Lagrangian polynomial shape functions, commonly used for low-order finite element methods, do not offer such flexibility. Finally, the hierarchical nature of this basis can also be used to accelerate the calculation of the system matrices when solving the problem repeatedly over a range of frequencies (this will be discussed in more detail in Section 2.4).

On each element Ω_e , the solution is expanded on the shape functions ϕ_n^e , as follows

$$u^h(\mathbf{x}) = \sum_{n=1}^{N_{\text{dof}}^e} \phi_n^e(\mathbf{x}) u_n^e = \Phi_e(\mathbf{x})^T \mathbf{u}_e, \quad (10)$$

where u_n^e represents the contribution of the n -th shape function to the local solution and N_{dof}^e denotes the total number of shape functions contributing on the element e .

In the usual way, each element is mapped onto a reference element with coordinates ξ , and the shape functions $\phi_n^e(\xi)$ are defined on this reference element. The construction of the Lobatto shape functions for different types of elements is well described in the literature [22, 25] and will not be repeated here. The discussion that follows will take the example of the tetrahedral element as it will be used in the subsequent sections.

The tetrahedral shape functions ϕ_n^e are constructed through a specific generalized coordinate mapping, which matches the corresponding hexahedral shapes, formed by a tensor product of the Lobatto shape functions $l_i(\xi)$ [26]. These functions can be classified into four categories, namely, vertex, edge, face, and bubble functions [22]. In the simple situation where the polynomial order p is the same for all elements, it is easy to show that the number of degrees of freedom contributing to the solution on any given element e is given by

$$N_{\text{dof}}^e = 4 + \sum_{i=1}^6 (p-1) + \frac{1}{2} \sum_{i=1}^4 (p-2)(p-1) + \frac{1}{6} (p-3)(p-2)(p-1), \quad \forall p \geq 1, \quad (11)$$

where the four terms correspond to the numbers of nodal, face, edge, and bubble functions, respectively.

In practice, as will be discussed in Section 5, it is often necessary to adjust the polynomial order over the mesh (p -refinement). A different order p_e is therefore associated with each element e (the actual process to select a suitable order p_e will be discussed in detail in Section 4). Due to

its hierarchical structure, the p -FEM can easily handle local order variations, meaning that two conforming adjacent elements can coexist with different polynomial orders, without requiring to modify the local approximation basis and without violating the H^1 conformity rule. From the order p_e defined for each elements, one can assign a unique interpolation order p_l to each edge and p_s to each face of the mesh, based on a given conformity rule. In this paper, we use the maximum rule, which states that the polynomial order assigned to a given face or edge in the physical mesh is equal to the maximum order of all adjacent elements sharing this face or edge. This rule guarantees that the polynomial order of any shape function associated with an element e verifies its order of approximation. Numerical experiments suggest that this rule leads to robust approximations in the presence of highly non-uniform meshes (Section 4). The number of shape functions on a given element e is

$$N_{\text{dof}}^e = 4 + \sum_{i=1}^6 (p_{l_i}^e - 1) + \frac{1}{2} \sum_{i=1}^4 (p_{s_i}^e - 2) (p_{s_i}^e - 1) + \frac{1}{6} (p_e - 3)(p_e - 2)(p_e - 1), \quad (12)$$

where again the four terms correspond to the numbers of nodal, face, edge, and bubble functions. $p_{s_i}^e$ (respectively $p_{l_i}^e$) denotes the order of the i -th face (resp. edge) of element e . Note that the case $p_e = p_{s_i} = p_{l_i} = 1$ corresponds to the linear finite elements, with $N_{\text{dof}}^e = 4$.

The bubble functions are purely internal, i.e. they vanish on the element boundaries and are therefore not coupled to the neighboring elements.

Static condensation consists in factoring out all bubble functions from the discrete system during the assembly. Bubble functions only affect the solution locally as they are not shared between elements. The corresponding DOFs are said to be internal, whereas the DOFs associated with the nodal, edge, and face functions are external because they contribute to the solution in the neighboring elements. These internal degrees of freedom can be statically condensed out at the assembly level, with the objective to remove them from the global system matrix. This procedure does not affect the final solution, but it can reduce significantly the size of the discrete problem to solve. The internal values of the solution can be recovered during the post-processing stage by solving element-wise local linear problems. This procedure is a very efficient tool to reduce the size and improve the condition number of hp finite element system matrices [22].

The number of internal degrees of freedom of a tetrahedral element of order p_e is equal to

$$\frac{1}{6} (p_e - 3)(p_e - 2)(p_e - 1), \quad \forall p_e \geq 3. \quad (13)$$

Hence, linear and quadratic tetrahedral elements do not involve any internal DOFs and cannot benefit from the static condensation. Besides, the effect of the condensation becomes more significant when the polynomial order becomes high because the number of bubble functions is a cubic function of the order. In the remainder of this paper, all computations are performed using static condensation of the bubble shape functions.

2.3. Standard solution procedure

The discretization of the weak form (4) yields the following algebraic system:

$$\mathbf{A}(\omega)\mathbf{u}(\omega) = \mathbf{b}(\omega), \quad (14)$$

where \mathbf{A} is a sparse, square, complex symmetric matrix of size N_{dof} with N_{nz} non-zero entries, \mathbf{u} is the vector of degrees of freedom, and \mathbf{b} is the complex-valued right-hand side vector.

The global matrix \mathbf{A} and the right-hand side \mathbf{b} are obtained by assembling the element contribution \mathbf{A}_e and \mathbf{b}_e . It is useful to make the dependence of the element matrices with frequency explicit by writing

$$\mathbf{A}_e = \mathbf{K}_e - k^2 \mathbf{M}_e + \gamma(\omega) \mathbf{M}_{\Gamma_e}, \quad (15)$$

where the following mass and stiffness element matrices have been introduced:

$$\mathbf{M}_e = \int_{\Omega_e} \Phi_e \Phi_e^T d\Omega_e, \quad \mathbf{K}_e = \int_{\Omega_e} \nabla \Phi_e \cdot \nabla \Phi_e^T d\Omega_e, \quad \mathbf{M}_{\Gamma_e} = \int_{\Gamma_e} \Phi_e \cdot \Phi_e^T d\Gamma_e, \quad (16)$$

of size $N_{\text{dof}}^e \times N_{\text{dof}}^e$. The column vector Φ_e contains the shape functions for the element e . The element matrices defined in (16) are independent of frequency. The right-hand-side element vector is also introduced

$$\mathbf{b}_e = \int_{\Omega_e} \Phi_e s(\omega) d\Omega_e + \int_{\Gamma_e} \Phi_e g(\omega) d\Gamma_e. \quad (17)$$

The integrals in Equations (16) and (17) are evaluated using numerical quadrature schemes [22, 25, 27]. A large number of quadrature points will invariably be needed because of the presence of high-order polynomials in the integrands. With curved elements and/or non-uniform coefficients, one has to increase the order of quadrature further to accommodate for the varying Jacobian and coefficients. Even with optimized quadrature rules, the cost of evaluating these integrals will grow very rapidly with increasing order p . As a consequence, and while solving the linear sparse system (14) remains the most costly task of the solution procedure, the calculation of element matrices is far from negligible for high-order FEM.

Once assembled, the sparse system (14) needs to be solved using a direct or an iterative approach. The development of robust iterative solvers for Helmholtz problems is still an area of active research. Various efficient short-term recurrence algorithms with dedicated preconditioners have been proposed for the algebraic systems arising from a linear FEM discretization (for a review, see [28]). However, the indefiniteness of the operator tends to hinder their performance at higher frequencies. Besides, the system matrices arising from a p -FEM discretization typically exhibits higher condition numbers than that of the linear FEM ones, which renders the problem even more challenging. Finally, iterative solutions are not suited for solutions involving a large number of right-hand sides, which are likely to be encountered in many applications. For all the reasons stated previously, iterative solutions will not be considered here, and instead, the Cholesky algorithm of the multifrontal sparse direct solver (MUMPS) [29] is employed to solve the linear system (14) at each frequency.

At this stage, it is useful to recall some results on the performance of high-order finite elements. For the hp version of the finite element method for the Helmholtz equation, Ihlenburg and Babuška [30] have obtained upper bounds for the global error in H^1 semi-norm

$$E < C_1 \left(\frac{kh}{2p} \right)^p + C_2 k \left(\frac{kh}{2p} \right)^{2p}. \quad (18)$$

The first term in this expression represents the interpolation error and has a dependency of order p with kh . This error can be controlled by using a sufficient number of elements per wavelength, given by $2\pi/(kh)$. The second term is associated with the dispersion error and the pollution effect. The dispersion error is the difference between the theoretical wavenumber k and the wavenumber \tilde{k} actually observed in the numerical model. Ainsworth [31] has obtained detailed results for the dispersion error of the hp -FEM. In particular for well-resolved models ($kh \ll 1$), the relative dispersion error can be approximated as

$$\frac{\tilde{k} - k}{k} = -\frac{1}{2} \left[\frac{p!}{(2p)!} \right]^2 \frac{(kh)^{2p}}{2p+1} + \mathcal{O}((kh)^{2p+2}). \quad (19)$$

This expression shows that the dispersion error decreases at a much faster rate than the interpolation error. However, the relative dispersion error (19) is a local effect because it represents the error incurred over one wavelength. In practice, a wave can travel over many wavelengths in the computational domain, leading to a significant build-up of the dispersion error. This is the so-called pollution effect, which is included in Equation (18) through the additional factor k in the second term. Therefore, in practice, the dispersion error still dominates the overall error for moderately and poorly resolved models.

