

Convected Acoustic Physics Interfaces for Acoustic Resonance Technology

Elisabetta Merico¹, Antonio Jimenez-Garcia²

1. Multiphysics Engineer, FT Technologies Ltd, Sunbury House, Sunbury-on-Thames, TW16 7DX, UK

2. Principal Research Engineer, FT Technologies Ltd, Sunbury House, Sunbury-on-Thames, TW16 7DX, UK

Abstract

Acoustic Resonance (Acu-Res®) Technology serves as the foundational principle behind FT wind sensors, enabling the measurement of fluid flow by detecting phase shifts between paired transducers within an acoustic resonator. This technology inherently couples acoustics with aerothermal fields, making precise modelling essential for optimal sensor performance. This paper presents a comprehensive modelling framework for convected acoustics using COMSOL Multiphysics®, focusing on Acoustic Resonance Technology applications. Two physics interfaces are evaluated to determine the best compromise between computational accuracy and efficiency. The study systematically explores mesh density, computational domain boundaries, solver types, and computational power, aiming to enhance the design lifecycle of Acu-Res® Technology through state-of-the-art simulation strategies and SciML-based methodologies.

Keywords: convected flow acoustics, acoustic resonance technology, finite elements method.

Introduction

As FT Technologies continues to advance Acu-Res® Technology [1], simulation tools that deliver both accuracy and speed are becoming crucial to keep up to speed with the latest industry standards.

This paper goes through multiple aspects of convected acoustic modelling for Acu-Res® Technology applications. It begins with a review of two available convected acoustic physics formulations in COMSOL Multiphysics® [2], outlining their theoretical foundations and practical implications. The next section investigates the impact of mesh density and computational domain setup, including the use of Perfectly Matched Layers (PML) and Acoustic Impedance boundary conditions, on solution accuracy and computational performance.

The following section examines the effects of solver selection, comparing direct and iterative approaches, on CPU time, emphasising the trade-offs encountered in computationally intensive simulations. The influence of various background flow conditions is then assessed, focusing particularly on the limitations of the Linearised Potential Flow (LPF) model and benchmarking its performance against more advanced formulations, such as the Linearised Navier-Stokes (LNS) model [3].

By providing a detailed evaluation across these sections, the study aims to offer clear guidance on how to optimise simulation strategies for industrial applications. Through this structured approach, the paper not only addresses the technical requirements of accurate and efficient computational modelling but also underscores its significance for accelerating design cycles and fostering innovation within the industry. This framework will serve as the foundation for future work integrating SciML-based techniques, ensuring that Acu-Res® Technology continues to evolve in alignment with industry needs.

Governing Equations

The study of convected flow acoustics aims to understand how background flow interacts with acoustic waves. While fully coupled simulations using compressible Navier-Stokes equations offer the most comprehensive insight, they are often computationally prohibitive for practical engineering problems. As a result, a two-step approach is generally preferred: first simulating the flow field, then analysing acoustic perturbations within that flow [4]. In this study, we concentrate on two key modelling strategies available in COMSOL Multiphysics®: the Linearised Potential Flow (LPF) and Linearised Navier-Stokes (LNS) models, which provide efficient and effective means for capturing the relevant aeroacoustic phenomena in Acu-Res® Technology applications. The first model is suitable for all the applications with a background flow well described by a compressible potential formulation, which translates into an inviscid, barotropic, and irrotational flow [5]. The governing equation in the frequency domain is:

$$-\frac{\rho_0}{c_0^2} i\omega(i\omega\phi + \mathbf{u}_0 \cdot \nabla\phi) + \nabla$$

$$\cdot \left(\rho_0 \nabla\phi - \frac{\rho_0}{c_0^2} (i\omega\phi + \mathbf{u}_0 \cdot \nabla\phi) \mathbf{u}_0 \right) = 0$$

with ϕ being the velocity potential, ρ the fluid density, \mathbf{u}_0 the mean background velocity and c_0 the speed of sound.

On the other hand, Linearised Navier-Stokes (LNS) formulation solves for the continuity, momentum, and energy equations, which allow accounting for compressible, viscous, and non-isothermal background flow [6].

