

UNIVERSITY OF SOUTHAMPTON

High-accuracy methods for frequency-domain flow acoustics

by

Alice LIEU

Thesis for the degree of Doctor of Philosophy

in the

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Institute of Sound and Vibration Research

October 2019

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Institute of Sound and Vibration Research

Doctor of Philosophy

High-accuracy methods for frequency-domain flow acoustics

by Alice LIEU

Limited by the current computer memory capacities, it is challenging to solve the convected wave equation in heterogeneous media and at high wave number. This thesis is concerned with the investigation of efficient discretisation methods and solvers for frequency-domain flow acoustics. The application of computational modelling to wave propagation problems is hindered by the dispersion error introduced by the discretisation. Two common strategies to address this issue are to use high-order polynomial shape functions, or to use physics-based methods where the shape functions are local solutions of the problem. Both strategies have been actively developed over the past decades and have demonstrated their benefits compared to conventional finite-element methods, but they have yet to be compared. In this work, a high-order polynomial method and the wave-based discontinuous Galerkin method are compared for two-dimensional Helmholtz problems. A number of different benchmark cases are used to perform a detailed and systematic assessment of the relative merits of these two methods. The results indicate that the differences in performance, accuracy and conditioning are more nuanced than generally assumed. The performance of a method relies heavily on efficient solving procedures for the resulting large, sparse, complex linear systems. An alternative to purely iterative or direct solving procedures is to resort to domain decomposition methods. The Finite Element Tearing and Interconnecting method (FETI-2LM) employs Lagrange multipliers to recover the connections between the non-overlapping sub-domains. An iterative solution procedure is formulated in terms of unknowns defined only on the interfaces between sub-domains. The FETI approaches have been used extensively for Helmholtz problems and their performance is well documented for conventional finite elements. In this work, the FETI-2LM formulation is extended to the linearised potential theory for sound waves propagating in a potential base flow. In each sub-domain, a high-order finite element method is used to solve the governing equations. The proposed approach is validated on a number of two-dimensional test cases. In addition to the dependency on the mesh size, frequency, or number of subdomains, the influence of the interpolation order and Mach number on the scalability of the method is also assessed. The memory requirements for solving a simple three-dimensional problem is also evaluated and compared to that of a direct solver. Finally, the proposed method is applied to the problem of propagation of fan noise from the inlet of an engine, considering a realistic three-dimensional geometry and flow field.

Contents

List of Figures	9
List of Tables	13
Declaration Of Authorship	15
Acknowledgement	17
Nomenclature	19
1 Introduction	25
2 Review of High-order Methods	31
2.1 Physical Models	31
2.1.1 Conservation Equations	31
2.1.2 Linearised Euler Equations	32
2.1.3 Linearised Potential Flow Theory	33
2.1.4 Uniform Flow	34
2.1.5 Non-reflecting Boundary Conditions	35
2.1.5.1 Characteristics Method	35
2.1.5.2 Perfectly Matched Layer	35
2.2 Classical FEM and its Limitations	37
2.3 Spectral and High-order Polynomial Methods	38
2.3.1 Literature Review	38
2.3.2 Higher-order FEM	40
2.3.2.1 Statement of the Problem	40
2.3.2.2 Weak Variational Formulation	40
2.3.2.3 Discretisation	41
2.3.2.4 High-order Shape Functions	41
2.3.2.5 Condensation	43
2.3.2.6 Error Estimates	43
2.4 Physics-based Methods	44
2.4.1 Literature Review	44
2.4.2 Wave-based Discontinuous Galerkin Method	46
2.4.2.1 Statement of the Problem	46
2.4.2.2 Variational Formulation	46
2.4.2.3 Numerical Flux	47
2.4.2.4 Boundary Conditions	48

