

An Adaptive, High-Order Finite-Element Method for Convected Acoustics

G. Gabard^{*a}, H. Bériot^{†b}, A.G. Prinn^{‡a} and K. Kucukcoskun^{§b}

^a*Institute of Sound and Vibration Research, University of Southampton, UK*

^b*Siemens Industry Software NV, Leuven, Belgium*

An improved finite element method for predicting sound propagation in non-uniform flows is proposed. Sound waves are described by the linearised potential theory solved in the frequency domain. Solutions are calculated using the p -FEM method with high-order, hierarchic shape functions which results in a drastic improvement in computational efficiency. The memory and time requirements for solving large-scale problems are significantly reduced compared to standard finite element methods. An additional feature presented in this paper is an adaptive scheme to select the optimal interpolation order in each element so as to achieve a prescribed accuracy. This greatly simplifies the preparation of the numerical model. When performing a series of calculations at different frequencies there is no need to adjust the mesh to maintain a sufficient resolution. Instead the adaptive p -FEM method adjusts the interpolation order automatically to ensure an accurate solution is obtained. The performance of the method is demonstrated for three-dimensional test cases for noise radiation from a turbofan intake. Guidelines for preparing models using this adaptive, high-order approach are also discussed.

I. Introduction

Computational models predicting the propagation and radiation of noise from aero-engines remain a crucial tool to support the acoustic design of various components of turbofan engines, a good example being the optimisation of acoustic liners on the nacelle of the engine.¹ Arguably the most widespread scheme for this purpose is the linearised potential equation solved in the frequency domain with standard finite elements. This approach forms the basis for commercial codes routinely used in industry. One of the main limitations of this approach is the computational resources required since the size of the problem and the upper frequency limit that can be considered are dictated by the memory available. In practice a range of frequencies are considered and a separate simulation has to be performed for each of these. In addition, for automated design optimisation it is necessary to perform a large number of calculations to explore the design space, multiplying even further the resources required.^{2,3} Improvements in the efficiency of current computational schemes are required if these tools are to be used more systematically to tackle large problems at high frequencies.

The present paper proposes an improved finite element method delivering a significant reduction in computational cost (both memory and time) and discusses its use for realistic, large-scale problems in an industrial context. The first feature of the proposed method is the use of high-order shape functions (in this work from order 1 to 10), which lead to a drastic reduction in the dispersion error in the numerical model, thus allowing the use of much larger elements. The second feature is to use a hierarchic family of Lobatto shape functions. In this way it is possible to re-use the same element matrices for a number of calculations at different frequencies. The third feature of the proposed method is to adapt the order of interpolation in

^{*}Senior Researcher. Now at LAUM, Le Mans Université, 72085 Le Mans, France. gwenael.gabard@univ-lemans.fr. Senior AIAA Member.

[†]Senior Project Leader. Siemens Industry Software NV, Interleuvenlaan 68, 3001, Leuven, Belgium. hadrien.beriot@siemens.com.

[‡]Research Fellow, ISVR, University of Southampton, SO17 1BJ, UK, AIAA Member.

[§]Product Manager. Siemens Industry Software NV, Interleuvenlaan 68, 3001, Leuven, Belgium.

each element so as to achieve a target error using a simple error estimator. This greatly simplifies the use of this method for non-uniform flows and for calculations over multiple frequencies.

While high-order methods have been used for time-domain models of flow acoustics, mainly with the finite-difference or discontinuous Galerkin methods, they bring about stability issues associated with non-uniform flows⁴ and impedance conditions.^{5,6} The frequency-domain approach is more robust but its use with high-order techniques has been limited⁷ and adaptive schemes have yet to be proposed and validated for this approach.

The propagation model for waves in a potential base flow is introduced in the next section, followed by a description of the high-order finite element methods in section III, including the adaptive scheme. A generic turbofan test case and the corresponding mean flow configurations and source models are introduced in section IV. Finally, examples of results are presented in section V to illustrate the benefit of the proposed method compared to the standard finite element method and to discuss best practice for this high-order approach.

II. Sound Propagation Model

We model the propagation of sound waves in a steady, potential base flow using the linearised potential equation:

$$\rho_0 \frac{D_0}{Dt} \left(\frac{1}{c_0^2} \frac{D_0 \phi}{Dt} \right) - \nabla \cdot (\rho_0 \nabla \phi) = 0, \quad (1)$$

where $D_0/Dt = \partial/\partial t + \mathbf{u}_0 \cdot \nabla$ represents the material derivative in the mean flow \mathbf{u}_0 . The mean sound speed and density are denoted by c_0 and ρ_0 . With this scalar equation the sound field is described by the potential ϕ from which the acoustic velocity is derived ($\mathbf{u} = \nabla \phi$). The acoustic pressure and density can also be calculated from ϕ using the relations $p = -\rho_0 D_0 \phi / Dt$ and $\rho = p / c_0^2$.

The wave equation (1) is solved in the frequency domain by assuming an implicit time dependence $e^{+i\omega t}$. This model represents the convective effect of the flow on the sound field, but does not include the refraction effect that would be induced by a sheared mean flow. It is however well suited for applications such as the radiation of sound from turbofan intakes.

For the purpose of predicting noise radiation from aeroengines various boundary conditions should be considered. For hard-wall surfaces the normal acoustic velocity vanishes: $\mathbf{u} \cdot \mathbf{n} = \partial \phi / \partial n = 0$ where \mathbf{n} is the unit normal pointing into the surface.

Surfaces that are acoustically treated are characterised by their acoustic impedance $Z(\omega)$, assuming a locally reacting liner. In the presence of a base flow we use the Myers⁸ impedance condition to describe the coupling between the liner and the sound waves in the presence of an infinitely thin boundary layer above the liner. This boundary condition reads:

$$\frac{\partial \phi}{\partial n} = \frac{-\rho_0}{i\omega Z} \left(\frac{D_0}{Dt} - \mathbf{n} \cdot \frac{\partial \mathbf{u}_0}{\partial n} \right) \frac{D_0 \phi}{Dt}. \quad (2)$$

To represent sound radiation to the far field, Perfectly Matched Layers (PML) are implemented around the boundary of the domain to absorb any outgoing waves and remove spurious reflections.^{9,10} The PML can also be used to generate a sound field entering the computational domain.^{11,12} This so-called ‘active’ PML is used to define the source of sound radiating from the engine, typically from the fan. These sources are generally described in terms of acoustic duct modes. Although not required for this particular application, another possibility is to extract the aerodynamic noise sources from an unsteady CFD calculation, using an appropriate acoustic analogy. The source models can either be based on flow wall pressure, which is used to define equivalent acoustic boundary conditions,¹³ or on volume flow velocity fluctuations, to define equivalent quadrupole sources.¹⁴

The calculation of the sound radiated to the far field is performed using a Kirchhoff surface that allows to extrapolate the near-field solution to the far field. This Kirchhoff formulation includes the effects of a uniform mean flow outside of the control surface, through the use of the Lorentz transformation.¹⁵

III. Adaptive High-Order Finite Elements

III.A. High-order FEM formulation

A variational formulation for the scalar wave equation (1) is obtained by applying the standard Galerkin approach and requiring that the residual of (1) is orthogonal to a test function ψ . After integration by parts the following weak formulation is obtained:

$$\int_{\Omega} \rho_0 \overline{\nabla \psi} \cdot \nabla \phi - \frac{\rho_0}{c_0^2} \frac{\overline{D_0 \psi}}{Dt} \frac{D_0 \phi}{Dt} d\Omega + \int_{\partial\Omega} \overline{\psi} \frac{\rho_0}{c_0^2} \frac{D_0 \phi}{Dt} \mathbf{u}_0 \cdot \mathbf{n} - \overline{\psi} \rho_0 \frac{\partial \phi}{\partial n} dS = 0, \quad (3)$$

where Ω is the computational domain and $\partial\Omega$ its boundary. The overbar denotes the complex conjugate.

The various boundary conditions described in the previous section are implemented by modifying the boundary integral accordingly. The implementation of the Myers condition (2) proposed by Eversman¹⁶ is used.