Numerical Simulation Setup

COMSOL Multiphysics® was utilised to configure the case by importing the geometry from a SolidWorks CAD model representing a simplified wind sensor, as shown in Figure 1.




Figure 1: Wind sensor's simplified geometry

The computational domain was then constructed as a block encompassing the resonator cavity. Its dimensions were defined according to the results of the study.




Figure 2: Computational domain encompassing half the resonator area and extending $rbc \cdot \lambda_{max}$ radially and $rbc_y \cdot \lambda_{max}$ vertically

As shown in the figure above, the computational domain dimension is determined by two variables, rbc (radial boundary condition) and rbc_y , which specify the extent of radial and vertical expansion beyond the resonator, respectively. Both coefficients are multiplied by the maximum wavelength under consideration, which corresponds to the minimum frequency that must be resolved. When implementing the PML as a boundary condition, an additional layer is incorporated into the computational domain, thereby increasing its overall dimensions.

After defining the dimensions of the computational domain, it has to be meshed. The maximum mesh element dimension was parametrised using the mesh density, defined as the number of elements per wavelength, with the maximum mesh element dimension being $md \cdot \lambda_{min}$, with md being the mesh density and λ_{min} the smallest wavelength corresponding to the highest frequency considered.

A distribution comprising six elements was established for the Perfectly Matched Layer (PML) mesh. Second-order Lagrange discretisation was utilised consistently across all cases. These two

parameters were held constant to restrict the scope of variable combinations under consideration.

Moreover, the use of a symmetry plane at the cavity's middle point was evaluated to take advantage of its computational efficiency while ensuring solution accuracy. The geometries of FT wind sensors exhibit a symmetry plane that bisects the cavity, a characteristic that can be leveraged in numerical simulations to significantly decrease computational time. Nevertheless, this assumption requires validation.

All these first steps of the process were run using the default direct solver for LPF simulations in the frequency domain. Only after determining the computational domain size and the mesh density, iterative solvers were considered, comparing GMRES with GMG and FGMRES with GMG in terms of computational time and solution accuracy.

Finally, a more powerful machine has been used to compare the best case found, in order to assess the improvements brought by an AMD Ryzen Threadripper PRO 7975WX 32-Cores @ 4GHz, 256GB memory, compared to an Intel® Core™ i9-9900K CPU @ 3.60GHz, 128GB memory.

Two physics modules were evaluated: Linearised Potential Flow (LPF) and Linearised Navier-Stokes (LNS). Both were applied across the computational domain with a uniform background flow. The only material considered was the air, as the walls of the wind sensor were modelled as sound-hard boundaries, and the acoustic source, a transducer, was simplified as a uniformly moving piston.

The test case considered to assess all the parameters was a study in the frequency domain, with a uniform flow of 75m/s passing through the cavity parallel to the symmetry plane.

Simulation Results

As outlined in the previous sections, this study followed multiple steps to define the computational domain, mesh, and the boundary conditions that would provide the best compromise between computational cost and accuracy, using the LPF model in COMSOL. After establishing these parameters, both direct and iterative solvers within COMSOL Multiphysics® were assessed to determine whether computational time could be reduced further without compromising the accuracy of the results. Subsequently, the simulation was executed on a higher-performance machine to evaluate the benefits of increased CPU cores and memory.

Finally, the focus is shifted to the LNS model to assess the improvements that it brings to the acoustic solution, studying also the effects of the mesh density on the accuracy of the solution.

The accuracy of the simulations will be assessed by analysing the magnitude and phase of acoustic

pressure recorded on the transducers acting as receivers.

This method aligns with how FT wind sensors detect wind speed and direction, making it a Key Performance Indicator (KPI) for evaluating simulation reliability.