2.4.2.5	Interpolation Basis	49
2.4.2.6	Error Estimates	51
3	Performance Study of p-FEM and Wave-Based DGM	53
3.1	Description of the Test Cases	53
3.1.1	Plane Wave in Free Field Problem	53
3.1.2	Spinning Wave Problems	55
3.1.3	Corner Problem	56
3.2	Measures of Accuracy and Costs	57
3.3	Comparison of the Interpolation Properties	58
3.3.1	Anisotropy	60
3.3.2	Convergence Properties	61
3.4	Performance of the Numerical Methods	62
3.4.1	Convergence	63
3.4.2	Conditioning	66
3.4.3	Computational Costs	67
3.4.3.1	Fixed Accuracy and Frequency	68
3.4.3.2	Influence of the Target Accuracy	69
3.4.3.3	Influence of the Frequency	71
3.5	Additional Study Involving Flow	71
3.5.1	Anisotropy	72
3.5.2	Performance of the Numerical Methods	74
3.6	Practical Considerations	75
3.7	Conclusions	76
4	Review of Domain Decomposition Methods	79
4.1	On the Difficulty to Solve the Helmholtz Equation with Classical Solvers	79
4.2	Domain Decomposition Methods for Wave Problems	82
4.2.1	Original Schwarz Methods	83
4.2.2	Optimised Schwarz Conditions and Schur Methods	84
4.3	The FETI-2LM Method for the Helmholtz Equation	89
4.3.1	Formulation of the Sub-problems	89
4.3.2	Introduction of the Lagrange Multipliers	90
4.3.3	Interface Problem and Iterative Solver	91
4.3.3.1	Algebraic Form	91
4.3.3.2	Interface Problem	92
4.3.3.3	Iterative Algorithm	93
4.3.3.4	Computation of the Matrix-Vector Product	95
4.3.3.5	Partitioning Algorithm	95
4.3.4	Summary	96
4.4	Objectives	96
5	FETI-2LM and High-Order FEM	99
5.1	Two-dimensional Scalability Analysis	99
5.1.1	Description of the Test Case	99
5.1.2	Influence of the Number of Sub-domains	100
5.1.3	Influence of the Mesh Size	102

5.1.4	Influence of the Wavenumber	103
5.1.5	Influence of the Polynomial Order	104
5.1.5.1	Varying the Polynomial Order Only	104
5.1.5.2	Varying the Polynomial Order for a Fixed Resolution	105
5.2	Comparison of Computational Costs	106
5.2.1	Cost Distribution	107
5.2.2	Influence of the Polynomial Order	108
5.2.3	Influence of the Number of Processes/Sub-domains	110
5.3	FETI-2LM and an Adaptive Order FEM	112
5.3.1	About FEMAO	112
5.3.2	Domain Decomposition and Adaptive Order	113
5.3.3	Plane Wave Scattered by Sphere	113
5.3.3.1	Description of the Test Case	113
5.3.3.2	Varying the Number of Sub-domains	114
5.3.3.3	Varying the Frequency at Fixed Accuracy	117
5.3.3.4	Varying the Target Accuracy	117
5.4	Realistic 3D Application: Plane Wave Scattered by a Submarine	121
5.4.1	Description of the Test Case	121
5.4.2	Results	121
5.5	Summary	123
6	FETI-2LM for the Linearised Potential Equation	125
6.1	Introducing the Flow	125
6.1.1	Formulation of the Sub-problems	125
6.1.2	Algebraic Form	127
6.2	Two-Dimensional Study	128
6.2.1	Description of the Test Case	128
6.2.2	Validation on Non-Homogeneous Flow	129
6.2.3	Scalability Analysis	131
6.3	Realistic 3D Application	132
6.3.1	Description of the Test Case	133
6.3.2	Results	134
6.4	Summary	135
7	Conclusions	137
	Annexe A Scattering of a Plane Wave by a Cylinder	143
	Annexe B Odd and Even Number of Plane Waves	145