The variational formulation (3) is discretised using the high-order finite element method.^{17,18} The so-called p -FEM approach leads to a drastic reduction of memory requirements and run times for two-dimensional as well as three-dimensional Helmholtz problems when compared to conventional FEM.^{19,20} The p -FEM method has also shown to deliver comparable, and in some cases even superior, performance compared to a high-order Trefftz method²¹. The key benefits of p -FEM (e.g. low dispersion error, exponential p -convergence) are also retained in the presence of background mean flows.²²

The domain Ω is approximated in the usual way with a tessellation of non-overlapping elements. On each element the approximation of the solution ϕ is constructed as a linear combination of high-order shape functions $\Psi_n(\mathbf{x})$ and the Lobatto shape functions are used in this work.¹⁷ In three dimensions, they combine four different types of shape functions and the solution is written:

$$\phi(\mathbf{x}) = \sum_{n=1}^{N_n(p)} \phi_n^n \Psi_n^n(\mathbf{x}) + \sum_{n=1}^{N_e(p)} \phi_n^e \Psi_n^e(\mathbf{x}) + \sum_{n=1}^{N_f(p)} \phi_n^f \Psi_n^f(\mathbf{x}) + \sum_{n=1}^{N_b(p)} \phi_n^b \Psi_n^b(\mathbf{x}), \quad (4)$$

where the degrees of freedom ϕ_n are the complex amplitudes of each shape function. p is the overall polynomial order of the approximation. The nodal shape functions $\Psi_n^n(\mathbf{x})$ are the standard linear shape functions whose associated degrees of freedom are the values of the solution at the nodes of the elements. The edge and face shape functions, $\Psi_n^e(\mathbf{x})$ and $\Psi_n^f(\mathbf{x})$, are defined on the corresponding geometric entities. The fourth category are the so-called ‘bubble’ shape functions $\Psi_n^b(\mathbf{x})$ which vanish on the boundary of each element. For instance for a tetrahedron, the numbers of shape functions are $N_n = 4$, $N_e = 6(p-1)$, $N_f = 2(p-2)(p-1)$ and $N_b = (p-3)(p-2)(p-1)/6$.

A key property of the high-order FEM is that it is possible to eliminate the degrees of freedom associated with the bubble shape functions since these degrees of freedom do not interact with neighboring elements. This technique, called condensation, allows to reduce the size of the global linear system before it is assembled and solved, hence reducing the memory requirements.¹⁸ The Lobatto shape functions provide a number of additional benefits. They exhibit good performance and conditioning for Helmholtz problems.³ Their ‘hierarchical’ nature implies that the shape functions at order p form a subset of the shape functions at order $p+1$. This key property implies that p -FEM can easily handle local order variations: two adjacent elements can coexist with different polynomial orders. This is in contrast with high-order Lagrange interpolation functions which cannot handle easily different orders for two contiguous elements. To ensure the continuity of the numerical solution the order p for a given edge or face shared by several elements is chosen based on a so-called conformity rule.¹⁷ From the polynomial orders defined for each element, one can assign a unique interpolation order to each edge and to each face of the mesh, based on a given conformity rule.¹⁷ In this work the so-called ‘maximum rule’ is used: the polynomial order assigned to a given face or edge in the mesh is equal to the maximum order of all elements sharing this face or edge. This approach guarantees that the polynomial order of any shape function associated with a given element is consistent with its order of approximation. This approach has been shown to yield robust approximations, even in the presence of highly non uniform meshes.²⁰

These features of p -FEM with Lobatto shape functions are particularly useful when simulating aero-engine noise propagation. Firstly, when applying the standard finite element method for problems with strongly non-uniform flows it is necessary to refine the mesh in the regions where the mean flow velocity is high since in these regions the acoustic wavelength can be significantly reduced by the flow convective




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

effect.¹ In practice this is addressed by manually refining the mesh in these regions. Secondly, typical use cases of these methods involves performing not one but several calculations for every frequency of interest (typically the center frequencies of the third-octave bands). The requirements in terms of mesh resolution can vary significantly between these frequencies. One has to generate several meshes to maintain a sufficient accuracy at all frequencies while minimising the computational cost, which is time consuming for the user. Another option is to generate a single fine mesh capable of resolving the complete range of frequencies, which drastically increases the time and computational resources needed for the analysis. The use of a variable polynomial order in each element provides a way to address these two issues. Using an adaptivity rule to define the polynomial order in each element it is possible to ensure that the numerical error on each element is controlled locally in terms of the element size, frequency and flow properties.

Calculating and assembling the matrices from every element yield a sparse linear system of equations which is solved using the fast, parallel multi-frontal solver MUMPS.^{23,24}

III.B. A-priori error indicator

Bériot et al. introduced an *a priori* error indicator for the p -FEM method applied to the Helmholtz equation.²⁰ Despite its apparent simplicity, this indicator was shown to efficiently control the global numerical error in realistic two-dimensional and three-dimensional problems with strong mesh non-uniformities.

With this approach, the numerical error incurred on a given element is estimated *a priori*, based on the numerical error measured on a single, one-dimensional element with equivalent wavenumber k and length h . For this single element, appropriate Robin boundary conditions are applied so that a single wave of the form $e^{i\omega t - ikx}$ is generated and leaves the computational domain without spurious reflection. The one-dimensional Helmholtz variational formulation is solved and the relative L^2 error, denoted ε_{L^2} , is measured. This error is a function of only the polynomial order p and the mesh resolution kh , as shown in Figure 1.

The adaptive approach consists in selecting the interpolation order p in each element so as to remain below a target L^2 error E_T in the 1D element. In practice, one does not need to solve the 1D model for every element of the mesh, since the order $p(kh, E_T)$ required to have $\varepsilon_{L^2} < E_T$ can be tabulated in terms of the mesh resolution kh and error level E_T . This is illustrated in Table 1 for three levels of accuracy $E_T = 15\%$, 5% and 0.5% . This table can be used to directly determine the element order based on the local values of k and h .

Note that different definitions of E_T have been considered (using the L^2 or H^1 norms), as well as different definitions of the equivalent length h for the 1D model (minimum, maximum or average values of the edge lengths). Based on the detailed comparisons in Reference 3, the length h of the 1D element is defined as the average edge length of the actual element and the target error is specified in terms of the L^2 error on the 1D element.

The proposed error indicator was originally designed for Helmholtz problems.²⁰ We show here that it

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

Table 1. Values of kh required to achieve a given accuracy at a given order p .

can be extended to the convected Helmholtz equation by generalizing the definition of the mesh resolution kh used in Table 1. When a base flow is present, the effective acoustic wavelength varies with the angle of propagation relative to the flow direction. For problems of practical relevance, upstream and downstream waves can coexist in the solution and the choice of a grid resolution should rely on the shortest possible wavelength, i.e. the upstream wave. The mesh resolution is therefore measured using

$$kh = \frac{\omega}{c_0(1 - M)} h \quad (5)$$

where $M = \|\mathbf{u}_0\|/c_0$ is the local Mach number of the base flow. This definition can lead to over-resolved solutions in areas where only waves propagating downstream are present. However, previous analysis^{22,25} indicates that a sufficient resolution for the upstream wave is required in order to ensure the accuracy of the discrete model, even when only downstream waves are present. Otherwise, the FEM solutions are likely to suffer from strong aliasing effects.

III.C. Efficient multi-frequency calculations

As mentioned above, a hierarchical set of shape functions is particularly useful 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 for each frequency one can extract 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.

This approach, combined with the error indicator presented in Section III.B leads to an efficient high-order adaptive strategy, which is outlined below. In what follows, the superscript i denotes the quantity computed at a given frequency f_i .

1. For each frequency f_i the *a priori* error indicator (1) allows to identify the order p_e^i required in each element e to achieve the user-defined error target E_T .
2. For each element e the highest polynomial order p_e^{\max} required across the frequency range is identified.
3. Using the maximum conformity rule described above, the highest polynomial orders used for each edge and face in the mesh is determined.
4. For each element e , the element matrices for the highest order p_e^{\max} are evaluated and stored. Thanks to the hierarchical structure of the Lobatto basis, these matrices also contain the smaller matrices that will be used for the lowest orders.
5. At each frequency i , the submatrices corresponding to the local order p_e^i are 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 construction of the discrete system negligible for problems involving a large number of frequencies.

III.D. Mesh requirements

As explained above the use of an adaptive high-order scheme changes significantly the requirements and best practice associated with the finite-element meshes. When using standard, fixed-order FEM methods for flow acoustics (typically with quadratic elements), the mesh resolution is dictated by the frequency, the mean flow velocity which influences the effective wavelength, and the geometry of the problem at hand. In practice this implies generating several meshes to cater for different mesh resolution at different frequencies, based

on a guideline such as 8 or 10 points per wavelength. One also has to adjust the local mesh size to the local mean flow properties to ensure that the sound field is sufficiently resolved throughout the computational domain. This is why generating meshes is often the most time consuming part of the simulation process.