Effect of Computational Domain on LPF

A summary of the cases evaluated to define the computational domain and the mesh density for the LPF physics interface is presented below:

caseID	Sym	rbc	rbc _y	elem/λ	BC
00-03	Yes	0.5,1,2,3	0.5	6	IMP
04-09	Yes	0.5	0.5	4,5,7,8,9,10	IMP
10	Yes	0.5	1	7	IMP
11	No	0.5	0.5	7	IMP
12-15	Yes	0.5,1,2,3	0.5	6	PML
16-21	Yes	0.5	0.5	4,5,7,8,9,10	PML
22	Yes	0.5	1	7	PML
23	No	0.5	0.5	7	PML

Table 1: List of cases tested to define the computational domain and mesh density

The first parameter studied was the radial boundary condition, with *rbc* varying from 0.5 to 3, both when using acoustic impedance (caseID: 00-03) and PML (caseID: 12-15) boundary conditions. To assess its effects, the plot in Figure 3 presents a comparison of the acoustic pressure magnitude measured on the receiving transducer, varying the *rbc*. The pressure magnitude is normalised to compare with the zero-flow case, where the maximum magnitude is 1. For clarity of the plot, only the case with an acoustic impedance boundary condition is presented, but similar results are obtained when PML is applied.

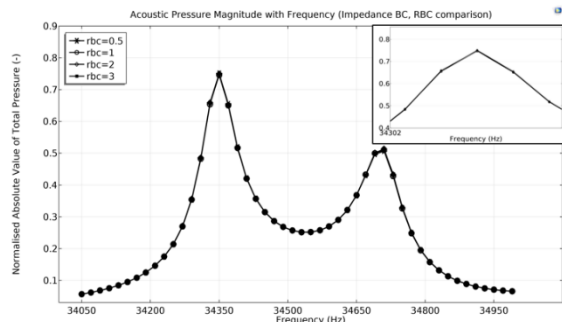


Figure 3: Effect of the radial dimension of the computational domain on the acoustic pressure magnitude on the receiver transducer

Figure 3 illustrates the variation in acoustic magnitude recorded with increasing computational

domain size in the radial direction. The figure demonstrates that expanding the domain in the radial direction results in minimal to no variation in acoustic magnitude. A similar observation is made when applying PML as a boundary condition, indicating that a radial dimension of 0.5λ is sufficient to prevent reflections from the boundaries, irrespective of the boundary condition used.

The plot in Figure 4 provides a comparison of computational times, highlighting the increase associated with a larger computational domain. As previously noted, the FT wind sensors determine speed and direction based on phase differences, making this metric the primary key performance indicator for evaluating both the accuracy of the solutions and their alignment with real-world scenarios. Hence, in the next figure, the relationship of the phase error with the domain size is presented alongside the increase in computational time. The phase error is calculated in relative terms by comparing it with the simulation run that shows the highest accuracy for each case. This approach aims to identify the point at which the solution does not depend on the domain or mesh density, as indicated by a minimal difference between cases with greater accuracy.

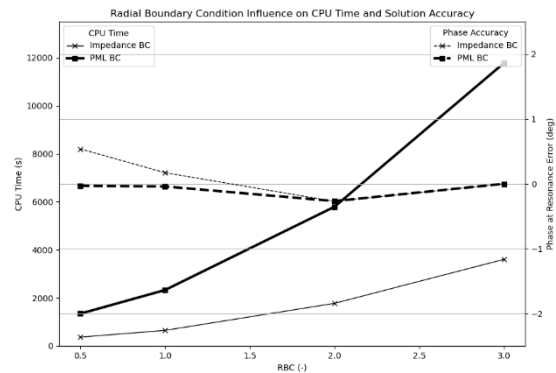


Figure 4: Computational time and error in the phase at resonance, when using impedance and PML BCs as a function of the rbc

The data presented in Figure 4 confirm that increasing the radial domain size does not offer significant benefits in solution accuracy, while it substantially increases computational time. Consequently, the decision was made to proceed with *rbc* = 0.5, as it brings computational efficiency without compromising the accuracy of the results.

On the other hand, the specification of mesh element size proved to be a critical factor influencing both the magnitude and phase of the received pressure signal. For both impedance (caseID: 04-09) and PML (caseID: 16-21) boundary conditions, mesh densities between 4 and 10 elements per wavelength (λ) were evaluated.