List of Figures

2.1	Example of triangular nodal, edge and bubble Lobatto shape functions. .	42
2.2	Plane waves discretisation in a triangle.	51
3.1	Examples of analytical solutions of the four benchmark problems ($kL = 50$). .	54
3.2	Examples of unstructured meshes used for the test cases.	54
3.3	Relative L^2 error associated with the best interpolation plotted against the direction of the incident plane wave with $h/a = 0.1$, $kL = 20$. Black solid lines: p -FEM, red dashed lines: wave-based DGM. Numbers indicates the polynomial order or the number of plane waves.	60
3.4	Relative best L^2 interpolation error (%) against D_λ at $kL = 50$; red dashed lines: wave-based method; black solid lines: p -FEM (with static condensation).	61
3.5	Relative L^2 -error (%) against D_λ at $kL = 50$; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation.	63
3.6	Optimality coefficient against D_λ for the propagating spinning wave problem at $kL = 50$; red solid lines: wave-based DGM; black solid lines: p -FEM with condensation.	64
3.7	Examples of unstructured non-uniform meshes used for the non-smooth test case with $h_{\max}/b = 0.17$	65
3.8	Relative L^2 -error (left) and condition number (right) against C_r at $kL = 50$ with $h_{\max}/b = 0.17$ for the non-smooth problem with locally refined meshes; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation.	66
3.9	Propagating spinning wave case with $kL = 50$ and $m = 10$. Left: Condition number as a function of D_λ . Right: relative L^2 error as a function of the condition number (left). Black solid lines: p -FEM with condensation, blue dotted lines: p -FEM without condensation, red dashed lines: wave-based DGM.	66
3.10	Factorisation time (s) against the element size required to obtain a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.	68
3.11	Factorisation memory against condition number for a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.	69
3.12	Number of non-zero entries against the number of degrees of freedom for a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.	70

3.13	Propagating spinning wave test case ($kL = 50$, $m = 10$); factorisation memory (left) and condition number (right) against D_λ to achieve different levels of accuracy; red dashed lines: wave-based DGM, solid black lines: p -FEM with static condensation.	70
3.14	Propagating spinning wave problem. Factorisation memory (left) and condition number (right) against D_λ to achieve 1% of accuracy at different frequencies; red dashed lines: wave-based DGM, solid black lines: p -FEM with static condensation.	71
3.15	Plane wave basis angles θ and α	72
3.16	Two approaches for the wave basis distribution with $N_w = 9$	73
3.17	Relative L^2 -error against the wave direction for $M = 0.5$, $h/a = 0.1$ and $kL = 50$; black lines: p -FEM with condensation, red lines: DGM with uniform wave distribution, magenta dashed lines: DGM with a non-uniform wave distribution. The number of plane wave or polynomial order is shown next to each curve.	73
3.18	Convected plane wave in free field problem with 1% of accuracy at $kL = 75$, $M = 0.5$; blue lines: p -FEM without condensation, black lines: p -FEM with condensation, red lines: DGM with uniform wave distribution, magenta dashed lines: DGM with non-uniform wave distribution. The numbers of plane wave or polynomial order are shown next to each point.	75
4.1	Original domain used by Schwarz; the complex domain is composed of two simple overlapping geometries, a disk Ω_1 and a rectangle Ω_2	83
4.2	Adjustment of the sign on both sides of the interfaces for a checkerboard partition.	88
4.3	Example of checkerboard splitting with $N_s = 4$	89
5.1	Example of different number of partitions for a given mesh $h/L = 1/60$	101
5.2	Influence of the number of sub-domains N_s . The interface residual and relative H^1 error against the iteration numbers. Fixed: $h/L = 1/60$, $kL = 100$, $p = 6$, $\varepsilon = 10^{-8}$	101
5.3	Influence of the mesh size h/L . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5$, $kL = 100$, $p = 6$, $\varepsilon = 10^{-8}$	103
5.4	Influence of the wavenumber kL . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5$, $h/L = 1/30$, $p = 6$, $\varepsilon = 10^{-8}$	104
5.5	Influence of the polynomial order p . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5$, $h/L = 1/60$, $kL = 100$, $\varepsilon = 10^{-8}$	105
5.6	Influence of the polynomial order p at $D_\lambda \approx 12$. The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5$, $kL = 100$, $\varepsilon = 10^{-8}$	106
5.7	Guided wave problem in 3D. Example of solution at $kL = 200$ (left) and mesh with $h/L = 1/26$ and $N_s = 8$ (right).	107
5.8	Guided wave problem at $kL = 200$ and mesh used to obtain $E_{L^2} \approx 1\%$ at $p = 8$ with $N_{\text{DOF}} = 5$ million, $N_s = 8$, $\varepsilon = 10^{-8}$	108
5.9	Characteristics of the FETI-2LM computation. Guided wave problem at $kL = 100$, $N_s = 4$, $\varepsilon = 10^{-8}$ with $E_{L^2} \approx 1\%$; $E_{H^1} \approx 2\%$	109
5.10	Comparison of the FETI-2LM with MUMPS. Guided wave problem at $kL = 100$, $N_s = 4$, $\varepsilon = 10^{-8}$ with $E_{L^2} \approx 1\%$; $E_{H^1} \approx 2\%$	110