With the proposed high-order approach the increase in resolution required for higher frequencies and flow speed is automatically provided by the adaptive change of interpolation order in each element. There are still constraints on the element size associated with the maximum order available (in this work $p_{\max} = 10$). The mesh resolution should be sufficient to obtain accurate results at the highest frequency and for the highest flow speed when using $p = p_{\max}$. Based on the definition of the mesh resolution (5), for a maximum frequency f_{\max} , the following guideline for the element size can be inferred

$$h \leq C_{E_T} \frac{c_0(1 - M)}{f_{\max}}, \quad (6)$$

where the constant C_{E_T} is adjusted according to the required accuracy and the maximum order available. From the tabulated values in Table 1 for $p_{\max} = 10$ one finds that $C_{5\%} = 15.4/2\pi = 2.44$ and $C_{0.5\%} = 12.2/2\pi = 1.94$. Following this guideline, a single mesh can be generated for a range of calculations at different frequencies.

Another constraint is to maintain an accurate representation of the problem geometry. For this purpose, using standard, linear shape functions for the geometry of the elements should be avoided. For the examples shown in this paper, elements with quadratic geometries proved sufficient to use coarse meshes while still representing the problem geometry accurately. The generation of finite-element meshes with high-order geometries is currently the subject of active research.^{26,27} An alternative approach is to incorporate the exact geometric description, based on B-splines for instance, directly into the finite element model^{28,29}.

IV. Application to a Generic Turbofan Intake

IV.A. Description of the test case

A generic turbofan intake is used as a test case to benchmark the proposed numerical method. This section provides the relevant information concerning this test case and more details are available in Prinn³ and Mustafi.²

The geometry of the axi-symmetric intake is shown in Figure 2. The main components are the spinner and the nacelle which have been positioned and scaled so that the fan radius is $R = 1.2\text{m}$ and the notional fan face is located at $z = 0$. The outer radius of the nacelle is 1.5m . For the purpose of the mean flow calculation and the acoustic simulations a straight, annular ‘fan duct’ is added downstream of the fan face. For the flow calculation this duct is used to obtain a uniform flow on the outflow condition.




Figure 2. Axi-symmetric geometry of the intake with the spinner (blue), nacelle (red), fan face (dashed line), fan duct (green), source plane (magenta), extension of the nacelle (cyan). The end points of the liner are shown by the black symbols

Three standard flow configurations are considered: sideline (at take-off when the engines are at full power), cutback (at take-off with reduced rate of climb and engine power) and approach (low speed condition before

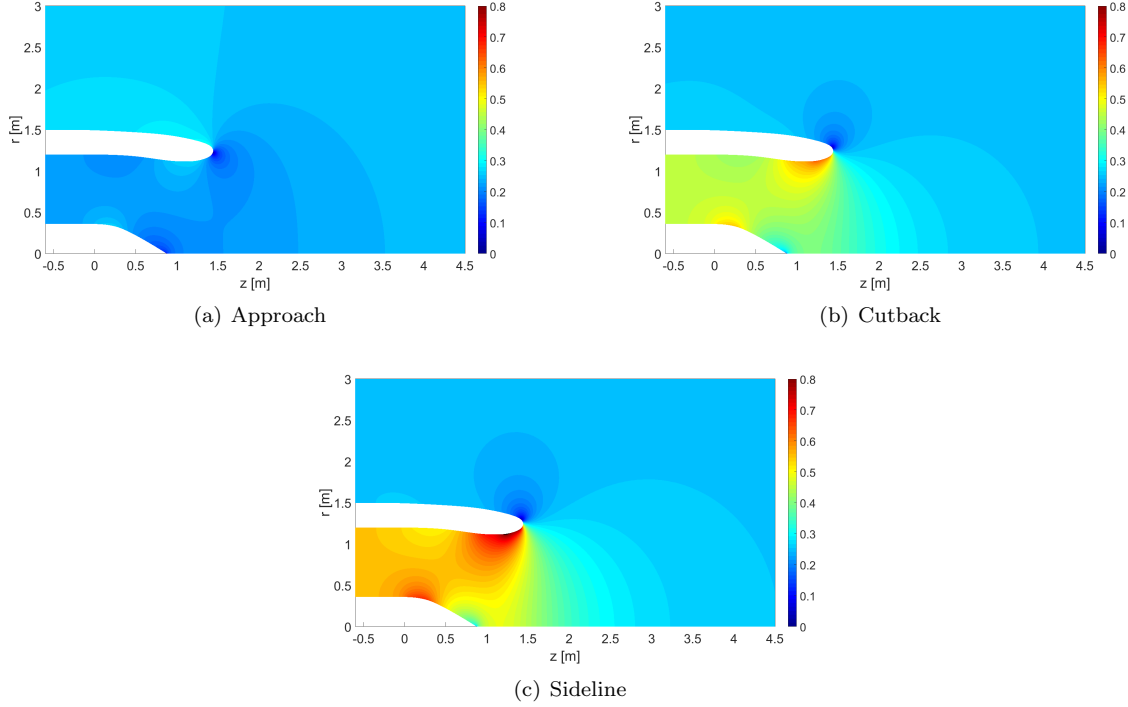


Figure 3. Mean flow Mach number for the different flow configurations.

landing). The three flow configurations differ only by the Mach number at the fan face with $M_f = 0.55$, 0.45 and 0.22 for the sideline, cutback and approach conditions, respectively. For simplicity, the conditions at infinity are as follows: sound speed $c_\infty = 340$ m/s, density $\rho_\infty = 1.2$ kg/m³ and Mach number $M_\infty = u_\infty/c_\infty = 0.25$. The local flow Mach number for these three configurations is shown in Figure 3. These have been obtained using a potential, compressible, inviscid flow model.

Tonal noise is present at the Blade Passing Frequency (BPF) and its harmonics. The Blade Passing Frequencies for the approach, cutback and sideline cases are 700 Hz, 1050 Hz and 1300 Hz, respectively. The BPF tones consist of a multi-mode component and a single-mode component. The multi-mode component consists of all the cut on modes at that frequency, while the single mode component consists of a single rotor-locked mode which is the first radial mode of the azimuthal order nB with $B = 24$ the number of fan blades and n the harmonic of the BPF.

The inner wall of the nacelle is acoustically lined with a locally reacting treatment. The end points of the liner are shown in Figure 2 and are defined by the axial positions $z = 0.08R = 0.096$ m and $z = 0.98R = 1.176$ m, corresponding to an axial length of $0.9R = 1.08$ m. The impedance function used for this liner is that of the standard single-degree-of-freedom perforate liner, see equation 5 in reference 30. The non-dimensional resistance and mass reactance of the perforate are $R_p = 2$ and $M_p = 0.023$. The cavity depth is 0.02 m.

IV.B. Numerical model

The computational domains used for the simulations are shown in Figure 4 where the different boundary conditions are also highlighted. The noise source is introduced through a duct mode boundary condition which is defined on the source plane located at $z = -0.4$ m. The unbounded region in front of the intake is truncated at a boundary Γ . On this boundary, an Automatically Matched Layer³¹ (AML) is used, resulting in the automatic addition of a PML region outside Γ .

For this particular application, the location of the truncation surface Γ is of particular importance as it influences the accuracy of the solution in two different ways. Firstly the performance of the PML as a non-reflecting region improves when it is located farther away from the source. Placing the absorbing region too close to the nacelle lips may lead to spurious reflections inside the domain. Secondly, the surface Γ is




Figure 4. Computational domain and boundary conditions: duct modes (red), liner (yellow), AML surface (blue). Two control lines are also shown.



Figure 5. Two examples of truncation surface Γ defining the short domain (red) and large domain (green).

	Sideline			Cutback			Approach		
	c_0	ρ_0	$\ \mathbf{u}_0\ $	c_0	ρ_0	$\ \mathbf{u}_0\ $	c_0	ρ_0	$\ \mathbf{u}_0\ $
short domain	333.7	1.09	79.82	335.4	1.12	79.29	340.5	1.21	82.49
large domain	338.4	1.17	83.82	339.0	1.18	83.98	340.5	1.21	84.50

Table 2. Flow properties obtained by averaging on the truncation surface. These should be compared against the far-field values $c_\infty = 340$ m/s, $\rho_\infty = 1.2$ kg/m³ and $u_\infty = 85$ m/s.

used as a control surface to apply the Kirchhoff integral and calculate the solution in the far field.¹⁵ This integral formulation is based on the assumption that the waves propagate in a uniform flow. Ideally, the mean flow should be uniform on the surface Γ and equal to the far field, but in practice this is not the case and the variation of mean flow properties on Γ can change slightly the far-field predictions. There is therefore a trade-off between reducing the size of the domain to minimise the computational cost while maintaining an sufficient level of accuracy in the predictions.