2.4. An efficient multi-frequency algorithm

There are several issues associated with a multi-frequency calculation for a large model at high frequencies. The methodology outlined in this section addresses these issues.

Firstly, solving the large linear system associated with the numerical model is the main bottleneck due to the time and memory required for this task. This is addressed here by using high-order p -FEM that leads to a drastic reduction of computational cost.

Secondly, assembling a high-order FE model repeatedly for a large number of frequency is rather costly because of the numerical quadrature of high-order polynomials. This can be avoided by using hierarchical shape functions such as the Lobatto polynomials used here. The element matrices can be calculated only once for the highest order, and from these, the element matrices for lower-order models can be extracted directly.

Thirdly, when using a fixed-order method (e.g., standard linear or quadratic FEM), one would have to produce an individual mesh for each frequency to achieve the best performance (fixed accuracy with the minimum computational cost). This is not practical or efficient as the end user would have to create a large number of meshes.

In addition, in practical applications, the element size is rarely uniform in the computational domain, often due to the need to accurately resolve the geometry of the problem. In such situations, using the same polynomial order is not optimal as the resolution of the numerical model will vary significantly over the domain. Instead, one would like to use the lowest polynomial order possible in each element so as to achieve a uniform level of numerical error. This requires to adjust p independently in each element, which suggests the use of a local error indicator.

With standard Lagrange interpolation, it would be difficult to change the order independently as interpolation points are shared between neighboring elements. In contrast, with the hierarchical Lobatto interpolation, the high-order shape functions are modal, and their order can be easily defined independently in each element [22]. With a variable-order method such as p -FEM, it is possible to use a single mesh for the complete range of frequencies and adjust the polynomial order for each frequency to ensure the required level of accuracy is maintained. This approach is followed here by developing an *a priori* adaptive scheme, in which the solution strategy is outlined below. In what follows, the superscript i is used to denote the quantity computed at a given frequency f_i .

- (1) Consider an *a priori* error indicator attributing an order p_e^i to each element at each frequency $f_i \in \mathcal{F}$, so as to achieve a user-defined target error (that will be discussed in detail in Section 4)
- (2) Based on this indicator, the highest polynomial order required for each element across the frequency range is identified. In the case of non-frequency dependent sound velocity, this amounts to considering the maximum frequency of interest $p_e^{N_f}$.
- (3) Using the maximum conformity rule described previously, the highest polynomial orders used for each edge and face in the mesh can be determined.
- (4) For each element, the element matrices in (16), $\mathbf{M}_e^{N_f}$, $\mathbf{K}_e^{N_f}$, $\mathbf{M}_{\Gamma_e}^{N_f}$ are evaluated and stored up to the highest order. Because of the hierarchical structure of the Lobatto basis, these matrices also contain the smaller matrices that will be required for all the lower frequencies

$$\mathbf{M}_e^i \subset \mathbf{M}_e^{N_f}, \quad \mathbf{K}_e^i \subset \mathbf{K}_e^{N_f}, \quad \mathbf{M}_{\Gamma_e}^i \subset \mathbf{M}_{\Gamma_e}^{N_f}, \quad \forall i \in \mathcal{F} = \{f_1, f_2, \dots, f_{N_f}\}. \quad (20)$$

- (5) At each frequency i , the submatrices corresponding to the local order p_e^i

$$\mathbf{A}_e^i = \mathbf{K}_e^i - k_i^2 \mathbf{M}_e^i + \gamma(\omega_i) \mathbf{M}_{\Gamma_e}^i \quad (21)$$

are then recovered and assembled. The resulting global linear system is solved, and the solution is post-processed.

With this simple approach, all the element matrices are evaluated only once, which renders the discrete system evaluation negligible for problems involving a large number of frequencies.

3. EVALUATION OF PERFORMANCE AT FIXED ORDER

Before discussing the use of an adaptive scheme to select the optimal polynomial order, it is important to assess the performance of the p -FEM method with a fixed order. To properly design such an adaptive approach, it is necessary to have a clear understanding of the relation between the numerical model (element size and polynomial order), the computational cost (memory and time), and the accuracy. This is performed in this section by considering a simple test case, and by illustrating the impact of mesh size and order on the accuracy and computational cost. These results will inform the use of an adaptive method (described in Section 4) and will also provide guidelines for the design of a mesh suitable for p -FEM.

3.1. Description of the test case

As a three-dimensional benchmark problem, the propagation of a time-harmonic plane wave is simulated using the p -FEM approach described previously. The exact solution is written

$$u_{\text{ex}} = \exp(-i\mathbf{k} \cdot \mathbf{x}) , \quad \text{with } \mathbf{k} = k (\cos \theta \cos \alpha, \sin \theta \cos \alpha, \sin \alpha) , \quad (22)$$

where \mathbf{k} is the wave vector and α and θ are the azimuthal and polar angle of the wave (Figure 1). For simplicity, the computational domain is a cube of size a represented by a mesh of tetrahedral elements with uniform size h .

The discussion will concentrate on tetrahedral elements as they are the most common elements used for this type of applications (results for other types of elements can be found in [32]). The inhomogeneous Robin condition (2) is used on all boundaries to generate the plane wave inside the computational domain by setting $\gamma = ik$ and $g = \partial u_{\text{ex}} / \partial n + \gamma u_{\text{ex}} = -i(\mathbf{k} \cdot \mathbf{n} - k)u_{\text{ex}}$. Note that the source term g is not interpolated on the mesh but is calculated exactly at the Gauss points for the evaluation of the right-hand side \mathbf{b}_e in Equation (17).

Finite element models for the Helmholtz equation suffer from two main sources of error: the interpolation error and the dispersion error [33]. Both errors are highly anisotropic, as they vary significantly with the orientation of the incident wave relative to each element. In order to minimize this effect and characterize the performance of the p -FEM in a way that is representative of real applications, a fully unstructured mesh and a large number of plane wave directions are considered. A total of 30 plane waves with randomly distributed directions in the range $\theta = [0, 2\pi]$ and $\alpha = [0, \pi]$ are used. This is to account for the fact that in real applications, the acoustic field is composed of a combination of waves with different directions, and the numerical error will be controlled by the wave direction inducing the largest error.

For each polynomial order from $p = 1$ (conventional linear finite elements) to $p = 10$, the mesh size is varied from $h = a$ to $h = a/50$. Two tetrahedral meshes, corresponding to $h = a/4$ and $h = a/35$ are shown in Figure 2.




Figure 1. Sketch of the fixed-order test case.




Figure 2. Examples of meshes used for the fixed-order test case with $h = a/4$ (left) and $h = a/35$ (right).

Only ten meshes were considered per order, leading to a total of 100 tetrahedral (h, p) combinations. For each (h, p) combination, the solution is calculated and the accuracy is measured by evaluating the relative error in the L^2 or H^1 norm defined by

$$E_{L^2} = \frac{\|u_{\text{ex}} - u\|_{L^2}}{\|u_{\text{ex}}\|_{L^2}}, \quad \text{with } \|f\|_{L^2} = \left(\int_{\Omega} |f|^2 \, d\Omega \right)^{1/2}, \quad (23)$$

and

$$E_{H^1} = \frac{\|u_{\text{ex}} - u\|_{H^1}}{\|u_{\text{ex}}\|_{H^1}}, \quad \text{with } \|f\|_{H^1} = \left(\int_{\Omega} |f|^2 + |\nabla f|^2 / k^2 \, d\Omega \right)^{1/2}. \quad (24)$$

This is repeated for the range of plane wave directions described previously, and the largest error over the wave directions is retained. In addition, the cost of solving the linear system is also measured. In this work, we use the multi-frontal solver MUMPS 4.10 [29] for large, sparse linear systems to solve Equation (14). The time and memory required were measured on a PC with 16 threads, 48 Gb of memory and 2.80 GHz CPU speed (all calculations were performed ‘in-core’, and the out-of-core method was not used). Note that the times recorded include both the factorization of the global matrix and the back-substitutions for the 30 wave directions.

3.2. Performance analysis

This section investigates the performance of p -FEM in the mid-frequency and high-frequency regimes, that is, when the cost of the conventional FEM increases significantly due to the pollution effect (Equation (18)). A high Helmholtz number $ka = 50$ provides a representative example of challenging cases for three-dimensional FE models. Other frequencies have also been considered ($ka = 25$ and 75), and the same conclusions can be drawn.

First, the properties of the sparse global matrix \mathbf{A} are examined by plotting the number of degrees of freedom as a function of the relative H^1 error in Figure 3(a) (results obtained with the relative L^2 error lead to similar conclusions and are not reported here for conciseness). Note that the bubble shape functions in each element have been factored out and are not included in the results in this graph. As expected, increasing the order p allows to significantly reduce the number of degrees of freedom required to achieve a given accuracy. For instance, for a relative error of 1 %, up to 1 million unknowns are needed with an order $p = 3$, but this is reduced by an order of magnitude when using $p = 10$.