5.11	Influence of the number of sub-domains N_s with $p = 8$. Guided wave problem at $kL = 200$, $\varepsilon = 10^{-8}$ with $E_{L^2} = 3.05\%$; $E_{H^1} = 4.56\%$	111
5.12	Distribution of the memory requirement over the $N_s = 16$ processes with $p = 8$. Guided wave problem at $kL = 200$, $\varepsilon = 10^{-8}$ with $E_{L^2} = 3.05\%$; $E_{H^1} = 4.56\%$	111
5.13	Example of an unstructured mesh (left) with a refinement coefficient of $1/3$ and a PML (in red) and solution at $kL = 200$	114
5.14	Influence of the number of sub-domains. Decomposition used. Plane wave scattered by a sphere with $h_{\max}/L = 1/30$	115
5.15	Influence of the number of sub-domains. Costs. Plane wave scattered by a sphere with $kL = 200$, $h_{\max}/L = 1/30$, $E_T = 5\%$	116
5.16	Influence of the Helmholtz number. Meshes and element orders (left) solution on sphere (right). Plane wave scattered by a sphere with $N_s = 8$, $h_{\max}/L = 1/30$, $E_T = 5\%$	118
5.17	Influence of the Helmholtz number. Plane wave scattered by a sphere with $N_s = 8$, $h_{\max}/L = 1/30$, $E_T = 5\%$	119
5.18	Influence of the target accuracy. Plane wave scattered by a sphere with $kL = 200$, $h_{\max}/L = 1/30$, $N_s = 16$	120
5.19	Shark-proof submarine, courtesy of Dav-dav (https://grabcad.com/library/submarine-shark-1).	121
5.20	Mesh and order distribution (left) and solution at $kL = 200$ (center and right). The black arrow shows the direction of the incident plane wave. . .	122
5.21	Submarine problem at $kL = 200$ with a target accuracy of 5%	123
6.1	Mean flow direction and magnitude with $M_\infty = -0.4$ (left); example of mesh $h = 0.125$ (right).	129
6.2	Solutions in the domain (left) and solutions on the cylinder (right) with $E_T = 1\%$	130
6.3	Influence of the number of sub-domains N_s with $M_\infty = -0.4$, $\omega = 25$, $h/R = 0.06$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$	131
6.4	Element orders for different Mach numbers M_∞ with $N_s = 8$, $\omega = 25$, $h = 0.06$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$	132
6.5	Influence of the Mach number M_∞ with $N_s = 8$, $\omega = 25$, $h = 0.06$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$	132
6.6	Influence of the frequency with $M_\infty = -0.4$, $N_s = 8$, $h = 0.025$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$	133
6.7	Geometry of the engine intake (left) and local Mach numbers for the sideline case (right); the black dotted lines indicate the start of the PML. . .	133
6.8	Generic intake at 1BPF, mode (1, 24): polynomial orders used (left); real part of the velocity potential (right).	134
6.9	Generic intake at 1BPF, mode (1,24): interface residual as a function of the number of iterations (left); memory requirements for each process (right).	135
A.1	Scattering of a plane wave by a cylinder at $kL = 50$ and $h/L = 1/20$	143
A.2	Scattering of a plane wave by a cylinder. Relative L^2 -error (%) against D_λ at $kL = 50$; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation. For each curve, D_λ is increased by refining the mesh while p or N_w is fixed.	144