For the results presented below, two truncation surfaces have been used, as shown in Figure 5, resulting in two different computational domains, with volumes 14.5 m³ and 30.5 m³. In practice, the mean flow properties used in the Kirchhoff formulation are obtained by averaging the nodal values on the surface Γ . Indicative averaged mean flow properties used in the Kirchhoff integral for the short and large domains are given in Table 2 for the three flow configurations. These results are indicative because they are based on a nodal averaging, which may slightly vary from one mesh to the other. However, for the meshes used in this study, the density of nodes is such that these differences are negligible.

These averaged quantities indicate that the short domain can be used for the approach mean flow configuration with less than 3% mismatch between u_∞ and the averaged value of $\|\mathbf{u}_0\|$. However, more significant discrepancies are observed for the short domain in the case of the cutback and sideline configurations, where the mismatch exceeds 6%. Numerical experiments indicate that this may lead to differences in the predicted sound directivity, and therefore, for these mean flow conditions, the larger domain should be preferred.

Once the truncation surface is created, the acoustic mesh can be generated based on the guideline given in (6) for the element size. The next step involves mapping and interpolating the mean flow values obtained from the potential, compressible, inviscid flow model onto the acoustic mesh using the existing Virtual.Lab interpolation procedure.³¹

V. Numerical Results

The objective of this section is to examine the benefits of the high-order adaptive solver in terms of performance (computational time and memory requirements), but also in terms of model creation process. All computations reported in this section were performed for fully three-dimensional models using the FEMAO solver (Finite-Element Method with Adaptive Order). This solver is an implementation of the present




Figure 6. Tetrahedral meshes for the short computational domain used for the approach configuration at 2 BPF.

adaptive p -FEM approach written in Fortran 2003 and using the MUMPS library to solve large, sparse linear systems. This solver is now used in an industrial context as part of the simulation package Virtual.Lab Acoustics.³¹ All computations reported hereafter were performed on a Dell Poweredge R730XD equipped with two Intel Xeon E5-2667 CPUs with 16 cores running at 3.20GHz and with 384 GB of memory.

V.A. Verification against standard FEM

The adaptive, high-order method is first benchmarked against a standard finite-element method using quadratic shape functions for all the elements. The approach configuration at the first harmonic of the blade passing frequency (1400 Hz) is considered. Based on the following definition of the Helmholtz number³²

$$kR = \frac{\omega R}{c_0(1 - M_f)} \quad (7)$$

this corresponds to a kR of 40, which provides a representative example of challenging cases for three-dimensional FE models. In this definition of the Helmholtz number, R and M_f are the fan radius and fan face Mach number, respectively, as defined in Section IV.A.

The short computational domain shown in Figure 5 is chosen, as the mean flow properties are deemed sufficiently uniform on the truncation surface in this case to calculate the far-field directivity. Three tetrahedral meshes are generated for this particular configuration. These include a refined mesh for the FEM reference solution, and two coarser meshes for the high-order simulations. For all three meshes, the tetrahedral elements use quadratic shape functions for their geometries so as to represent more accurately the curved boundaries of the computational domain.

Axial cross sections of the three meshes are shown in Figure 6. For the FEM mesh, a minimal resolution of 8 nodes per wavelength is applied, which leads to a typical element size of $h = 46$ mm based on the wave propagating upstream in the far field ($M = 0.25$). The resulting mesh contains 2 996 168 nodes and 2 208 549 quadratic tetrahedral elements. The first coarse mesh, denoted FEMA0 I, is generated from the same surface mesh of the nacelle as the FEM model. Larger elements of typical size $h = 300$ mm are then used in the bulk of the domain and on the outer AML surface. This first mesh will allow to demonstrate the benefits of using higher-order approximation while retaining exactly the same geometric accuracy as the FEM model. A second coarse mesh, denoted FEMA0 II, is created with a typical mesh size of $h = 300$ mm on the surface boundaries and inside the volume. Mesh refinements are applied close to the hard-soft and soft-hard lining transitions ($h = 80$ mm) in order to better capture the strong pressure gradients that may appear in this region. The main properties of the three meshes are outlined in Table 3. The high-order coarse meshes contain far less nodes and elements than their FEM counterpart.

To assess the results of the high-order models against the standard finite element solution, two duct modes are considered individually: namely the mode (12,1) and the mode (24,1), each with unit power. Their cut-off frequencies are 611 Hz and 1160 Hz, which makes them representative of a well cut-on and a close to cut-off mode at 2BPF (1400 Hz), respectively.

Mesh	Nodes	Elements	h_{\max}	h_{\min}
FEM	2 996 168	2 208 549	46 mm	40 mm
FEMAO I	320 978	215 693	300 mm	40 mm
FEMAO II	50 186	33 247	300 mm	80 mm

Table 3. Properties of the meshes used for the comparison using the approach condition at 2 BPF.

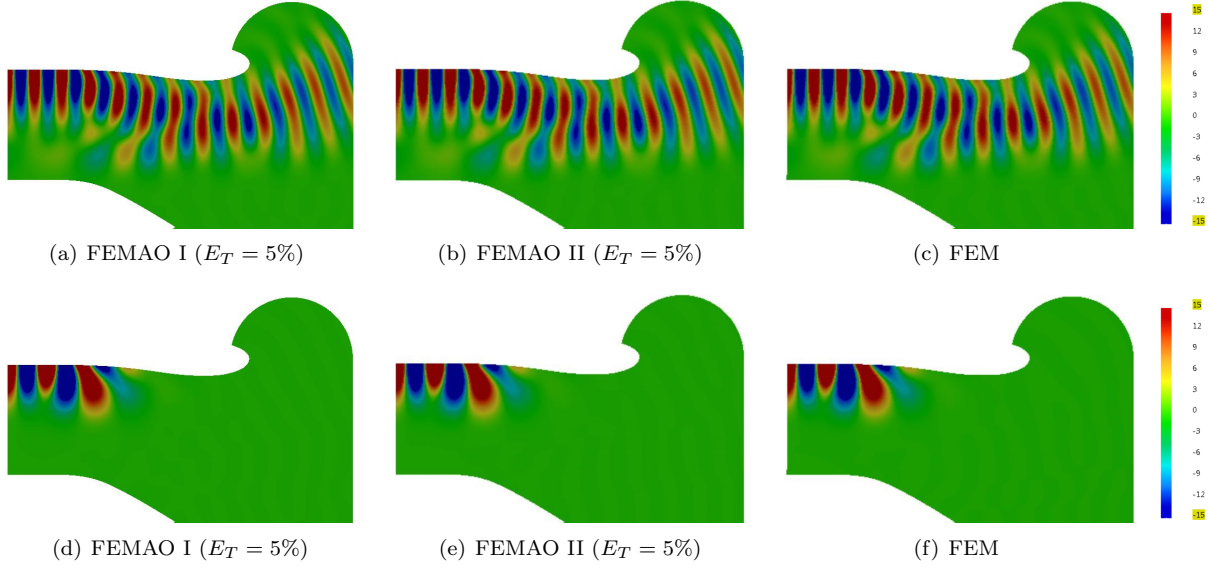


Figure 7. Real part of pressure on a cross-section of the three-dimensional solutions for the approach condition at 2 BPF. Mode (12,1) (top) and mode (24,1) (bottom) with unit power.

The target accuracy for the adaptive scheme is first set to intermediate accuracy ($E_T = 5\%$ as described in section III). Figure 7 shows the pressure fields obtained with the high-order adaptive method using the two coarse meshes alongside the solution from the quadratic FEM solution. The pressure distributions are in excellent agreement for both modes. While the mode (12,1) efficiently propagates to the far field, the mode (24,1) is strongly attenuated by the liner and does not radiate outside the duct.

In order to examine the numerical solutions in more details, Figure 8 presents the acoustic pressure along a control line following the inner wall of the nacelle (see Figure 4 for an illustration of the position of this control line). The pressure distributions between the FEM and the two adaptive solutions with $E_T = 5\%$ exhibit excellent visual agreement for both modes. This is confirmed by computing the relative L^2 -norm of the error along this control line when considering the FEM solution as a reference. The FEMAO I relative error is found to be 1.62% for the mode (12,1) and 1.26% for the mode (24,1). The FEMAO II relative errors are slightly higher in comparison but still close to the target error, with respectively 2.78% and 2.21% for the two incident modes, which could be attributed to the use of a coarser geometric description.