However, the global matrices originating from high-order models are more densely populated. Figure 3(b) shows the number of non-zero entries against accuracy, and it is clear that increasing the polynomial order has a limited impact on the non-zero entries necessary to achieve a given accuracy. This is because, for a given accuracy, the number of degrees of freedom is significantly reduced by increasing p , but at the same time, the number of interactions between degrees of freedom increases rapidly with p . As a consequence, the total numbers of terms in the matrices resulting from the

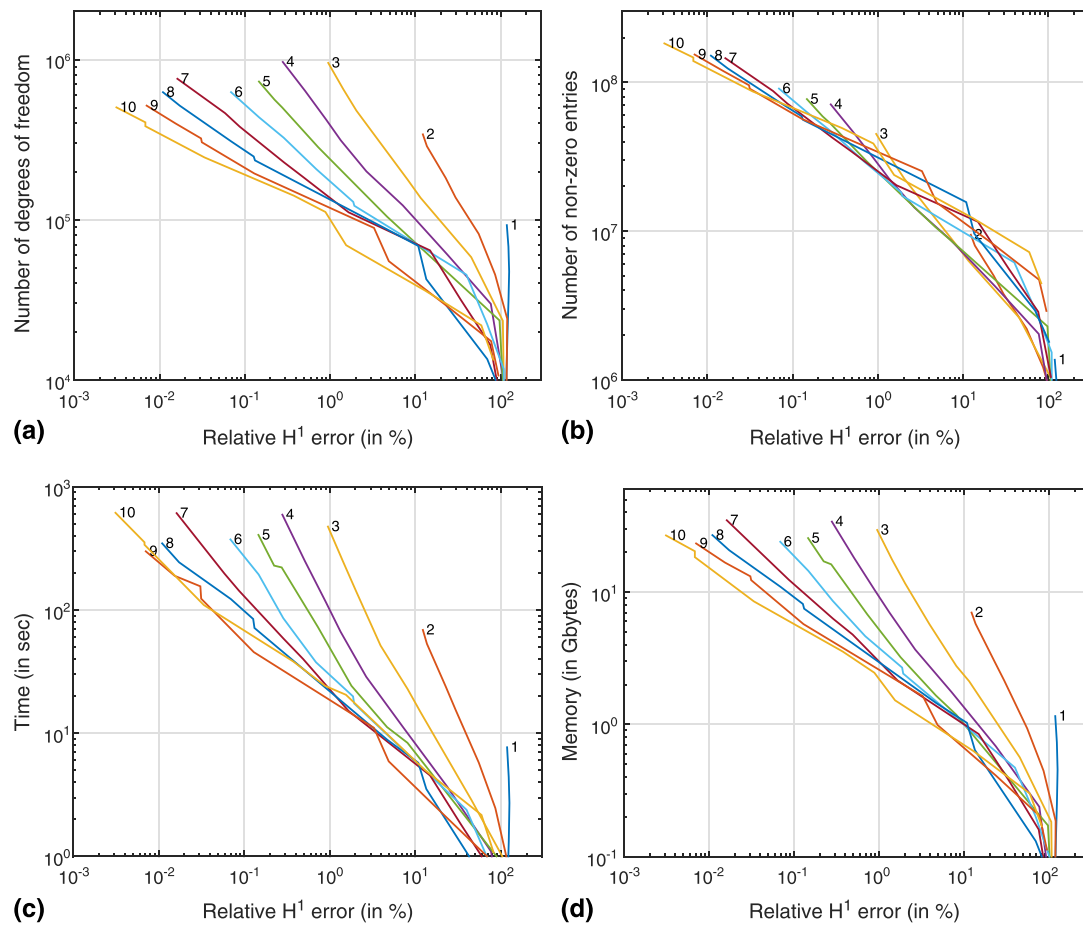


Figure 3. Computational cost to reach a given accuracy for the 100 tetrahedral (h, p) combinations at $ka = 50$. (a) Number of degrees of freedom, (b) number of non-zero entries, (c) factorization time and (d) factorization memory.

low-order and high-order discretizations are of the same order of magnitude. For instance, for a 1 % accuracy, the average number of entries per row in the matrix varies from 14.9 with $p = 1$ to 364.2 with $p = 10$. The total number of non-zero entries needed to reach a 1 % accuracy varies only slightly with the polynomial order and remains close to 30 million.

The argument that the improvement in accuracy of high-order p -FEM are significantly, or completely, balanced by the rapid increase in cost is often put forward in the literature. But the results in Figure 3(c) and (d) for the time and memory used by p -FEM suggest that this argument is not applicable for the present solver. At an engineering accuracy of 1%, increasing the polynomial order does yield significant benefits for both measures of cost. With $p = 10$, the solution procedure is two orders of magnitude faster than the conventional FEM solution ($p = 2$). A similar drastic improvement is observed for the memory requirements. A quadratic solution requires around 200 Gb to converge to 1% error but will require only 2 Gb when increasing the order to $p = 10$. This represents the difference between an intensive simulation mobilizing a specialized computational platform and a routine calculation on a basic workstation.

The results in Figure 3 also illustrates that the link between the number of non-zero entries and the cost of solving the linear system is not straightforward, and an increase in non-zero entries does not translate into a comparable increase in computational cost. At constant number of entries, it appears that the structure of high-order matrices is more favorable for factorization, by requiring less fill-in.

A key observation from this test case is that the optimal (h, p) combination is almost always obtained with the maximum order $p = 10$ (one exception being for low levels of accuracy, typically

$E_{H^1} > 20\%$). While the benefit observed by increasing the order tends to diminish for large orders, we did not observe trends suggesting that a further increase in p would lead to a deterioration in performance.

This is in contrast with the conclusion reported in Vos *et al.* [10] where runtimes for a large number of (h, p) combinations were measured for a two-dimensional Helmholtz problem. For an engineering accuracy of 10% at $ka = 2\pi$, the minimal computational time was obtained using the sixth-order expansion $p = 6$. Several reasons can explain this difference with the present work. Firstly, the previous study focused on the performance of the p -FEM in the lower frequency range, with $ka = 2\pi$. In this regime, the pollution effect is negligible, and the interpolation error dominates, and in Equation (18), it is clear that increasing the order in p -FEM has a much stronger impact on dispersion error compared with interpolation error. Secondly, only a two-dimensional Helmholtz problem was examined in [10], and it is known that resorting to higher polynomial orders is even more beneficial in three dimensions [11]. Thirdly, it could be argued that the gains observed depend to some extent on the value of the targeted error E_{H^1} . In particular, a higher target error would tend to favor lower-order models. However, even if one considers a much larger error of, say 20%, the high-order FEM solutions are still computationally cheaper, although to a lesser extent.

It is also useful to examine the link between element size h , polynomial order p , and accuracy, which is shown in Figure 4. As expected, the requirement in terms of element size required to achieve a prescribed accuracy is greatly reduced when using higher-order elements. As an example, for a fixed accuracy of 1%, there is a factor 7 between the element size required with cubic elements and with elements of order 10. For real applications, there is a practical limit to this trend: one cannot use too large elements because the geometry of the problems has to be described with sufficient accuracy. This is an important aspect of high-order FEM that will be discussed in more detail in Section 5 for a realistic application.

We can now summarize the main conclusions that will guide the strategy for the p -adaptive scheme presented in the next section.

To achieve a given accuracy, resorting to high-order elements is consistently beneficial in terms of performance, as shown by the results in Figure 3. Therefore, the guideline for mesh generation is to use elements as large as possible while still capturing the geometry accurately (this guideline will be discussed in more detail in Section 5).

For a given frequency and mesh, the objective is to identify a trade-off between accuracy and efficiency. In this work, one specifies a target level of accuracy, and the smallest polynomial order required to achieve this accuracy is then identified. How this is achieved automatically using a simple error indicator is the topic of the next section.




Figure 4. Mesh resolution required to reach a given accuracy for the 100 tetrahedral (h, p) combinations at $ka = 50$.

4. DESIGN OF AN *A PRIORI* ERROR INDICATOR

Following the strategy outlined in the previous section, a simple error indicator used to adjust the polynomial order in each element is now described and verified.

4.1. Description of the adaptive scheme

In each element, the relevant parameters are the sound speed c , the frequency ω , and the edge lengths h_i for this element. Based on this information, the aim is to select the polynomial order p to ensure that the relative numerical error remains below a prescribed target E_T based on E_{L^2} or E_{H^1} , as defined by (23) and (24).

To achieve this, one needs an error estimator, or indicator, to identify the smallest order sufficient to meet the error target. Error estimators approximate a measure of the actual error in a given norm (like in Equation (18)). On the other hand, error indicators are based on heuristic considerations [34]. It is shown here that a simple error indicator can be used effectively to provide sufficient control on the actual numerical error to be used in practice. The reasons for using this indicator instead of other alternatives are also discussed.

The error incurred in the element is estimated *a priori* using a single, one-dimensional element with the same wavenumber k and a length h calculated from the edge lengths h_i of the actual element. On this domain $[0, h]$, the variational formulation (4) is solved with the Robin boundary condition (2). At $x = 0$, the choice of parameters $\gamma = ik$ and $g = 1$ generates a single traveling wave e^{-ikx} . The parameters $\gamma = ik$ and $g = 0$ are used at $x = h$ so that this traveling wave leaves the computational domain without spurious reflection. The numerical error ε over this one-dimensional element can be calculated in the L^2 or H^1 norm and is denoted by ε_{L^2} or ε_{H^1} . This error is a function of only the polynomial order p and kh (the ratio between the element size h and the acoustic wavelength), and it is plotted in Figure 5. For a given frequency, sound speed, and element size, it is then straightforward to identify the smallest order p required to maintain the error ε_{L^2} or ε_{H^1} below the target error E_T .