-
- B.1 Spinning wave problem with $kL = 50$, $m = 10$. The relative L^2 error against D_λ (left) and against the condition number (right). Black lines: p -FEM with condensation; blue lines: p -FEM without condensation; red dashed lines: wave-based DGM with $N_w = 2n$, magenta dashed lines wave-based DGM with $N_w = 2n - 1$ 146
- B.2 Spinning wave problem with $kL = 50$, $m = 10$. The relative L^2 error against D_λ (left) and against the condition number (right). Black lines: p -FEM with condensation; blue lines: p -FEM without condensation; red dashed lines: wave-based DGM with $N_w = 2n$, magenta dashed lines wave-based DGM with $N_w = 2n - 1$ 146

List of Tables

5.1	Influence of the number of sub-domains N_s . Fixed: $h/L = 1/60, kL = 100, p = 6, \varepsilon = 10^{-8}$	101
5.2	Influence of the mesh size h/L . Fixed: $N_s = 5, kL = 100, p = 6, \varepsilon = 10^{-8}$	102
5.3	Influence of the wavenumber kL . Fixed: $N_s = 5, h/L = 1/30, p = 6, \varepsilon = 10^{-8}$	104
5.4	Influence of the polynomial order p . Fixed: $N_s = 5, h/L = 1/60, kL = 100, \varepsilon = 10^{-8}$	105
5.5	Influence of the polynomial order p . Fixed: $N_s = 5, D_\lambda \approx 12, kL = 100, \varepsilon = 10^{-8}$	106
6.1	Influence of the number of sub-domains N_s with $M_\infty = 0.4, \omega = 25, h/R = 0.06, E_T = 0.5\%$ and $\varepsilon = 10^{-8}$	131

Declaration of Authorship

I, Alice Lieu, declare that the thesis entitled High accuracy methods for frequency-domain flow acoustics and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
 - A comparison of high-order finite element method and wave-based discontinuous Galerkin method for Helmholtz problems, A. Lieu, G. Gabard and H. Bériot. *The 12th International Conference on Mathematical and Numerical Aspects of Wave Propagation*, Karlsruhe, Germany, 2015.
 - A performance study of high-order finite element and wave-based discontinuous Galerkin methods, A. Lieu, G. Gabard and H. Bériot. *The 22th International Congress on Sound and Vibration*, Florence, Italy, 2015.
 - A comparison of high-order polynomial and wave-based methods for Helmholtz problems, A. Lieu, G. Gabard and H. Bériot. *Journal of Computational Physics*, 321:105-125, 2016.

- Performance of a domain decomposition method for high-order FEM, A. Lieu, H. Bériot, G. Gabard and F.-X. Roux, *European Congress on Computational Methods in Applied Sciences and Engineering*, Crete Island, Greece, 2016.
- A domain decomposition method with high-order finite elements for flow acoustics, A. Lieu, G. Gabard, H. Bériot, *25th AIAA/CEAS Aeroacoustics Conference*, Delft, The Netherlands, 2019.

Acknowledgement

My first thoughts go to my supervisors Dr. Gwénaél Gabard and Dr. Hadrien Bériot. Not only knowledgeable but also kind and extremely patient, they taught me a great deal about scientific research and I would like to thank them for their guidance. They were always available for me, motivating and enthusiastic. I am deeply grateful for having them as mentors.

My gratitude also goes to Prof. François-Xavier Roux who kindly took the time to tell me about the intricate world of solvers and domain decomposition techniques. I am taking this opportunity to thank Prof. Jeremy Astley and Dr. Michel Tournour for their constructive comments and feedback on my work; and Dr. Alan McAlpine for his help towards the end.

I would also like to thank my friends at the ISVR in Southampton and at Siemens in Leuven for brightening up my days. In particular, thanks to Marine and Célia who were always there to make me laugh, even during the most difficult moments.

I want to express my gratitude to my huge family, always truthful and caring. I am especially grateful for my brothers and sisters, who are always loving in their own special way. On ne le dit jamais parce que “*c’est la famille*”, mais juste cette fois-ci : merci. I am thankful for Dr. A.G.P., for his technical advice and emotional support: thank you for always laughing “with me” and not “at me”.