The impact of the liner is clearly visible on these graphs. The incoming pressure field is efficiently attenuated with a decay rate of approximately 15 dB/m for the mode (12,1) and 65 dB/m for the mode (24,1). The effect of the hard-soft wall transitions, indicated by vertical lines on the figure, can also be observed.

The pressure field outside the duct is also examined. Figure 9(a) presents the comparison of the pressure field for the mode (12,1) on the control line #2 positioned along the truncation surface Γ , which is also the control surface for the Kirchhoff integral formulation of the far-field solution (see Figure 4 for an illustration of the position of this control line). The quadratic FEM solution appears to differ from the two high-order models. The main lobe located around 22 degrees is shifted and the pressure amplitude is slightly higher. To confirm that it is indeed the standard FE model that is less accurate, calculations have also been performed with the high-order scheme with a much more stringent accuracy target $E_T = 0.5\%$ for the adaptive model, resulting in higher polynomials orders used in all the elements. When compared to the results obtained with $E_T = 5\%$, see figure 9(b), it is clear that there is very little difference and that the high-order solutions




Figure 8. Real part of pressure along the control line #1 for the approach test case at 2 BPF.

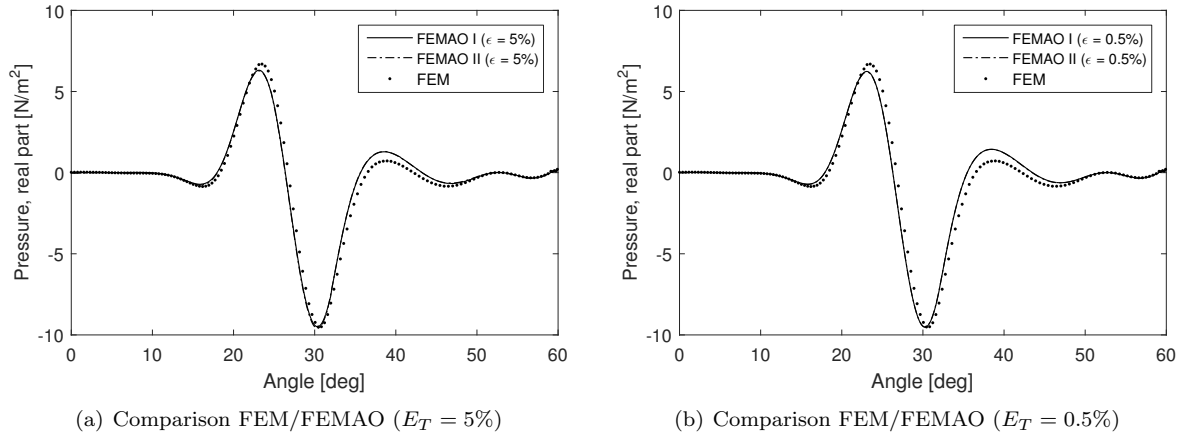


Figure 9. Real part of pressure for the mode (12,1) along the control line #2 for the approach test case at 2 BPF.

have already converged with $E_T = 5\%$. Furthermore, the fact that the results obtained with the two coarse meshes are in excellent agreement also indicates that, at this particular frequency, enforcing a minimum element size of $h = 80$ mm is sufficient to maintain an excellent geometric accuracy. This qualitative analysis is also confirmed by evaluating the relative difference in L^2 -norm along the control line #2, considering the refined FEMAO I solution ($E_T = 0.5\%$) as reference. This yields relative differences of 3.05% and 3.75% for the coarser ($E_T = 5\%$) FEMAO I and II solutions respectively. In comparison, the error on the FEM solution is found to be 18.9%.

The directivity plots in the far field are now examined for the mode (12,1), as shown in Figure 10. The results of the high-order solutions with $E_T = 5\%$ and at $E_T = 0.5\%$ are in excellent agreement, which confirms again that the high-order solution has converged. Differences are observed between the high-order converged solutions and the quadratic FEM solution. In the region of the main directivity lobe, close to 35° , the typical discrepancies are close to 2 dB.

For this type of applications, the numerical models should provide a sufficient dynamic range, that is the numerical noise should be sufficiently small compared to the pressure amplitude of the source. This is important because many of the dominant duct modes are well attenuated by the liner, and the radiated sound field is very small compared to the initial mode amplitude. Yet it is important to capture these properly and for this the dynamic range of the numerical model should be particularly large. The most accurate high-order solution, computed using the mesh FEMAO I with a target accuracy $E_T = 0.5\%$, exhibits a large dynamic range of 50 dB. This solution benefits from the combination of an accurate geometry representation and an efficient high-order polynomial basis. The dynamic range of the solutions computed with $E_T = 5\%$ are between 40 dB and 50 dB which is satisfactory. Figure 10 illustrates how changing the accuracy target




Figure 10. Pressure directivity obtained at 10 m from the intake at the approach condition at 2 BPF for the mode (12,1) with unit power.

	Accuracy	1	2	3	4	5	6	7	8	9	10
FEMAO I	$E_T = 5\%$	0.00	9.65	14.19	15.11	9.21	12.00	14.79	15.92	6.73	2.39
	$E_T = 0.5\%$	0.00	0.00	9.24	11.35	13.73	9.67	9.44	11.22	15.60	19.74
FEMAO II	$E_T = 5\%$	0.00	0.36	2.32	12.10	26.72	21.42	18.55	16.65	1.87	0.00
	$E_T = 0.5\%$	0.00	0.00	0.33	1.40	7.74	18.31	25.05	16.04	18.76	12.37

Table 4. Element order repartition, in percentage of the total volume, for the different high-order models.

E_T allows to control the dynamic range of the numerical model.

As explained above the proposed adaptive approach automatically adjusts the polynomial order across the mesh, based on the target accuracy E_T , the local mean flow properties and the element size. This mechanism can be clearly seen in Table 4, where the orders used for each high-order model are given. The values reported in this table do not refer to the proportion of elements with a given order, but indicate the relative predominance of each order measured in terms of volume (that is the volume occupied by all elements of order p compared to 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.²⁰ With the mesh FEMAO I, the range of orders contributing to the solution is larger than with the mesh FEMAO II. This is expected since there is a wider range of sizes h for the element in the mesh FEMAO I (from $h = 40$ to 300 mm).

It can be noted that setting the more stringent accuracy target $E_T = 0.5\%$ leads to a significant increase in the polynomials orders used in the elements, with a significant portion of the mesh using orders 9 and 10.

V.B. Comparison of performance

The computational cost of the different numerical models is now investigated. The different metrics used to assess the cost are outlined in Table 5. The factorization of the sparse matrix was performed in-core using the MUMPS library.^{23,24}

It is clear that the high-order models outperform the conventional quadratic FEM approach by an order of magnitude. The solution computed from the coarse mesh FEMAO II with a target accuracy of $E_T = 5\%$ is 16 times faster and requires less than a tenth of the memory needed for the standard finite element method. This drastic improvement does not come at the price of a loss in accuracy, since the results in Figures 9 and 10 show that this solution is in fact more accurate than the FEM solution. The benefits obtained from the high-order model based on the mesh FEMAO I are lower but it is still 8 times faster and requires 1/6th of the memory. With this mesh, the model does not fully exploit the accuracy of the high-order approximation basis. This is because a significant portion of the mesh still relies on small elements where a low polynomial order is used.

Model	Accuracy	Number of DOFs*	Factorization time	Factorization memory
FEM		2 996 168	24 min	302 Gb
FEMAO I	$E_T = 0.5\%$	2 219 597	11 min 44 sec	153 Gb
	$E_T = 5\%$	1 012 894	2 min 58 sec	50 Gb
FEMAO II	$E_T = 0.5\%$	985 493	4 min 45 sec	76 Gb
	$E_T = 5\%$	547 276	1 min 30 sec	30 Gb

Table 5. Computational costs of the different finite-element models for the approach condition at 2 BPF. *After static condensation of the internal degrees of freedoms.




Figure 11. High-order acoustic meshes used for the sideline condition, generated from the short and large computational domains defined in Figure 5.

As expected, solutions with the high accuracy target of $E_T = 0.5\%$ are more computationally intensive. However, they still significantly outperform the FEM model, with speed up factors of 2 and 5, and a memory reduction of a half and a quarter, for the meshes FEMAO I and FEMAO II, respectively.

The CPU time for the creation of the mass, stiffness and damping matrices are also monitored. With the mesh FEMAO I this required up to 2 min 58 sec and up to 49 sec with the mesh FEMAO II. In comparison, the FEM quadratic solution required only 23 sec. These results illustrate the fact that evaluating the element matrices represents a larger proportion of the overall computational time for higher-order solutions. But this cost is completely offset by the large reduction in time and memory required for the solution of the linear system. In addition this cost can also be factored out of the frequency loop through the use of hierarchical shape functions, as explained in Section III.C, allowing the matrix evaluation to be performed only once.