In fact, it is possible to tabulate the order $p(kh, E_T)$ in terms of the mesh resolution kh and the target error E_T such that in practice, one does not need to solve the 1D model. This is illustrated in Table I for three levels of error, namely, $\varepsilon_{L^2} = 15\%$, 5% , and 0.5% . This table can be used to directly determine the element order based on the knowledge of the local values of k and h .

Different definitions of E_T have been compared (based on ε_{L^2} or ε_{H^1}), as well as different definitions of the length h for the 1D model (minimum, maximum, or average values of the edge lengths). All these tests are not reported here for the sake of brevity but can be found in [35]. In what follows, the length h of the 1D element is defined as the average edge length of the actual element as it was




Figure 5. Errors $\varepsilon_{L^2}(kh, p)$ (a) and $\varepsilon_{H^1}(kh, p)$ (b) measured on the single 1D element with the polynomial order p varying from 1 to 10.

Table I. Tabulated values of kh to achieve a given accuracy at a given order p .

Order p	1	2	3	4	5	6	7	8	9	10
$\varepsilon_{L^2} = 15\%$	1.5	2.9	4.6	6.4	8.1	10.1	11.8	13.7	15.5	17.4
$\varepsilon_{L^2} = 5\%$	0.8	2.0	3.4	5.0	6.6	8.4	10.1	11.9	13.7	15.4
$\varepsilon_{L^2} = 0.5\%$	0.2	0.9	1.9	3.1	4.5	5.9	7.4	8.9	10.6	12.2

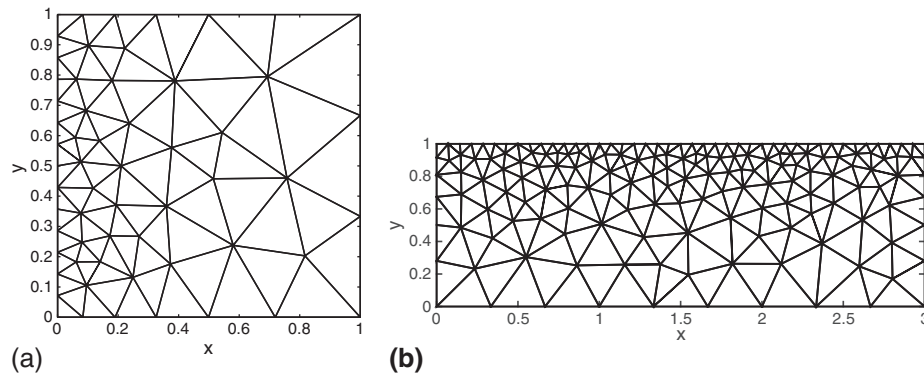


Figure 6. (a) Non-uniform mesh used for the 2D verification of the adaptive order scheme. (b) Mesh used to illustrate the effect of the pollution effect.

found to be more robust. Besides, the target error E_T is specified in terms of the error based on the $\varepsilon_{L^2}(kh, p)$ measure calculated on the 1D element.

There are several reasons for using this error indicator. First of all, error estimators such as (18) rely on constants that are strongly problem-dependent. In addition, when the dispersion error dominates, Equation (18) tends to overestimate the rate of convergence for large order (more precisely when kh is not very small).

The coefficients C_1 and C_2 in Equation (18) are independent of the element size h , the frequency, and are only weak functions of the polynomial order p . However, they are still problem dependent. Using simple tests, these coefficients have been found to vary significantly between different test cases.

4.2. Verification for 2D models

To show that this simple error indicator is able to control the numerical error, a 2D model of a single plane wave propagating in a duct is first considered. The square computational domain, discretized with a fully unstructured mesh, is shown in Figure 6(a). Parameters are made non-dimensional by setting the domain size, the sound speed, and density to unity. A single plane wave is introduced on the boundary $x = 0$, while the unit impedance is imposed at $x = 1$, so that the plane wave leaves the domain without reflection. A hard-wall condition is applied on the other boundaries.

In Figure 7, the actual numerical error E_{L^2} calculated over the computational domain is plotted for four different frequencies against the target error E_T used to adjust p in each element. It can be seen that in almost all cases, the actual error is lower but close to the target error, thus maximizing the efficiency of the numerical model.

The main exception is found at high frequencies (Figure 7(c) and (d)) for small values of the target error E_T . In this situation, all the elements are using the highest order available, in this case $p_{\max} = 10$, but this is not sufficient to achieve the requested level of accuracy. This can be easily addressed by implementing higher-order shape functions or by refining the mesh.

But Figure 7 illustrates clearly that the mesh resolution has to be selected so that the target accuracy can be achieved by the adaptive method over the whole range of frequencies of interest. Table I shows, for instance, that for a target error of 0.5%, the resolution should be chosen such that $kh < 12.2$ across the whole mesh, for a highest available order $p_{\max} = 10$.




Figure 7. Numerical error as a function of the target error E_T for $\omega = 1$ (a), 10 (b), 20 (c), and 50 (d). Actual global relative error E_{L2} in the numerical model (solid blue line). Minimum and maximum values of the error indicator ε_{L2} (red dashed lines). The dotted black line shows when the actual error matches the target error.

Note also that for large E_T and low frequencies (Figure 7(a)), all the elements use linear shape functions, and yet the actual error is significantly smaller than the target. This is not an issue because the cost of solving the problem at such low frequencies is insignificant compared with high frequencies.

When the polynomial order has been chosen in each element, one can use the error indicator ε_{L2} to measure the local numerical error incurred within each element. It is then interesting to compare the minimum and maximum values of the local error given by ε_{L2} against the global error E_{L2} actually obtained. This is shown in Figure 7 where the global error (blue line) is almost always within the bounds of the local error (red line), the only exception being for very high values of E_T . This suggests that a simple local error indicator such as ε_{L2} can be used to effectively control the global numerical error (a caveat is when the pollution effect is very significant; Section 4.4).

4.3. Verification for 3D models

We illustrate that this simple error indicator is applicable to three-dimensional problems as well. The computational domain is a unit cube with the same boundary condition as in the 2D test case mentioned previously.

Different meshes have been considered and are shown in Figure 8. They include a uniform coarse mesh and three meshes with different levels of local refinement to illustrate more realistic situations when the boundary of the domain is often refined to capture the geometry accurately.




Figure 8. Unit cube with unstructured meshes with $h_{\max} = 0.2$. Uniform element size $h_{\min} = h_{\max}$ (a), non-uniform element sizes: $h_{\min} = 0.05$ (b), $h_{\min} = 0.04$ (c), and $h_{\min} = 0.02$ (d).

To illustrate the use of the p -FEM method in practice, each of these meshes is used to solve this problem over a wide frequency range from $\omega = 2\pi/5$ to $\omega = 26\pi$. At the highest frequency, this corresponds to 13 acoustic wavelengths in the computational domain. As the frequency is varied, the method relies solely on the use of the adaptive polynomial order to improve the resolution of the numerical model at high frequency because the same mesh is used for all frequencies. Three different error levels are prescribed for the selection of the polynomial orders, with $E_T = 15\%$, 5% , and 0.5% .

The global error E_{L^2} obtained in all these cases is shown in Figure 9, while Figure 10 shows the evolution of the minimum, mean, and maximum order as the frequency increases. It can be seen that the adaptive method is indeed able to maintain the numerical error at a relatively constant level over most of the frequency range. As expected, this is achieved by progressively increasing the polynomial order. For strongly non-uniform meshes (Figure 10(b) corresponding to Figure 9(b)), the range of orders used in the model is much larger because p remains low in the small elements. The actual errors obtained in 3D are somewhat lower than the target accuracy obtained from the 1D model. For instance, a target error of $E_T = 5\%$ yields an average L^2 error close to 1% for all the meshes in Figure 8.

For low frequencies, some oscillations are visible. This is because the transition from order 1 to 2 has a much stronger impact on the accuracy compared with, say, the transition from order 7 to 8. For the uniform mesh, this transition happens simultaneously for almost all the elements, as seen in Figure 10(a), resulting in the sharp changes in accuracy visible in Figure 9(a).




Figure 9. Global numerical error for the meshes shown in Figure 8 with target accuracy E_T 15% (dashed line), 5% (solid line), and 0.5% (dotted line).

At high frequencies, it can be seen that for each mesh, the error increases above a certain frequency, as already observed above for the 2D results. This is because the adaptive order allocates the highest available order $p_{\max} = 10$ to almost all the elements, but this is not sufficient to maintain the same target level of global error.

4.4. Pollution effect

It is important to note that this error indicator is used in each element independently and that therefore, it can only control the numerical error incurred within this element. As a consequence, it is not designed to account for the pollution effect, which corresponds to the build-up of phase error as a wave propagates over several wavelengths.

This can be illustrated by using the same test case as in Section 4.2 by keeping all the parameters constant except for the length of the domain, which varies from 1 to 20. A non-uniform mesh is used for this purpose with an element size varying from $h = 0.3$ at $y = 0$ to $h = 0.07$ at $y = 1$ (see Figure 6(b) for a domain length of 3). Different frequencies are tested such that in the worst configuration, for a domain length of 20 at $\omega = 40$, more than 127 wavelengths are found in the domain. The target error is fixed to $E_T = 1\%$.