Also for this quantity, the initial focus is on the effect on magnitude by examining the pressure response across frequencies for simulations conducted with varying mesh densities. As previously, only cases using impedance boundary conditions are included to maintain clarity in the plot.

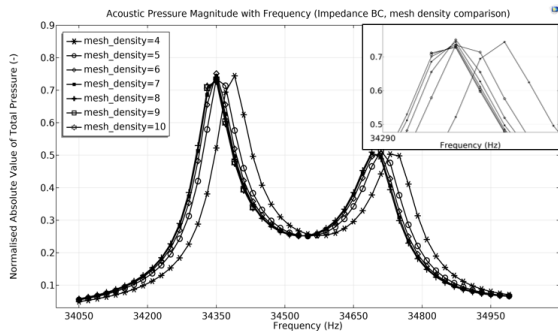


Figure 5: Effect of the mesh density on the acoustic pressure magnitude on the receiver transducer

Figure 5 demonstrates that the mesh density significantly influences the acoustic pressure signal within the resonator cavity, particularly affecting the readings of the receiver transducer. Inadequate selection of mesh density can result in inaccurate outcomes, thereby compromising the reliability of the computed solution. Focusing on the accuracy of the phase and computational time, a similar observation emerges: mesh density plays a critical role in ensuring the reliability of the computed solution, but it must be balanced with the computational time required.

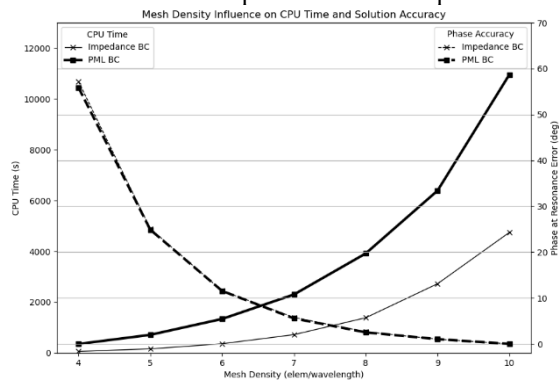


Figure 6: CPU time and phase error as a function of the mesh density

The plot above reaffirms the findings presented in Figure 5, demonstrating the significant impact of mesh density on accuracy outcomes. As expected, simulation time increases considerably with the number of degrees of freedom, which are directly associated with mesh density. Consequently, it is necessary to balance the acceptable error margin with the available computational resources. In this instance, a compromise was achieved by selecting 7 elements per wavelength. This resulted in an error reduction to 8.5% with a computation time of 12 minutes for the impedance boundary condition, and

an error of 8% with a computational time of 38 minutes to obtain the solution using PML.

After the mesh density and radial dimension of the domain were set, the vertical extension was evaluated by varying rb_{c_y} from 0.5 to 1, using both impedance (caseID:10) and PML (caseID: 22) boundary conditions. This adjustment resulted in minimal changes in the solution, with less than 1% difference in magnitude and phase at resonance for both PML and Impedance BCs, while computational time nearly doubled. In addition to enabling shorter simulation times, this approach provides valuable insight into the propagation of the acoustic field beyond the resonant cavity. The results indicate that the majority of energy remains confined within the resonator area, so even a relatively limited computational domain is sufficient to prevent reflections from the boundaries.

The impact of symmetry was analysed by comparing simulations of the full cavity with those using only half of the domain. The study found that, regardless of whether acoustic impedance (caseID: 11) or PML (caseID: 23) boundary conditions were applied, the symmetry assumption did not influence the results. However, the computational time more than doubled when simulating the full cavity. These findings confirm that modelling the entire cavity is unnecessary for achieving accurate outcomes, thereby reducing computational costs and offering important insight into the acoustic symmetry present in the cavity.

After evaluating the computational domain and mesh effects for both boundary conditions, the optimal case for each will be compared to assess whether the computational time savings from using Acoustic Impedance as a boundary condition impact solution accuracy. Figure 7 presents a comparison of acoustic pressure magnitudes.




Figure 7: Comparison of the pressure magnitude obtained using impedance (caseID:06) and PML (caseID:18) BCs

Figure 7 shows minimal differences between PML and impedance boundary conditions. Consequently, it is appropriate to employ acoustic impedance as the boundary condition, as it offers a significant computational advantage, reducing the computation time by more than threefold.