V.C. Efficient use of the model

In the previous section, the high-order adaptive approach has been benchmarked against a reference quadratic FEM model for the approach configuration at 2 BPF. In this section, we examine the ability of the high-order method to tackle more challenging configurations at higher frequencies.

For this purpose we first consider the most challenging flow configuration, namely the sideline condition with the aim to reach the second harmonic of the BPF (i.e. a frequency of 2600 Hz). Given the high frequency and the Mach number at the fan face ($M_f = 0.55$), this corresponds to a very challenging Helmholtz number of $kR = 128$. Following the meshing guideline in equation (6), with a target accuracy $E_T = 5\%$, this requires a maximum element size of 0.14 m at the fan plane and $h = 0.24$ m in the far field. A small refinement $h = 0.075$ m is also applied in the region of the strongest geometric curvature, i.e. close to the nacelle lips and at the spinner tip. Finally, a refinement of $h = 0.04$ m is applied at the hard-soft impedance transition, where large gradients are typically observed for the velocity potential.

Figure 11 presents cross-sections of the short and large domain quadratic meshes generated from these requirements. They differ only in the positioning of the truncation surface, as discussed in Section IV.B. It should be emphasized again that, compared to conventional FEM, no particular mesh refinement is required inside the discretization volume. Besides, the resulting meshes contain only 178 834 and 213 707




Figure 12. Real part of pressure on a cross-section of the three-dimensional solutions obtained for the tonal noise for the sideline condition.

nodes respectively, compared to several million nodes for classical FEM, which facilitates the pre- and post-processing of the model.

Solutions are calculated for the rotor-alone tonal noise: namely the first radial modes (24,1) and (48,1), respectively at 1BPF ($kR = 64$) and 2BPF ($kR = 128$). Figure 12 shows the corresponding pressure fields. The results for the large domain and short domain are in close agreement. No noticeable spurious reflections can be seen, confirming the accuracy of the AML method. The mean flow is characterized by a localised region of high speed flow near the throat of the inlet where, for the sideline condition, the Mach number exceeds 0.8 (see Figure 3(c)). The effect of the mean flow in this region on the sound field is well captured by the high-order model. This region of high velocity in the mean flow has a significant impact on the propagation and radiation of sound from the intake, as well as on the sound attenuation by the liner.³³

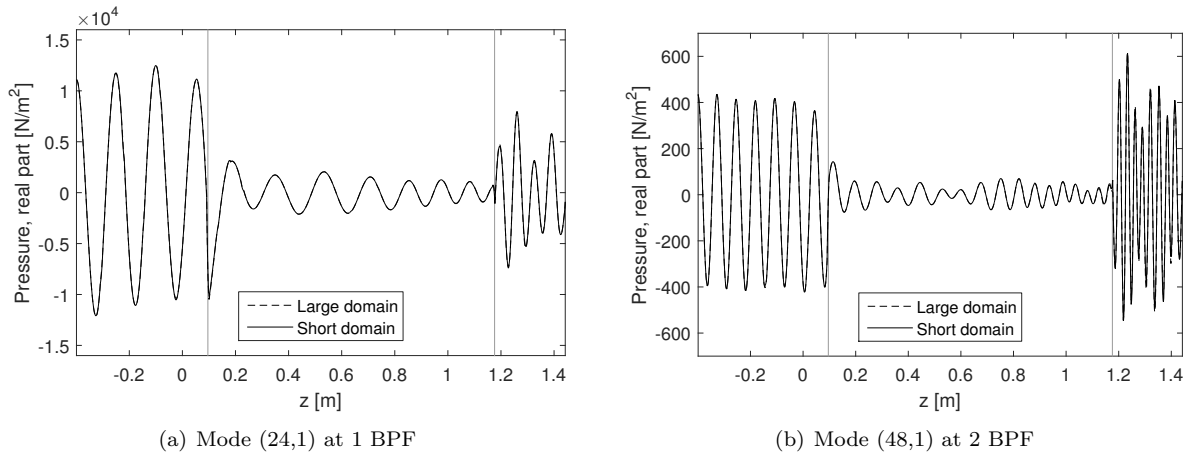


Figure 13. Real part of pressure obtained along the control line #1 for the sideline condition.

In order to better assess the quality of the predictions, the acoustic pressure along the inner wall is




Figure 14. Directivity predicted at 10 m for the sideline condition.

plotted in linear scale in Figure 13. Again, the pressure distributions between the two high-order models with $E_T = 5\%$ are in good agreement: the relative difference in L^2 -norm is found to be 1.95% for the mode (24, 1) and 4.01% for the mode (48, 1) respectively on this graph, taking the large domain solution as reference. The change in wavelength induced by the mean flow is also clearly visible: the wavelength varies from $\lambda \approx 0.08$ m close to the fan, down to 0.025 m near the throat where the flow velocity is high.

Let us now examine the results obtained in the far field. The tones radiate at a similar angle at 1BPF and 2BPF, but the sound radiation pattern of the mode (48,1) is characterized by the presence of two secondary lobes at higher angles instead of one for the mode (24,1), see Figure 14. Larger differences are observed between the large domain and the short domain on these graphs. Typical mismatch in the main lobes is of the order of 1 dB, however, it can reach up to 2.5 dB for the mode (48, 1) at 2 BPF. This difference can be explained by the fact that the mean flow properties computed on the surface Γ of the short domain deviates significantly from the far-field properties, as illustrated in Table 2. For the sideline case, if accurate predictions are required in the far field, the use of the large domain is therefore advocated.

Figure 15 shows the repartition of the element orders in a cross section of the short domain mesh at 1BPF and 2BPF (with $E_T = 5\%$). The high-order adaptive approach automatically adapts the polynomial order based on the frequency, the local element size and the local mean flow properties. At 1BPF with a standard accuracy, the typical order is $p = 5$. In the region close to the hard-soft wall transition, where the mesh is refined, mostly cubic elements are used. At 2BPF, the mesh approaches its maximum frequency limit, the typical order is $p = 9, 10$.

	BPF	kR	Nr of DOFs*	Solving Memory	% in-core	Computational time
Short domain	1	64	2 285 111	182 Gb	100%	17min
	2	128	7 217 171	1410 Gb	18.1%	4h58min
Large domain	1	64	2 423 871	202 Gb	100%	20min
	2	128	7 452 959	1620 Gb	16.9%	5h30min

Table 6. Computational costs for the short and large domain meshes with $E_T = 5\%$ for the sideline configuration.

* After static condensation of the internal degrees of freedoms.

The computational costs for the two different meshes are reported in Table 6. For the large domain mesh, solving at the blade passing frequency ($kR = 64$) with the MUMPS in-core solver required 20 minutes using 202Gb of memory. The corresponding linear system involved over 2.4 million DOFs (after static condensation) and contained over 333 million non-zero entries. This represents a memory requirement of 5.3Gb for the storage of the matrix. The second BPF is much more demanding in comparison. The total size of the linear system, including the bubble shape functions, exceeds 13.5 million DOFs. After static condensation of the internal DOFs, the matrix size is reduced to 7.45 million. The resulting profile contains 2.4 billion terms, hence requiring 39Gb of memory available solely for the matrix storage. The factorization would have required 1620Gb to solve in-core and was solved using the MUMPS out-of-core capabilities




Figure 15. Element order repartition p_e as obtained from the a-priori error estimator on the short domain mesh for the sideline mean flow configuration.

in around 5h30min. Note finally that the element matrix evaluation was performed only once using the multi-frequency algorithm described in III.C in approximately 17min. Expectedly, limiting the size of the domain by moving the truncation surface closer to the nacelle allows reducing the computational cost. The reported timings and memory footprints for the short domain mesh are around 15% lower in comparison. An appropriate trade-off should therefore be found when defining the fictitious truncation surface, in order to obtain accurate far-field predictions without compromising the model performance.

It is also worth commenting on the scalability of the high-order adaptive solver with respect to the frequency. Numerical experiments indicate that the cost scales roughly like ω^3 , which amounts to multiplying the cost by 8 when the frequency is doubled. This is the expected behaviour.²⁰ Note that a standard 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 model size or frequency, but it renders the computational process more efficient and robust.