As shown in Figure 11, the global numerical error steadily increases with domain length, and this is directly attributed to the pollution effect. Also shown in Figure 11 are the minimum and maximum values of the error indicator over all the elements. As expected, these remain almost constant with ω because the same target accuracy has been specified for all frequencies. For problems with a very large number of wavelengths in the domain, it is advised to reduce the target error E_T so as to maintain a sufficient level of accuracy for the global solution.




Figure 10. Minimum (green), average (blue), and maximum (red) orders for the meshes (a) and (b) in Figure 8 as a function of frequency.

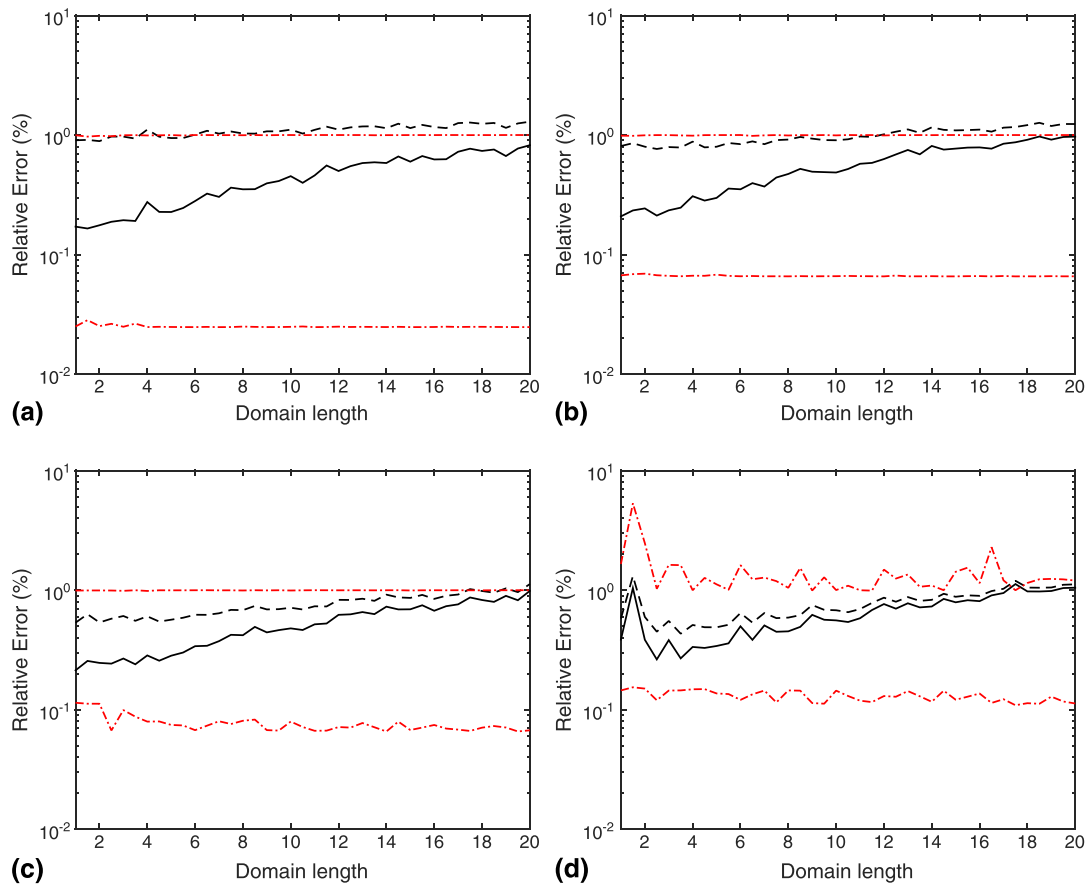


Figure 11. Actual error as a function of the length of the computational domain for (a) $\omega = 5$, (b) $\omega = 10$, (c) $\omega = 20$, and (d) $\omega = 40$. E_{L^2} error (solid line), E_{H^1} error (dashed line), minimum and maximum values of the error indicator ε_{L^2} (red line) over all the elements.

5. APPLICATION TO A REALISTIC CASE

The application of the p -FEM approach for a realistic 3D problem is now discussed. The first objective is to illustrate the benefits of the high-order FEM solver on a real-life engineering test case, by comparing its performance with a generic commercial FEM solution. The second objective is to highlight best practice that should be followed when preparing models for adaptive high-order solvers.




Figure 12. Schematic of the car cabin together with the associated boundary conditions.

All computations reported in this section were performed using the FEMAO solver, which is an industrial implementation of the high-order FEM method with the adaptive order approach described in Section 4. Developed in Fortran 2003 and using the MUMPS library [29] to solve sparse linear systems, the FEMAO solver is part of the vibro-acoustic simulation package Virtual.Lab Acoustic, commercialized by the company Siemens Industry Software [36]. Simulations were performed on a platform with 126 Gb of RAM and two Intel Xeon processors operating at 2.53 GHz with a total of 16 threads.

5.1. Description of the test case

Predicting the sound field in the cabin of a car is a typical example of numerical modeling regularly performed to support the acoustic design of new vehicles. Acoustic simulation tools are used in the automotive industry to reduce the booming noise corresponding to high sound pressure levels near the resonance frequencies of the cabin. This is also an example of applications where the use of numerical modeling is still limited by the cost of standard FE methods.

A realistic geometry is used for the cabin of the car, as illustrated in Figure 12. The overall dimensions of the model are $L_x = 2.25$ m, $L_y = 1.62$ m, and $L_z = 1.23$ m, and the total volume is 3.42 m³. The sound speed and fluid density are considered homogeneous in the cavity with $c = 340$ m/s and $\rho = 1.225$ kg/m³. Several surfaces of the cabin are acoustically treated: the floor, the ceiling, and the so-called firewall (the lower front of the cabin). The impedance condition representing these boundaries is obtained by setting $\gamma = i\omega\rho A_z$ and $g = 0$ in the Robin condition (2). The acoustic admittance A_z of the surface is assumed to be a cubic function of frequency [37]. The resulting linear system in (15) is therefore not of second order. In this situation, the application of the modal analysis is difficult, and resorting to a direct solution procedure is justified. Two separate external excitations are considered. First, a point mass source of amplitude $A_s = 10^{-3}$ is positioned at $\mathbf{x}_s = [0.4, 0., 0.2]$ (Figure 12). This source is enforced by applying $s = A_s\delta(\mathbf{x}_s)$ in Equation (6). Secondly, a normal velocity $V = 10^{-3}$ m/s is imposed on the firewall to simulate an acoustic excitation due to the engine vibrations. This is achieved using the Robin condition (2) with $\gamma = 0$ and $g = -i\omega\rho V$.

The aim with this model is to obtain accurate predictions of the pressure field over the frequency range from 20 Hz to 2 kHz. This corresponds to a Helmholtz number in the range $ka = [0.6, 60]$ based on a model characteristic length of $a = L_y = 1.62$ m.

5.2. Mesh creation process

We begin by discussing the issues influencing the design of a finite element mesh suitable for high-order FEM, and the meshes used for the present test case are then described.

As already outlined in Section 3, the computational performance of p -FEM is maximized when large elements with high-polynomial orders are used. An added benefit of this is that it makes the process of generating and manipulating these meshes much easier and faster, compared with the

fine meshes required for conventional finite element models. There are however a number of other considerations that should be addressed in any practical problem.

Firstly, while the polynomial order is automatically selected, there is necessarily a limit to the order available, in this case $p_{\max} = 10$. The mesh should be sufficiently fine so that at the highest frequency of interest, the maximum order will be sufficient to maintain a sufficient level of accuracy. Based on the definition of the error indicator given in Section 4, the following recommendation for the element size can be inferred

$$h < C_{ET} \frac{c}{f_{\max}} \quad (25)$$

where the constant C_{ET} is adjusted according to the required accuracy and the maximum available order. From the tabulated values in Table I for $p_{\max} = 10$, one finds that $C_{5\%} = 15.4/(2\pi) = 2.44$ and $C_{0.5\%} = 12.2/(2\pi) = 1.94$.

Secondly, the geometry of the problem should also be accurately represented by the finite element mesh. Capturing the relevant geometrical features in the model will usually require to use smaller elements near the boundaries of the domain. Improved geometry representation in FE model is an active research area, including high-order curvilinear meshes [38] or NURBS-enhanced method [39]. This goes beyond the scope of the present paper, and in this paper, the element geometry is described with classical linear interpolation.

The level of refinement of the geometry is determined by the ratio between the characteristic length l of a given geometric detail and the smallest acoustic wavelength λ_{\min} . If $l \ll \lambda_{\min}$, then this geometrical feature can be considered acoustically compact, and it does not need to be accurately described. However, if $l \approx \lambda_{\min}$, then the mesh should be refined locally such that $h \ll l$.

With conventional iso-parametric FEM, the geometry and the solution are approximated with the same interpolation. The mesh resolution selected for the solution (based on the number of points per wavelength) is implicitly providing a selection of the geometrical length scales (in other words, any feature not captured by the mesh is unlikely to be significant for the wavelength of interest). With p -FEM, the discretization of the geometry and of the solution are decoupled, and one should ensure that all the relevant geometrical features of the problem are appropriately resolved.