I am especially indebted to my parents who showed me unconditional love and support throughout the past 29 years. They have crossed oceans to provide my brothers and me with a better life. They have pushed me towards long studies hoping that it would bring me an effortless life. A special thought goes to my dad who woke up every exam morning to make me breakfast and always treated me with my favourite diners afterwards. My parents have taught me endurance and optimism which were much needed during the long journey that led to this thesis. I did my best and I would not be here if it was not for them.

Last but not least, I thank Dr. Alexander Pretchel for helping me get onto my ‘two’ feet enabling me to complete this work.

This research would not have been possible without the financial support of the European Commission under the 7th Framework Programme People 2013 through the CRANE Marie Skłodowska-Curie Initial Training Network (606844). I acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

Nomenclature

Superscripts

- $(\cdot)^{\text{opt}}$ Optimal quantities (minimising the L^2 error)
- $(\cdot)^{\text{T}}$ Transpose
- $(\cdot)^h$ Discretised variable
- $\overline{(\cdot)}$ Complex conjugate
- $(\tilde{\cdot})$ Global quantities (sum of mean value and perturbation)

Subscripts

- $(\cdot)_0$ Base flow quantities
- $(\cdot)_\infty$ Far field quantities
- $(\cdot)_{\text{ex}}$ Exact solution
- $(\cdot)_e$ Quantity associated to element e

Roman Letters

- a Side length of the square domain
- \mathbf{A} Coefficient matrix for the 2D convected Helmholtz equation
- B Number of fan blades
- b Side length of the L-shape domain
- $\mathbb{B}_{i,j}$ Restriction operator from Ω_i to $\Gamma_{i,j}$
- \mathbf{B} Coefficient matrix for the 2D convected Helmholtz equation
- \mathbf{b} Right-hand side vector of the discretised problem
- $\mathcal{B}_{i,j}$ Transmission operator on $\Gamma_{i,j}$
- b_w Bandwidth of a sparse matrix

c	Speed of sound
C_{opt}	Optimality coefficient
C_r	Refinement factor
\mathbf{d}	RHS of the interface problem
d	Number of space dimension
D_λ	Number of degrees of freedom per wavelength
E_X	Relative error using the X norm
f	Frequency
\mathbb{F}_I	Interface matrix
$\mathbf{f}_{e,e'}$	Numerical flux on the edge between element e and e'
g	Source term
\mathbf{g}	Source term vector
\mathbf{F}	Flux matrix
h	Average element size
H_m	Hankel function of order m
\mathbb{I}	Identity tensor
i	Imaginary unit
\mathfrak{J}	Cost function
J	Bessel function of the first kind
k	Wavenumber
\mathbb{K}	Matrix arising from discretisation of the governing equations
\mathcal{K}_m	Krylov subspace
L	Characteristic length
\mathbb{L}	Lower triangular matrix
\mathcal{L}_q	Legendre polynomial of order q
l_q	Lobatto polynomial of order q
M	Mach number

m	Order of the spinning waves
$\mathbb{M}_{i,j}$	Mass matrix on $\Gamma_{ji,j}$
\mathcal{N}_i	Set of sub-domains indices that are neighbours to Ω_i
N	Dimension of the linear system to be solved
N_{DOF}	Total number of degrees of freedom
$N_{\text{DOF}}^{\text{I}}$	Size of the interface problem
N_{w}	Number of plane waves in the basis
N_{NZ}	Number of non-zero entries in the sparse global matrix
N_{s}	Number of sub-domains
\mathbf{n}	Outward normal vector with coordinate n
n	Normal coordinate
n_{h}	Number of harmonics
n_{c}	Number of iterations used to converge
p	Polynomial order
R	Radius of the inner cylinder
r_{g}	Residual on the global problem
r_{I}	Residual on the interface problem
s	Entropy
t	Time
u	Pressure
\mathbf{u}	Vector of linear perturbation of density, x -momentum and y -momentum
\mathbb{U}	Upper triangular matrix
\mathbf{v}	Velocity
\mathcal{V}	Volume
w	Test function associated to u
\mathbf{w}	Test function associated to \mathbf{u}
\mathbf{W}	Matrix of eigenvectors