To conclude, this comparison indicated that using acoustic impedance as boundary conditions, a domain extending half a wavelength in all directions, and a mesh density of 7 elements per wavelength, while modelling only half of the resonating cavity to reduce computational time, resulted in a total computational time of 12 minutes per frequency sweep, making it the best compromise between cost and accuracy.




Figure 8: Final computational domain and mesh setup

Effect of Solver

Upon completing the domain and mesh definitions, achieved through a balance of accuracy and computational cost, various solver options were assessed. Up to this point, all investigations have utilised the default direct solver without modification. According to reference [7], when employing MUMPS, disabling low-rank factorisation can yield speed improvements of up to 25%. Furthermore, as the number of degrees of freedom increases, iterative solvers may significantly reduce both computational time and memory requirements. For LPF models operating in the frequency domain, the two standard iterative solvers are *GMRES with GMG* and *FGMRES with GMG*. Both leverage a geometric multigrid preconditioner performing two iterations in each V-cycle. Using these iterative solvers on the best case identified earlier (caseID: 06) significantly reduced computational time with minimal changes to the solution.

Solver	CPU Time
Direct Solver (MUMPS)	12min 1s
Direct Solver (MUMPS with low-rank factorisation disabled)	10min 50s
GMRES with GMG	2min 57s
FGMRES with GMG	2min 40s

Table 2: Computational time for different solvers

As indicated in the table above, disabling the low rank factorisation resulted in a 10% reduction in computational time. However, the implementation of iterative methods significantly enhanced the efficiency of the simulations without sacrificing accuracy. Specifically, the GMRES and FGMRES methods reduced computational time by factors of 4 and 4.5, respectively, with discrepancies in the computed acoustic field below 0.015% in both cases. The acoustic magnitude and phase at the receiving transducer were nearly identical across methods; thus, errors were evaluated by comparing the entire acoustic field within the domain.

Additional improvements were observed when the simulation was conducted on a more powerful machine, as specified in the previous section. Only the best-case scenario was evaluated, resulting in a reduction of computational time from 160 seconds to 90 seconds.

Validation against Experimental Data

Following the model setup, a comparison was performed to validate the simulation results against experimental data. The primary KPI for this analysis is the predicted resonance frequency across various speeds. The resonance frequency of the cavity is affected by gap size, speed of sound (which relates to temperature), and flow velocity [8]. To assess this, a test using LPF in COMSOL simulated speeds from 0 to 75 m/s to determine whether changes in the resonance frequency predicted by COMSOL matched measurements obtained by averaging the outputs of multiple sensors from the FT wind tunnel.

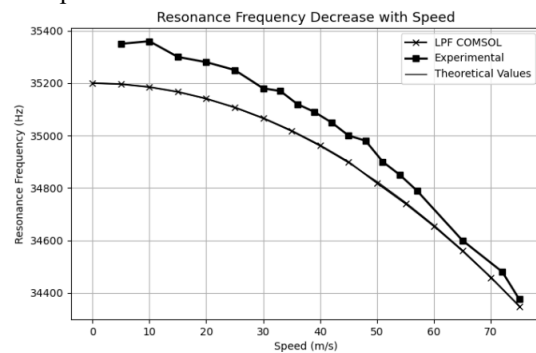


Figure 9: Variation of resonance frequency with speed

A second and really important KPI that was used to determine the accuracy of the LPF simulations was the speed that could be computed considering the phase differences among the axes, exactly in the same way as in the real sensor. To achieve that, the full domain was simulated, and three simulations were run for each speed, exciting one transducer at a time. This way, the wind speed could be resolved and compared with the imposed speed for the uniform background flow:

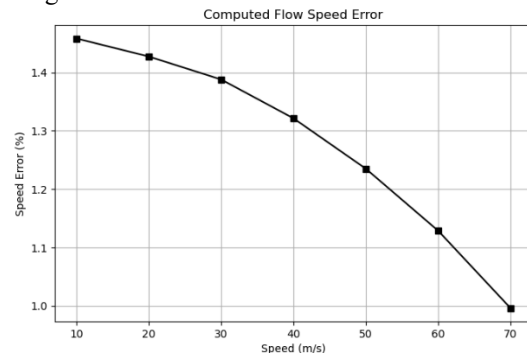


Figure 10: Percentage error in speed from LPF COMSOL simulations computed using phase differences, as a function of the prescribed background flow

Figure 9 and Figure 10 illustrate that the LPF model within COMSOL reliably predicts resonance frequency and speed in the cavity, demonstrating strong agreement with experimental results and perfect overlap with theoretical values, confirming that it is robust and reliable enough to be used for iterative cycles. Additionally, the LPF model supports rapid simulations, typically completed in approximately ninety seconds, which makes it well-suited for the development of Design of Experiments (DoE) as well as surrogate models, incorporating the effects of variable speed and geometric parameters. This capability offers significant advantages for efficient prototyping.

Comparison of LPF and LNS

The final phase of this study involved comparing two physics interfaces for convected acoustics available in COMSOL Multiphysics®. Employing the Linearised Navier-Stokes model facilitates the incorporation of multiple physical effects and enables a more accurate representation of fluid flow. However, this approach significantly increases computational time. For analysis using the LNS model, the same mesh density evaluation procedure was applied, with domain dimensions held constant, as prior results indicated it had minimal impact on the solution. Due to the substantial computational demands, even with relatively low mesh densities and using iterative solvers, a standard PC proved insufficient; consequently, all simulations were performed on the previously specified high-performance machine.

The mesh densities considered ranged from 4 to 13 elements per minimum wavelength. In all the cases considered, quadratic Lagrange discretisation was used for the pressure, and linear for both velocity and temperature. However, the adiabatic condition was enabled, which reduces the number of degrees of freedom, deriving the temperature from pressure and density.




Figure 11: Effect of the mesh density on the acoustic pressure magnitude measured at the receiving transducer

The table below indicates that LNS simulations require more computational time compared to what was observed for LPF simulations. Indeed, the LPF

physics interface resolves only for velocity potential, while LNS computes pressure, all three components of velocity, and temperature. As a result, the degrees of freedom, and consequently the computational time and memory requirements, are increased in LNS simulations.

Mesh Density	DoF	CPU Time
4	72k	1min 21s
6	231k	5min 56s
7	361k	15min 9s
8	539k	33min 40s
9	763k	1hr 8min
10	1047k	2hr 23min
12	1810k	5hr 50min
13	2295k	9hr 10min

Table 3: Correlation between mesh density, Degrees of Freedom (DoF), and CPU time for LNS simulations

Furthermore, Figure 11 shows that a higher mesh density is required in the case of LNS to achieve good accuracy of the results. Looking into more detail at the error in the magnitude and phase compared to the mesh density:

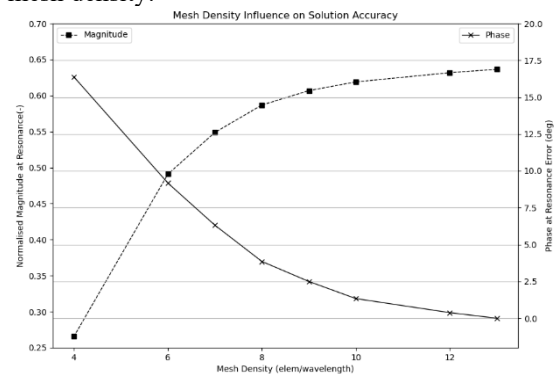


Figure 12: Mesh density effect on solution accuracy in terms of acoustic pressure magnitude and phase values at resonance

For the LPF model, 7 elements per wavelength proved to be a good compromise between accuracy and computational cost. On the other hand, for the LNS physics interface, 8 elements per wavelength were selected, as this configuration enables tracking of the correct frequency for the peak magnitude with minimal error in the actual magnitude value and maintains phase measurement uncertainty below 4°. Utilizing a mesh with at least nine elements per wavelength would have further enhanced the solution, reducing the phase error to 2.5° and achieving a magnitude equal to 96% of that calculated with the maximum tested resolution (13 elements per wavelength). However, these improvements in accuracy would not offset the doubled computational time required by this resolution. Consequently, a resolution of eight elements per wavelength was deemed sufficient for the required level of accuracy.