In order to demonstrate the robustness of the proposed adaptive approach, several meshes were generated. Two surface meshes of the car cabin are used with a uniform element size of $h = 17$ mm or $h = 22$ mm, as shown in Figure 13. From these surfaces meshes, the volume of the domain is discretized with linear tetrahedral elements.

For the reference solution based on the conventional FEM, the fine surface mesh is used ($h = 17$ mm), and the tetrahedral elements inside the domain also have a size $h = 17$ mm. Solutions obtained with this mesh will be denoted 'FEM' in what follows.

To illustrate the guidelines discussed previously, the meshes used for the FEMAO solver use larger elements inside the domain so that the benefit of using higher-order approximation can be demonstrated. Three different meshes are considered for the high-order adaptive method:

- A mesh based on the surface mesh with $h = 17$ mm with a ratio between smallest and largest element size of 8, denoted 'FEMAO I'.




Figure 13. Envelope surface meshes used to generate the volume meshes (a) coarse grid, $h = 22$ mm, and (b) refined grid $h = 17$ mm.

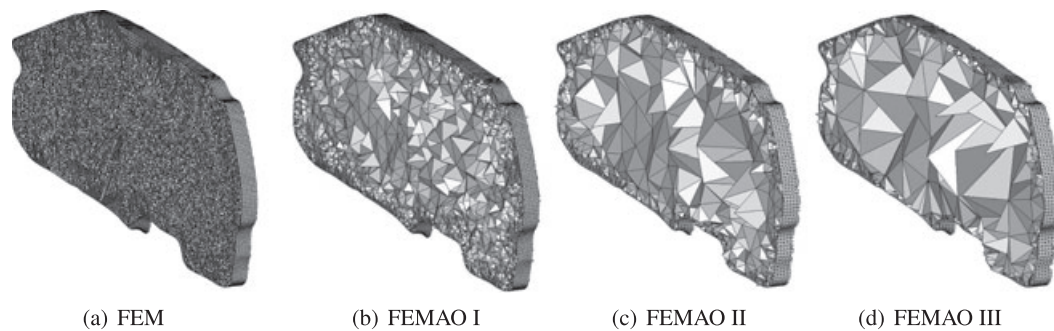


Figure 14. Cross section of the four meshes used for the car cavity test case.

Table II. Characteristics of the car cavity meshes used for the computations.

Mesh	No. of nodes	No. of elements	Max frequency (Hz)
FEM	1 173 716	6 902 675	2000*
FEMAO I	177 803	869 540	5215 [†]
FEMAO II	62 865	261 626	2987 [†]
FEMAO III	53 601	200 938	2020 [†]

*Computed using the classical FEM rule of thumb (27).

[†]Computed using the *a priori* error indicator rule (26) with $E_T = 5\%$.

- A mesh based on an envelope mesh of size $h = 22$ mm with a ratio between smallest and largest element size of 12, denoted ‘FEMAO II’.
- A mesh based on an envelope mesh of size $h = 22$ mm with a ratio between smallest and largest element size of 20, denoted ‘FEMAO III’.

Longitudinal cross sections of these four different meshes are shown in Figure 14. Their properties are outlined in Table II. The high-order coarse meshes contain far less nodes and elements than their FEM counterpart. The maximum frequency of a given high-order mesh is easily obtained from Equation (25), it reads

$$f_{\max} = C_{E_T} \frac{c}{h_{\max}} \quad (26)$$

where h_{\max} designates the largest element average edge length in the mesh. For the FEM mesh, the maximum frequency is obtained from the usual rule of thumb:

$$f_{\max} = \frac{1}{10} \frac{c}{h_{\max}} \quad (27)$$

which guarantees a resolution rate of 10 degrees of freedom per wavelength. According to these rules, all the meshes presented here are valid up to 2 kHz.

5.3. Comparison of results

The numerical predictions at several distinct frequencies are first compared. For each of these frequencies, the reference linear FEM solutions is computed using the mesh described previously, and the high-order adaptive method is then used with the other three coarse finite element meshes (Table II) with a fixed target accuracy of $E_T = 5\%$.

The numerical predictions are compared on a horizontal control plane of size 1×1.8 m, located in the center of the car cavity (Figure 12).

The real part of the acoustic pressure obtained with the monopole source is shown in Figure 15 for 400, 1400 and 2000 Hz. Overall, the pressure distribution is in good agreement for all models across the whole frequency range. For 2 kHz, the differences between the standard FEM model and




Figure 15. Real part of pressure on the control plane for the monopole source. Models (from left to right): FEM, FEMAO I, FEMAO II, and FEMAO III. Top: 400 Hz (16 contours levels between $\pm 8.86 \cdot 10^{-4}$ Pa); Middle: 1400 Hz (16 contours levels between $\pm 1.38 \cdot 10^{-3}$ Pa); Bottom: 2000 Hz (16 contours levels between $\pm 8.85 \cdot 10^{-4}$ Pa).

the high-order solutions appear to be more visible, and this is because the mesh resolution for the standard FEM is close to its limit.

The results obtained with the imposed panel velocity instead of the monopole source yield similar conclusions and are not reported here for brevity.

A second comparison is performed in terms of the frequency response function (FRF) by covering the frequency range from 20 to 2000 Hz with a fine resolution of 5 Hz (corresponding to 397

calculations). Note that for such a fine frequency increment, fast frequency reduction methods involving reduced-order models could also be considered to accelerate the frequency sweep [37, 40], but this is left for further study.

The FRFs in the center of the control plane predicted by the four numerical models are shown in Figure 16 for the velocity boundary condition and for the monopole source (with a dB reference of $2 \cdot 10^{-5}$ Pa).

Below 1200 Hz, the four models yield very similar predictions (except near resonance frequencies). Above 1200 Hz, there is very little difference between the three high-order models, but the standard FEM solutions begin to differ significantly and the resonance frequencies tend to be overestimated, which is a typical symptom of the pollution effect.

To confirm that it is indeed the standard FE model that is less accurate at high frequencies, calculations have been performed with the high-order scheme with a much more stringent target of




Figure 16. Frequency response at the center of the control plane for the velocity boundary condition (top) and the monopole source (bottom). Straight line: FEMA0 I; dashed line: FEMA0 II; dash-dotted line: FEMA0 III, blue line: FEM.




Figure 17. Frequency response at the center of the control plane for the monopole source (5 Hz increment). Black line: FEMA0 II with standard accuracy $E_T = 5\%$; magenta line: FEMA0 II with refined accuracy $E_T = 0.5\%$.

$E_T = 0.5\%$ for the numerical accuracy using the mesh FEMA0 II. When compared with the results obtained with $E_T = 5\%$ (Figure 17), it is clear that there is very little difference and that the high-order solutions have already converged with $E_T = 5\%$ over the whole range of frequencies of interest.

Another observation in Figure 16 is that the adaptive scheme is able to adjust the polynomial order to obtain reliable results over a wide range of frequencies despite the fact that the three meshes used with the FEMA0 solver are very different (with a change in characteristic element size of the order of 10 to 20).

Although the proposed error indicator is only local and does not account for the pollution effect, these results indicate that for this test case, the adaptive scheme is able to maintain a sufficient level of accuracy even at high frequencies when the waves are propagating over several wavelengths within the computational domain.

5.4. Computational performance

The efficiency of the different numerical models is now investigated.

The computational cost of the high-order adaptive approach is inherently varying with frequency because the error indicator will adjust the polynomial order with frequency. This mechanism can be clearly illustrated by plotting the different orders p used for each frequency. This is shown in Figure 18 for 20 frequencies between 100 and 2000 Hz by plotting the relative predominance of each order measured in terms of volume percentage (that is the volume occupied by all elements of order p compared with the total volume).

This is a convenient metric for highly inhomogeneous meshes where a small number of elements can represent a large portion of the computational domain.

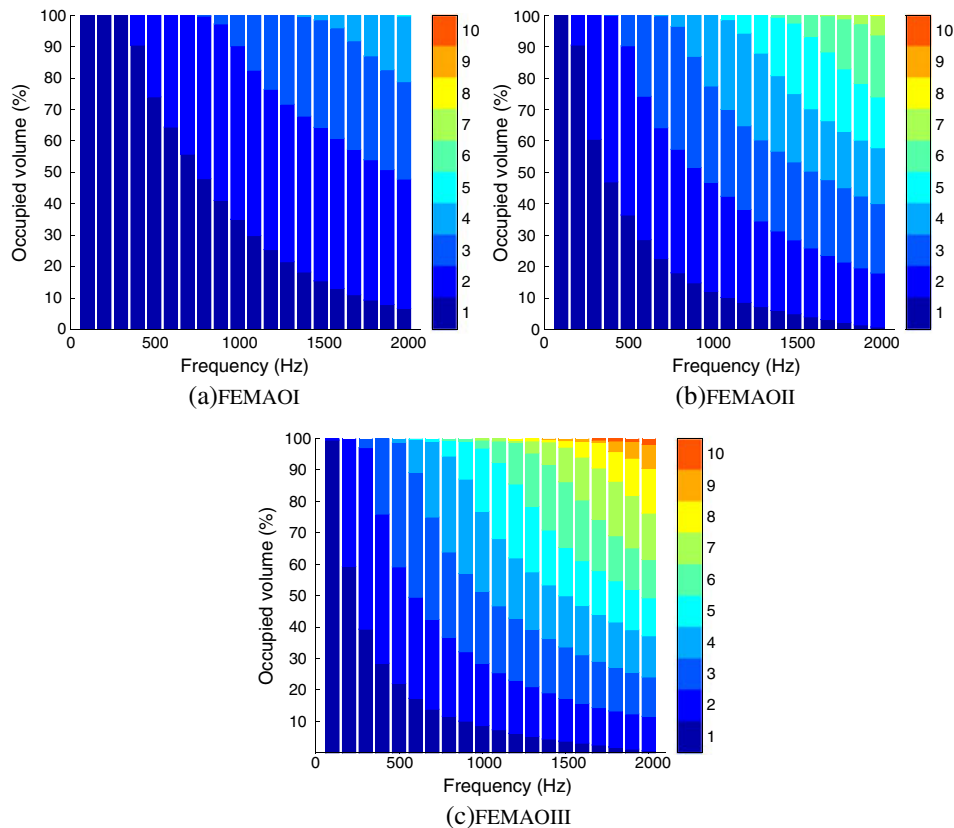


Figure 18. Element order repartition, in percentage of the total volume, as a function of frequency.