\mathbf{x} Cartesian coordinates (x, y, z) in the physical space

Greek Letters

δ_X	Relative X norm of the differences between the iterative and direct solutions
ε	Tolerance on the interface residual
η_i	i^{th} affine coordinate
γ	Ratio of specific heats
Γ_i	Section i of the boundary $\partial\Omega$
$\Gamma_{e,e'}$	Edge between the element e and e'
$\Gamma_{i,j}$	Interface between sub-domain i and sub-domain j and belonging to sub-domain i
κ	Condition number
λ	Wavelength
$\mathbf{\Lambda}$	Matrix of eigenvalues
$\lambda_{i,j}$	Lagrange multiplier at the interface $\Gamma_{i,j}$
Ω	Numerical domain
ω	Angular frequency
Ω_i	Sub-domain i
ϕ	Velocity potential
φ	Shape function
φ^b	Bubble shape function
$\varphi_q^{e_i}$	Edge shape function of q order associated to the i^{th} edge
φ^{v_i}	Vertex shape function associated to the i^{th} node
ψ	Test function associated to ϕ
Ψ_q	Kernel function of order q
ρ	Density
σ	Stress tensor
τ	Tangent vector
\mathbb{T}	Viscous stress tensor

θ_0	Wave propagation direction
θ_0	Wave propagation angle
$\xi_{i,j}$	Test function associated with $\lambda_{i,j}$

Other Symbols

(ζ, ν)	Cartesian coordinates in the reference space
$\partial\mathcal{V}$	External boundary of the volume \mathcal{V}
$\partial\Omega$	Boundary of the numerical domain
$D(\cdot)/Dt$	Material derivative
(r, θ)	Polar coordinates in the physical space

Chapter 1

Introduction

Aviation is a vital sector of European society and economy. It brings people together and connects Europe to the rest of the world. Involving millions of jobs, it is as well an important generator of prosperity and wealth as it eases the development of new markets and international trade. However, aviation has significant impacts on the environment: relying mainly on fossil fuel, it vastly participates to climate change. An even more perceptible problem is the impact of community and ramp noise caused by aircraft during landing and take-off on the health of populations living and working close to airports [9].

In Europe, nowadays, hundreds of millions of people are exposed to noise levels which can induce serious health problems. Epidemiological studies have shown that frequent exposure to road and air traffic noises may cause sleeping disturbances, tinnitus, hearing loss, or cardiovascular diseases [64]. Noise pollution also highly affects wildlife [101, 65, 149, 135].

The environmental impact of aircraft noise occurs mainly during take-off and landing. Diverse strategies have been put in place to mitigate it. For example, steeper angles of descent may be used such that the time spent at low altitudes is limited; the lowering of the landing gears, which increases the airframe noise, may also be delayed. Other approaches aim at reducing engine noise, which embodies at least four different sources of disturbance, namely fan noise which propagates through the intake and the bypass duct; compressor noise which radiates from the intake; turbine and combustor noise which propagate to the rear of the engine; and jet noise downstream from the engine [7]. It is generally admitted that fan and jet noise are dominant at take-off, whereas fan and airframe noise are more dominant at approach. Intake and bypass noise can be mitigated using acoustic liners or by modifying the design, while jet noise is typically reduced by using lower jet speeds or more recently, by incorporating chevrons at the bypass and core duct exhausts [123, 112]. Since the advent of commercial jet aircraft in the sixties, noise reducing technologies have advanced considerably [91]. Current aircraft

are significantly quieter than first generation turbofan aircraft. However, air traffic is expected to increase at a rate of 5% per year. Simultaneously, environmental and noise regulations are becoming stricter. For instance, the European Union's objectives for the aviation sector is to reduce the aircraft noise emission by 65% before 2050 compared to 2000 [132]. Innovative and disruptive noise reduction techniques therefore need to be introduced in order to meet these new requirements. The understanding, predicting and reducing aircraft noise has hence become one of the major challenges faced by the aviation industry.