Focusing on the comparison of the results obtained using the two different physics interfaces:




Figure 13: Normalised pressure magnitude obtained using LPF and LNS physics interfaces

Figure 13 indicates that the LNS physics interface incorporates additional dissipation mechanisms, evidenced by the reduction of acoustic magnitude compared to the LPF solution. Nonetheless, the predicted resonance frequency remains unchanged, and the correspondence between the two local maxima exhibits a high degree of similarity. In conclusion, the LNS physics interface serves as a valuable resource for obtaining further insights into cavity response, while the LPF model delivers sufficiently accurate solutions suitable for rapid iteration cycles.

Conclusions

Optimising both the computational domain and mesh density has proven to be pivotal in achieving an efficient balance between accuracy and computational resource expenditure in the numerical simulations presented. The results clearly demonstrate that carefully selecting mesh parameters enables substantial reductions in CPU time without compromising the fidelity of the solutions. Specifically, for the LPF physics interface, a mesh density of 7 elements per wavelength with a second-order scheme was found to offer an excellent compromise, yielding accurate results at a significantly lower computational cost compared to higher resolutions.

In the case of the LNS simulations, which inherently demand more computational effort due to the inclusion of pressure and all three velocity components, optimising the mesh density to 8 elements per wavelength allowed the capture of resonance frequency and magnitude with minimal error, while maintaining the phase uncertainty below 4°. While further refinement to 9 elements per wavelength could marginally improve accuracy, the associated doubling of computational time was not justified by the modest gains.

A brief comparison of the LPF and LNS methodologies highlights the trade-off between computational efficiency and physical detail. The

LPF approach, by resolving only the velocity potential, achieves substantial savings in computational resources, making it suitable for scenarios where phase and magnitude accuracy can be achieved with lower degrees of freedom. In contrast, the LNS methodology, though more demanding, provides a more comprehensive representation of the acoustics by accounting for additional physical phenomena. Nevertheless, targeted optimisation of the mesh and computational domain ensures that the additional computational cost of LNS can be effectively managed.

In summary, the strategic optimisation undertaken in this study underscores the importance of tailoring simulation parameters to the chosen modelling approach and desired accuracy, enabling robust and efficient acoustic analysis across both LPF and LNS frameworks.

References

- [1] “FT Technologies,” [Online]. Available: <https://fttechnologies.com/technology>. [Accessed 8 September 2025].
- [2] “COMSOL Multiphysics®,” [Online]. Available: <https://www.comsol.com/>. [Accessed 9 September 2025].
- [3] COMSOL Multiphysics®, “Acoustics Module User's Guide,” Stockholm, Sweden, COMSOL AB, 2022, pp. 347-442.
- [4] “Theory Background for the Aeroacoustics Branch,” [Online]. Available: https://doc.comsol.com/5.5/doc/com.comsol.h elp.aco/aco_ug_aero.08.100.html#570134. [Accessed 5 September 2025].
- [5] "Linearized Potential Flow," COMSOL, [Online]. Available: https://doc.comsol.com/5.5/doc/com.comsol.h elp.aco/aco_ug_aero.08.105.html. [Accessed 4 September 2025].
- [6] "Linearized Navier-Stokes Model," COMSOL, [Online]. Available: https://doc.comsol.com/5.5/doc/com.comsol.h elp.aco/aco_ug_aero.08.026.html. [Accessed 3 September 2025].
- [7] “Solving Large Acoustics Problems Using Iterative Solvers,” COMSOL, [Online]. Available: https://doc.comsol.com/5.6/doc/com.comsol.h elp.aco/aco_ug_pressure.05.109.html. [Accessed 5 September 2025].
- [8] S. Kapartis, “Ultrasonic flow velocity sensor and method of measuring the velocity of a fluid flow”. Europe Patent 97202423.5, 9 April 1997.