Predictably, at 100 Hz, all the FEMAO solutions rely on a linear approximation as this is sufficient to achieve the target accuracy. At higher frequencies, as the proportion of higher-orders progressively increases, the volume occupied by the linear elements decreases, until they represent less than 10% of the volume at 2000 Hz (they might still be dominant in terms of number of elements). These remaining linear elements are found in the refined areas of the mesh close to the boundary of the computational domain. In terms of the highest polynomial order, to cover the full frequency range, the FEMAO I model uses at most elements of order $p = 5$, while the FEMAO II and III models require elements of order $p = 8$ and 10, respectively. It should also be noted that on the mesh FEMAO III, between 1700 and 2000 Hz, the complete range of approximations is used, with polynomial orders ranging from $p = 1$ to $p = 10$. This highlights the fact that the error indicator adjusts the order at each frequency to the local properties of the mesh.

In terms of performance, for the 20 frequencies shown in Figure 18, the high-order FEM adaptive solutions were considerably faster than the linear FEM model. The full computation time, including the matrix assembly and the post-processing of the results on the control plane (shown in Figure 15), took 6641, 1399, 658, and 553 s for the FEM, FEMAO I, FEMAO II, and FEMAO III models, respectively. This corresponds to speed-up factors of respectively 4.7, 10.1, and 12.0 for the three high-order models.

The computational cost of computing the frequency response functions presented in Figure 16 is shown in Figure 19 in terms of the runtime and memory required for each frequency. The computational time per frequency includes the matrix assembly, the static condensation of all internal degrees of freedom, the factorization and the solving of the two right-hand sides through forward-backward substitution. The computational cost of the FEM reference is also shown for comparison. Note that the calculation of the element matrices is not included as the mass, stiffness, and damping matrices are computed prior to the frequency loop (this applies as well to the FEM reference solution).

When compared with the reference FEM solution, the high-order adaptive models are faster and require less memory across the whole frequency range. This is also true above 1200 Hz where Figures 16 and 17 show that the high-order models are more accurate.

Two different regimes can be observed in Figure 19. First, for low frequencies the computational cost does not increase. The solution is approximated using mostly linear elements, and although it would be beneficial, the approximation space can only be coarsened by changing the mesh to use larger elements. This regime is therefore sub-optimal because large parts of the solutions are over-resolved.




Figure 19. Computational cost of the numerical models as a function of frequency. Dotted line, third-order frequency dependency.

Table III. Runtime for the calculation of a FRF for the car cavity test case on the frequency range [20, 2000] Hz with a resolution of 5 Hz.

Model	Matrix evaluation timing (s)	Total time (s)	Speed-up factor
FEM	12.0	122 780	—
FEMAO I	20.9	18 824	6.5
FEMAO II	21.4	7 467	16.4
FEMAO III	24.6	6 125	20.0

Secondly, an asymptotic regime is observed for high frequencies where the computational cost increases with frequency as $\mathcal{O}(\omega^3)$. Ideally, the input mesh should be designed for a given frequency range so as to operate completely in this asymptotic regime (that is, by avoiding mesh refinements, which may lead to over-discretized regions). Note that the $\mathcal{O}(\omega^3)$ scaling of the computational cost in the asymptotic regime is as expected. An adaptive scheme based on an h -FEM approach in which a separate input mesh is designed for each frequency by keeping kh constant would yield a similar frequency dependency. In essence, the adaptive approach does not modify the complexity of the standard finite element model with respect to the size of the model or the frequency, but it renders the computational process more efficient and robust.

Table III provides the overall runtime associated with the calculations of the FRFs shown in Figure 16. The high-order adaptive solutions are an order of magnitude faster than the standard FEM model. The FEM solution required 34 hours for the full frequency range, whereas the fastest adaptive solution (FEMAO III) required only 1 h and 42 min, corresponding to a speed-up factor of 20. The reported speed-up factors for the refined models FEMAO I and FEMAO II (16.4 and 6.5) are lower. In terms of memory requirements, the standard FEM model requires 33.8 Gb, while the optimal high-order model (FEMAO III) can handle the full frequency range in-core with only 7.9 Gb, which represents a reduction by a factor 4.3.

The CPU time for the creation of the mass and stiffness matrices is also reported in Table III. These results illustrate that for higher-order solutions evaluating the element matrices represents a larger proportion of the overall computational time. For the linear FEM, the matrix evaluation is a small operation in comparison with the factorization, with respectively 12 s against approximately 290 s, which represents a ratio of 1/25 (4 %). In contrast, the cost of one matrix evaluation for FEMAO III represents 40% of the overall solving time at 2 kHz (24.6 s against 60 s). Note however that this is an increase in relative term compared with the overall solution procedure, and the total cost of the model is reduced when increasing the polynomial order for a fixed accuracy (as shown in Figure 3). If a matrix evaluation was required for each individual frequency, it would represent a significant portion of the overall cost of the method. However, this cost can be factored out of the frequency loop through the use of hierarchical shape functions, allowing the matrix evaluation to be performed only once.

It could be argued that using a single mesh when calculating an FRF with the standard FEM is not optimal, and that at least several meshes should be prepared to cover the full frequency range. While this is true, in practice, it is not uncommon that a single mesh designed for the highest frequency is used for the full FRF in order to save time for the end user. And even with a mesh optimized for each frequency in the FRF, the high-order finite element would still outperform the standard FEM.

Because the high-order method leads to significant improvements in accuracy, it is possible, with the same computational resources, to tackle larger problems or higher frequencies. As an example, the solution at 5000 Hz (corresponding to a Helmholtz number of 150) was calculated using the mesh FEMAO I with a target accuracy of $E_T = 5\%$. The real part of pressure obtained from this computation is shown on a refined visualization mesh in Figure 20. This model involved 4 314 740 degrees of freedom, and the corresponding sparse matrix contained 173 533 294 non-zero entries. The matrix factorization required 210 Gb of memory and was performed in-core in approximately 1 h.




Figure 20. Colormap of the real part of the pressure (between $\pm 10^{-3}$ Pa) for the monopole source at 5000 Hz, obtained with the mesh FEMAO I ($E_T = 5\%$).

6. SUMMARY AND CONCLUSIONS

An efficient implementation of a high-order finite element method for the Helmholtz equation was presented. It is designed to tackle realistic acoustic engineering problems, typically large-scale 3D models to be solved over a large number of frequencies. The key aspect of this approach can be summarized as follows.

- To reduce the computational cost for large problems and/or at high frequencies, it is crucial to minimize the dispersion error of the discrete model so as to mitigate the pollution effect. This is achieved in the present method with the p -FEM approach using Lobatto shape functions, which allow to drastically improve the accuracy of the numerical model compared with conventional FEM. An assessment of the cost and accuracy of this method over a range of frequency, element sizes, and polynomial orders was conducted to identify the optimal parameters. Contrary to what has been reported elsewhere for smaller 2D models, the computational efficiency for large 3D models does not drop as the order is increased up to $p = 10$ (although it could be worth considering even higher polynomial orders).
- An added benefit of using a hierarchical set of shape functions such as the Lobatto polynomials is found when solutions are required over a range of frequencies. It is possible to compute the element matrices only once for the highest order, and then at each frequency, one can extract only the required portion of these matrices to assemble the global matrix. This is crucial as the cost of calculating high-order element matrices is significant due to the large number of quadrature points involved.
- A key feature of the proposed method is the use of a simple local error indicator to select *a priori* the polynomial order in each element. This is based on the numerical error for a 1D element. Despite its apparent simplicity, a number of tests have shown that it is able to maintain the accuracy of a p -FEM within specified target accuracy. However, because it is purely local, it does not account for the pollution effect. Addressing this issue would require more complex adaptive methods, in particular, *a posteriori* schemes.
- For calculations over a range of frequencies, the proposed approach requires only a single mesh, and the error estimator is able to select the suitable polynomial order in each element to maintain the efficiency and accuracy as the frequency is changed by identifying the lowest order required for the accuracy target. Best practice for preparing meshes suited for p -FEM were discussed. The element size is dictated by the high frequency of interest and the high

order available, and a simple guideline was proposed in (26). When using linear or quadratic description of the element geometry, as is the case with conventional FEM, smaller elements are still required to describe the geometry, which penalizes the performance of the high-order scheme. Several options for improving the geometry description of high-order methods have been proposed and will be the subject of future work [38, 39].