The same meshes are now used for the cutback configuration using the rotor-alone mode (24,1) at 1BPF and the mode (48,1) at 2 BPF. The pressure field is plotted along the control line #1 on the inner walls of the nacelle in Figure 16 for both modes. The BPF tone modes are strongly attenuated by the liner only a small portion of the initial acoustic power is radiated to the far field. Note that a logarithmic scale was preferred here in order to better examine the decay of the solution along the liner. An attenuation rate of around 120 dB/m is found for both modes. The short and large domain high-order solutions are in good agreement and exhibit a dynamic range superior to 80dB. In terms of computational cost, solving for the large domain at 2BPF ($kh = 84$) involved 3.9 million DOFs and required 543 Gb of memory. It was solved using the MUMPS out-of-core libraries in 1h25min. In comparison, on the same mesh, the first BPF ($kh = 42$) required only 72Gb and was factorized in-core in only 5min 30sec.

Finally, the high-order adaptive solver has been used to predict the tonal noise radiated at the approach condition. The noise source at the fan plane is modelled as an equal energy ensemble of uncorrelated cut-on modes, as explained in Section IV. Results are presented for all the tones from the BPF ($kh = 20$) up to 4BPF ($kh = 80$) in Figure 17 and the computational costs are reported in Table 7.

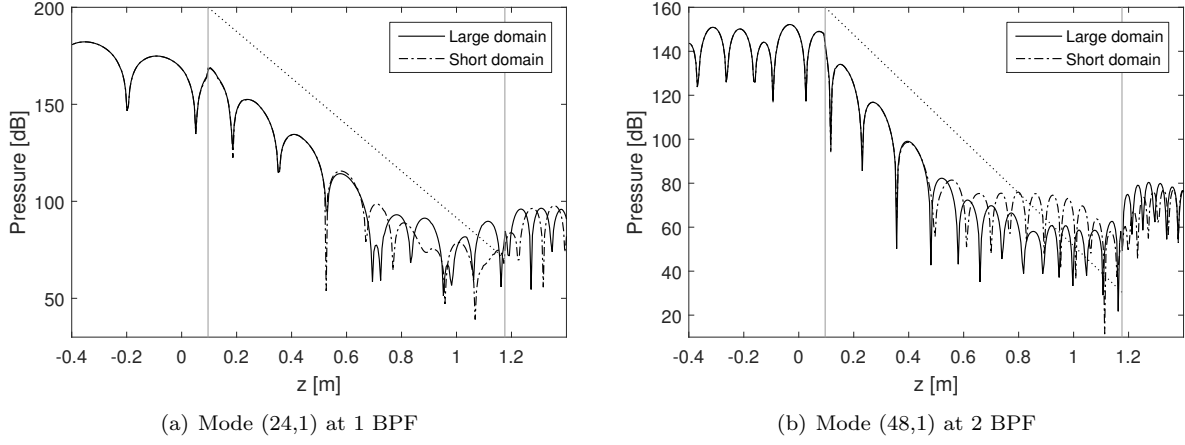


Figure 16. Pressure obtained along the control line #1 of the tonal noise single mode contribution for the cutback test case. Dotted line: 120 dB/m attenuation rate.

	Number of rhs	Number of DOFs*	Factorization Memory	% in-core	Factorization time
1 BPF ($kR = 20$)	68	318 418	8.4Gb	100%	1min30sec
2 BPF ($kR = 40$)	250	1 035 415	52Gb	100%	2min31sec
3 BPF ($kR = 60$)	549	2 029 883	160Gb	16.8%	39min
4 BPF ($kR = 80$)	963	3 322 557	377Gb	16.7%	1h37min

Table 7. Computational costs of the short domain mesh with $E_T = 5\%$ for the multi-mode tonal noise prediction at approach. *After static condensation of the internal degrees of freedoms.

A very large number of modes need to be considered particularly at higher frequencies, with up to 963 cut-on modes at 4BPF. The total computation for the four frequencies took a total of 60 hours, most of the time being spent in the forward-backward substitutions and in the post-processing of each individual modes. Each mode being computed as an independent right-hand-side, the challenge here lies in the way the solver handles a large number of right-hand sides.

Another aspect to consider would be the conditioning of the linear system as this has an impact on the robustness of the direct solver. Unfortunately, for such large matrices, it is extremely resource intensive to accurately calculate the condition number. Instead we point to Reference 21 where a detailed analysis of p -FEM is reported for 2D models, including the link between accuracy and conditioning. It was found that, even for very accurate solutions, the conditioning remained acceptable, see Figure 9 in Ref. 21.

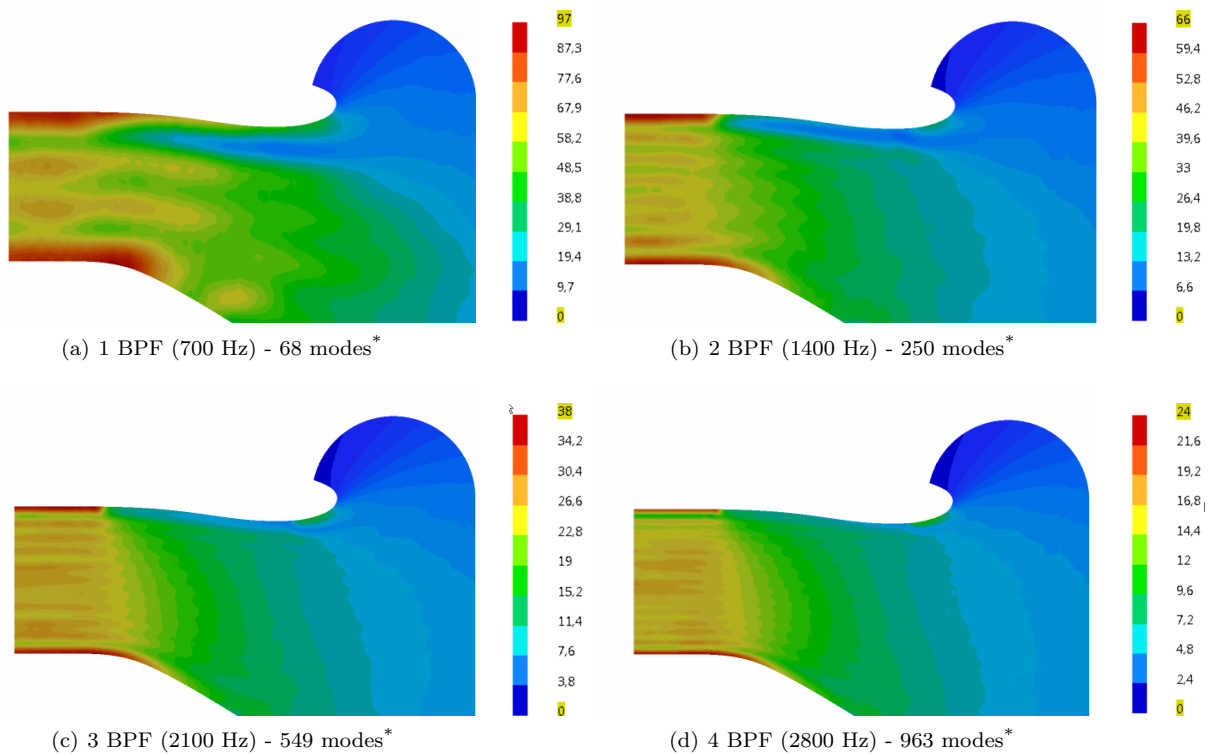


Figure 17. Pressure amplitude obtained for the multi-mode component of the tonal noise for the approach condition.

*Number of cut-on modes included for the multi-mode source.

VI. Conclusions

A high-order, adaptive finite-element method has been described to predict sound propagation in potential base flows. With frequency domain methods, the main limitation is the memory required to solve the large, sparse systems of linear equations resulting from the finite element models. This is addressed in the present method by resorting to high-order shape functions (up to order 10 in this work), instead of the standard quadratic FE models currently in use, so as to reduce of the dispersion error in the numerical model. The benefits in terms of computational costs compared to a traditional quadratic FEM model are very significant, in some cases more than an order of magnitude in memory and runtime. As an indication of the potential of the proposed method for nacelle acoustics, a fully three-dimensional analysis at a Helmholtz number kR of 40 can be solved in a little more than a minute compared to 20 minutes with the standard approach. Analyses were performed for frequencies corresponding to $kR = 128$ within acceptable CPU timings.

Calculating the element matrices when using high-order shape functions represents only a small increase in computational cost compared to the overall computational cost, and it is completely offset by the significant reduction in cost for the linear system solution. Moreover the use of hierarchic shape functions (Lobatto polynomials) allows to render the element matrix calculations negligible. When a number of frequencies are considered, it is possible to calculate the element matrices only once for the highest order and then reuse these for the other frequencies.