The application of this approach to a demanding industrial problem was discussed in detail by considering the prediction of the sound field in a 3D model of a car cabin over a large frequency range. Significant gains in performance were observed compared with a standard finite element model. For this test case, the memory requirements is reduced by a factor 4.3. The calculation of a complete FRF with a fine frequency resolution is performed 10 times faster and remains accurate over a wider range of frequencies. It was found also that the numerical model is particularly robust to significant changes in the element size and that the error indicator is able to compensate for these changes to maintain the same level of accuracy. Finally, the computational cost (memory and time) is found to scale like $\mathcal{O}(\omega^3)$, which is similar to an h -FEM approach where a separate mesh would be designed specifically for each frequency of interest (by keeping kh constant). With the p -FEM approach used here, this scaling is achieved with a single mesh and the proposed adaptive scheme, which simplifies the solution procedure significantly.

The current study focuses solely on direct solution procedure, and future work should consider the application of robust iterative solution algorithms as an alternative to limit the memory footprint of the method. In particular, the conditioning of high-order p -FEM models and its impact on the performance of iterative schemes should be investigated. Other possible research directions include reduced-order methods to accelerate the computation of frequency responses over a broad frequency range.

ACKNOWLEDGEMENTS

The provision of EngD funding from EPSRC grant EP/G036896/1, the Industry Doctoral Training Centre in Transport and the Environment is gratefully acknowledged. The authors also acknowledge the European Commission (EC) for the support of the CRANE project (www.crane-eid.eu), supervised by Michel Tournour from Siemens Industry Software.

REFERENCES

1. Babuška IM, Sauter SA. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? *SIAM Review* 2000; **42**(3):451–484.
2. Thompson LL, Pinsky PM. A Galerkin least-squares finite element method for the two-dimensional Helmholtz equation. *International Journal for Numerical Methods in Engineering* 1995; **38**(3):371–397.
3. Harari I, Magoulés F. Numerical investigations of stabilized finite element computations for acoustics. *Wave Motion* 2004; **39**(4):339–349.
4. Ortiz P, Sanchez E. An improved partition of unity finite element model for diffraction problems. *International Journal for Numerical Methods in Engineering* 2001; **50**(12):2727–2740.
5. Farhat C, Harari I, Franca LP. The discontinuous enrichment method. *Computer Methods in Applied Mechanics and Engineering* 2001; **190**(48):6455–6479.
6. Cessenat O, Després B. Application of an ultra weak variational formulation of elliptic PDEs to the two-dimensional Helmholtz problem. *SIAM Journal in Numerical Analysis* 1998; **35**:255–299.
7. Gabard G, Gamallo P, Huttunen T. A comparison of wave-based discontinuous Galerkin, ultra-weak and least-square methods for wave problems. *International Journal for Numerical Methods in Engineering* 2011; **85**(3):380–402.
8. Wang D, Tezaur R, Toivanen J, Farhat C. Overview of the discontinuous enrichment method, the ultra-weak variational formulation, and the partition of unity method for acoustic scattering in the medium frequency regime and performance comparisons. *International Journal for Numerical Methods in Engineering* 2012; **89**(4):403–417.
9. Petersen S, Dreyer D, von Estorff O. Assessment of finite and spectral element shape functions for efficient iterative simulations of interior acoustics. *Computer Methods in Applied Mechanics and Engineering* 2006; **195**(44-47):6463–6478.
10. Vos PEJ, Sherwin SJ, Kirby RM. From h to p efficiently: implementing finite and spectral/hp element methods to achieve optimal performance for low- and high-order discretisations. *Journal of Computational Physics* 2010; **229**(13):5161–5181.
11. Cantwell CD, Sherwin SJ, Kirby RM, Kelly PHJ. From h to p efficiently: strategy selection for operator evaluation on hexahedral and tetrahedral elements. *Computers and Fluids* 2011; **43**(1):23–28.

12. Giorgiani G, Modesto D, Fernández-Méndez S, Huerta A. High-order continuous and discontinuous Galerkin methods for wave problems. *International Journal for Numerical Methods in Fluids* 2013; **73**:883–903.
13. Huerta A, Angeloski A, Roca X, Peraire J. Efficiency of high-order elements for continuous and discontinuous Galerkin methods. *International Journal for Numerical Methods in Engineering* 2013; **96**:529–560.
14. Babuška I, Guo BQ. The h , p and h - p version of the finite element method: basis theory and applications. *Advanced Engineering Software* 1992; **15**(3-4):159–174.
15. Thompson LL. A review of finite-element methods for time-harmonic acoustics. *The Journal of the Acoustical Society of America* 2006; **119**(3):1315–1330.
16. Bouillard P, Ihlenburg F. Error estimation and adaptivity for the finite element method in acoustics: 2D and 3D applications. *Computer Methods in Applied Mechanics and Engineering* 1999; **176**(1-4):147–163.
17. Irimie S, Bouillard P. A residual a posteriori error estimator for the finite element solution of the Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering* 2001; **190**(31):4027–4042.
18. Rachowicz W, Demkowicz L. An hp -adaptive finite element method for electromagnetics – part II: A 3D implementation. *International Journal for Numerical Methods in Engineering* 2002; **53**(1):147–180.
19. Demkowicz L, Kurtz J, Pardo D, Paszynski M, Rachowicz W, Zdunek A. *Computing with hp-Adaptive Finite Elements, Vol. 2: Frontiers: Three Dimensional Elliptic and Maxwell Problems with Applications*. Chapman and Hall/CRC, 2007.
20. Demkowicz L, Kurtz J, Qiu F. hp -adaptive finite elements for coupled multiphysics wave propagation problems. In *Computer Methods in Mechanics*, vol. 1, Kuczma M, Wilmski K (eds.), Advanced Structured Materials. Springer: Berlin Heidelberg, 2010; 19–42.
21. Díez P. A unified approach to remeshing strategies for finite element h -adaptivity. *Computer Methods in Applied Mechanics and Engineering* 1999; **176**(1-4):215–229.
22. Šolín P, Segeth K, Doležal I. *Higher-Order Finite Element Methods*. Chapman & Hall, 2003.
23. Abramowitz M, Stegun IA. *Handbook of Mathematical Functions* (10th edition), Applicable Mathematics Series. National Bureau of Standards, US Government Printing Office, 1972.
24. Szabó B, Babuška I. *Introduction to Finite Element Analysis: Formulation, Verification and Validation*. Wiley, 2011.
25. Karniadakis G, Sherwin S. *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, 2013.
26. Dubiner M. Spectral methods on triangles and other domains. *Journal of Scientific Computing* 1991; **6**(4):345–390.
27. Grundmann A, Moeller HM. Invariant integration formulas for the n -simplex by combinatorial methods. *SIAM Journal Numerical Analysis* 1978; **15**:282–290.
28. Osei-Kuffuor D, Saad Y. Preconditioning Helmholtz linear systems. *Applied Numerical Mathematics* 2010; **60**(4):420–431.
29. Amestoy PR, Duff IS, L'Excellent JY. Multifrontal parallel distributed symmetric and unsymmetric solvers. *Computer Methods in Applied Mechanics and Engineering* 2000; **184**(2-4):501–520.
30. Ihlenburg F, Babuška I. Finite element solution to the Helmholtz equation with high wave number – Part II: the h - p -version of the FEM. *SIAM Journal of Numerical Analysis* 1997; **34**(1):315–358.
31. Ainsworth M. Discrete dispersion relation for hp -version finite element approximation at high wave number. *SIAM Journal on Numerical Analysis* 2004; **42**(2):553–575.
32. Beriot H, Gabard G, Prinn A. On the performance of high-order FEM for solving large-scale industrial acoustic problems. *International Congress on Sound and Vibration*, Bangkok, Thailand, 2014.
33. Deraemaeker A, Babuška I, Bouillard P. Dispersion and pollution of the FEM solution for the Helmholtz equation in one, two and three dimensions. *International Journal for Numerical Methods in Engineering* 1999; **46**(4):471–499.
34. Huerta A, Rodríguez-Ferran A, Díez P, Sarrate J. Adaptive finite element strategies based on error assessment. *International Journal for Numerical Methods in Engineering* 1999; **46**(10):1803–1818.
35. Prinn A. Efficient finite element methods for aircraft engine noise prediction. *Ph.D. Thesis*, University of Southampton, 2014.
36. Manual U. *LMS Virtual.Lab Acoustics r10*. LMS International. (Available from: <http://www.lmsintl.com>, 2011).
37. Lenzi MS, Lefteriu S, Beriot H, Desmet W. A fast frequency sweep approach using padé approximations for solving Helmholtz finite element models. *Journal of Sound and Vibration* 2013; **332**(8):1897–1917.
38. Toulorge T, Geuzaine C, Remacle JF, Lambrechts J. Robust untangling of curvilinear meshes. *Journal of Computational Physics* 2013; **254**(0):8–26.
39. Sevilla R, Fernández-Méndez S, Huerta A. NURBS-enhanced finite element method (NEFEM). *International Journal for Numerical Methods in Engineering* 2008; **76**(1):56–83.
40. Modesto D, Zlotnik S, Huerta A. Proper generalized decomposition for parameterized Helmholtz problems in heterogeneous and unbounded domains: Application to harbor agitation. *Computer Methods in Applied Mechanics and Engineering* 2015; **295**:127–149.