A key aspect of the proposed method is the use of an *a priori* adaptive technique to automatically select the order in each element individually. This technique simplifies the preparation of the computational models in several ways. The mesh does not need to be locally refined by the user in the areas with high speed flow (although the geometry still has to be represented accurately, see below). With a uniform mesh, the polynomial order will be automatically increased in regions where the flow velocity is high. When performing several calculations for a number of frequencies, there is no need to prepare several meshes because the adaptive scheme will progressively increase the polynomial order for high frequencies, thus ensuring that a comparable level of accuracy is achieved at each frequency. This changes some of the best practice for the preparation of the models in order to obtain the best performance, and these have been discussed in this paper.

A direction for further work is the description of highly curved elements to allow the use of large elements while still providing an accurate representation of the problem geometry. Preliminary results³⁴ show there is potential in using curved elements with p -FEM but one of the issues remaining is the untangling of these curved elements.

Acknowledgements

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

References

- ¹Astley, R., “Numerical methods for noise propagation in moving flows, with application to turbofan engines,” *Acoustical Science and Technology*, Vol. 30, No. 4, 2009, pp. 227–239.
- ²Mustafi, P., *Improved turbofan intake liner design and optimization*, Ph.D. thesis, University of Southampton, 2013.
- ³Prinn, A., *Efficient finite element methods for aircraft engine noise prediction*, Ph.D. thesis, University of Southampton, 2014.
- ⁴Williamschen, M., Gabard, G., and Bériot, H., “Performance of the DGM for the Linearized Euler Equations With Non-Uniform Mean-Flow,” *21st AIAA/CEAS Aeroacoustics Conference*, 2015, AIAA Paper 2015-3277.
- ⁵Gabard, G. and Brambley, E., “A Full Discrete Dispersion Analysis of Time-Domain Simulations of Acoustic Liners with Flow,” *Journal of Computational Physics*, Vol. 273, 2014, pp. 310–326.
- ⁶Brambley, E. J. and Gabard, G., “Time-domain implementation of an impedance boundary condition with boundary layer correction,” *Journal of Computational Physics*, Vol. 321, 2016, pp. 755–775.
- ⁷Rao, P. and Morris, P., “Use of finite element methods in frequency domain aeroacoustics,” *AIAA Journal*, Vol. 44, No. 7, 2006, pp. 1643–1652.
- ⁸Myers, M., “On the acoustic boundary condition in the presence of flow,” *Journal of Sound and Vibration*, Vol. 71, No. 3, 1980, pp. 429–434.
- ⁹Bermúdez, A., Hervella-Nieto, L., Prieto, A., et al., “An optimal perfectly matched layer with unbounded absorbing function for time-harmonic acoustic scattering problems,” *Journal of Computational Physics*, Vol. 223, No. 2, 2007, pp. 469–488.

- ¹⁰Bécache, E., Dhia, A. B.-B., and Legendre, G., “Perfectly matched layers for the convected Helmholtz equation,” *SIAM Journal on Numerical Analysis*, Vol. 42, No. 1, 2004, pp. 409–433.
- ¹¹Casalino, D. and Genito, M., “Achievements in the numerical modeling of fan noise radiation from aero-engines,” *Aerospace Science and Technology*, Vol. 12, No. 1, 2008, pp. 105 – 113, Aircraft noise reduction.
- ¹²Hamiche, K., Gabard, G., and Bériot, H., “A High-Order Finite Element Method for the Linearised Euler Equations,” *Acta Acustica united with Acustica*, Vol. 102, No. 5, 2016, pp. 813–823.
- ¹³Martínez-Lera, P., Schram, C., Bériot, H., and Hallez, R., “An approach to aerodynamic sound prediction based on incompressible-flow pressure,” *Journal of Sound and Vibration*, Vol. 333, No. 1, 2014, pp. 132–143.
- ¹⁴Martínez-Lera, P., Kucukcoskun, K., Shur, M., Travin, A., and Tournour, M., “Hybrid aeroacoustic computations for flows in ducts with single and tandem diaphragms,” *22nd AIAA/CEAS Aeroacoustics Conference*, 2016.
- ¹⁵Rienstra, S. and Hirschberg, A., *An introduction to acoustics*, Self published, 2014, Available at <http://www.win.tue.nl/~sjoerdr/papers/boek.pdf>.
- ¹⁶Eversman, W., “The boundary condition at an impedance wall in a non-uniform duct with potential mean flow,” *Journal of Sound and Vibration*, Vol. 246, No. 1, 2001, pp. 63–69.
- ¹⁷Šolín, P., Segeth, K., and Doležal, I., *Higher-order finite element methods*, Chapman & Hall, New York, 2003.
- ¹⁸Karniadakis, G. and Sherwin, S., *Spectral/hp element methods for computational fluid dynamics*, Oxford University Press, 2013.
- ¹⁹Vos, P. E. J., Sherwin, S. J., and Kirby, R. M., “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*, Vol. 229, No. 13, 2010, pp. 5161–5181.
- ²⁰Bériot, H., Prinn, A., and Gabard, G., “Efficient implementation of high-order finite elements for Helmholtz problems,” *International Journal for Numerical Methods in Engineering*, Vol. 106, No. 3, 2016, pp. 213–240, nme.5172.
- ²¹Lieu, A., Gabard, G., and Bériot, H., “A comparison of high-order polynomial and wave-based methods for Helmholtz problems,” *Journal of Computational Physics*, Vol. 321, 2016, pp. 105–125.
- ²²Bériot, H., Gabard, G., and Perrey-Debain, E., “Analysis of high-order finite elements for convected wave propagation,” *International journal for numerical methods in engineering*, Vol. 96, No. 11, 2013, pp. 665–688.
- ²³Amestoy, P. R., Duff, I. S., Koster, J., and L’Excellent, J.-Y., “A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling,” *SIAM Journal on Matrix Analysis and Applications*, Vol. 23, No. 1, 2001, pp. 15–41.
- ²⁴Amestoy, P. R., Guermouche, A., L’Excellent, J.-Y., and Pralet, S., “Hybrid scheduling for the parallel solution of linear systems,” *Parallel Computing*, Vol. 32, No. 2, 2006, pp. 136–156.
- ²⁵Gabard, G., Astley, R., and Ben Tahar, M., “Stability and accuracy of finite element methods for flow acoustics. II: Two-dimensional effects,” *International Journal for Numerical Methods in Engineering*, Vol. 63, 2005, pp. 974–987.
- ²⁶Johnen, A., Remacle, J.-F., and Geuzaine, C., “Geometrical validity of curvilinear finite elements,” *Journal of Computational Physics*, Vol. 233, 2013, pp. 359–372.
- ²⁷Ziel, V., Bériot, H., Atak, O., and Gabard, G., “Comparison of 2D boundary curving methods with modal shape functions and a piecewise linear target mesh,” *Procedia Engineering*, Vol. 203, 2017, pp. 91–101.
- ²⁸Sevilla, R., Fernández-Méndez, S., and Huerta, A., “NURBS-enhanced finite element method (NEFEM),” *International Journal for Numerical Methods in Engineering*, Vol. 76, No. 1, 2008, pp. 56–83.
- ²⁹Khajah, T., Antoine, X., and Bordas, S., “Isogeometric finite element analysis of time-harmonic exterior acoustic scattering problems,” *arXiv preprint arXiv:1610.01694*, 2016.
- ³⁰Motsinger, R. and Kraft, R., “Design and performance of duct acoustic treatment,” *Aeroacoustics of Flight Vehicles: Theory and Practice. Volume 2: Noise Control*, edited by H. Hubbard, 1991.
- ³¹LMS, *Virtual Lab Acoustic R13 user’s manual*, LMS, Leuven, Belgium, 2014.
- ³²B. Van Antwerpen, R. Leneveu, S. C., “New advances in the use of Actran/TM for nacelle simulations,” *AIAA 14th Aeroacoustics Conference*, Vancouver, British Columbia, Canada, 2008.
- ³³Astley, R., Hamilton, J., Baker, N., and Kitchen, E., “Modelling tone propagation from turbofan inlets,” *Proceedings of the 8th AIAA/CEAS Aeroacoustics Conference, Breckenridge, Colorado*, 2002, AIAA paper 2002-2449.
- ³⁴Ziel, V., Bériot, H., Atak, O., and Gabard, G., “High-order 2D mesh curving methods with a piecewise linear target,” *Computer-Aided Design*, 2018, Submitted.