Project motivations

One obvious way to develop quieter aircraft engines is through empirical research. However, this type of investigation is very expensive and inefficient as it requires manufacturing, testing, redesigning... Numerical simulations are a recognised alternative to those costly flight and static tests. They are used to support and optimise the acoustic designs of aircraft engines. Traditional Computational Fluid Dynamics (CFD) methods were found inadequate to model sound propagation. The acoustic perturbations and the flow field have different length scales and amplitudes. Therefore, solving both scales would be computationally expensive. Hence, at the end of the eighties, for the purpose of understanding the physics of noise generation and propagation, Computational Aero-Acoustics (CAA) emerged as a specific research area. Nowadays, most noise prediction methods uses a two-stage calculation. The sources are first characterised using CFD methods, then, the propagation of the sound generated by these sources is computed using CAA. Of these two stages, the present project focuses on the computational methods for sound propagation.

Generally, three-dimensional simulations are required for multiple frequencies, sources, operating conditions (take-off, landing ...). But the time and computational resources needed depend on the problem size and the prescribed accuracy. Even though computational power has considerably increased during the last decades, performing large-scale CAA simulations taking into account the full physical and geometrical complexity of the aircraft remains a challenge. For instance, a full three-dimensional turbofan engine problem has to be split into sub-systems (inlet/exhaust) to obtain a solution within reasonable time. Yet, the resolution of these smaller solutions can be computationally demanding. Hence, novel methods that are faster and relatively inexpensive in terms of computational resources have to be developed. These are the main goals of the project Community Ramp Aircraft NoiseE (CRANE), which is an European Marie Skłodowska-Curie action. The project intends to improve the efficiency and reliability of CAA predictions and involves the Institute of Sound and Vibration Research at the University of Southampton (United Kingdom) and Siemens Industry Software (Belgium) as main

partners and the Katholieke Universiteit Leuven (Belgium) and Rolls-Royce Deutschland (Germany) as associated partners. The main applications are sound radiation from the exhausts and intakes of turbofan engines, usually associated with challenging strong flow inhomogeneities and complex geometries.

For most aero-acoustic engineering applications, a common mathematical model used is the linearised Euler equations. This model describes the propagation of linear perturbations on an inviscid mean flow field. As this model admits acoustic, entropy and vortical disturbances, the refraction of acoustic waves by the rotational mean flow are taken into account (important for applications such as the mixing layer of a jet). But when solved in the time domain, the model supports hydrodynamic instabilities such as the Kelvin–Helmholtz instability: these disturbances grow as they move away from their sources and may at some point overcome the acoustic perturbations. An alternative is to solve the equations in the frequency-domain [2]. This relatively general approach is costly to accurately solve due to the large number of variables involved. In the framework of the CRANE project, the present study is aimed at improving the description of sound propagation from the inlet of the engine. Therefore, another more specific model can be used: the convected Helmholtz equation. The mean flow field and perturbations are considered irrotational and homentropic and only one variable is involved. This equation is cheaper to solve and although it does not include the full physics of the problem, it still captures most of the flow effects of the inlet. It is therefore a relevant prototype problem to test numerical methods.

Scope of the thesis

High-frequency acoustic problems call for fine meshes which lead to the solutions of large linear systems. Limited by the current computer memory capacities, the convected wave equation in heterogeneous media and at high wavenumber is therefore challenging to solve. To bypass this issue, the discretisation methods and solvers have to be improved. The aim of this work is to investigate efficient modelling techniques and solvers that outperform the ones that are currently employed in commercial aero-acoustic prediction codes. Realistic three-dimensional test cases will be solved in addition to more conventional academic cases.

Most of the computational aero-acoustics tools used in industry still rely on low-order discretisation methods which suffer from the so-called pollution effect at high frequencies [42, 14]. This effect is the consequence of the accumulation of dispersion error over the whole numerical domain. It has been demonstrated that the efficiency with which one models short wavelength problems can be increased by using higher order methods. For example, to achieve a solution with a target accuracy, it is well-known that higher-order approximations allow a reduction of the resolution requirements and therefore the total

number of degrees of freedom needed. The principle is the use either of higher-order polynomials [158, 29, 159] or of local solutions of the physical problem [122, 53, 35, 66]. Of the numerous advanced methods developed to prevent high dispersion errors for short waves problems, we choose to focus on a high-order finite element method using Lobatto functions and a wave-based method employing a set of plane waves to interpolate the solution. Previous studies have demonstrated that both methods lead to an effective control of the pollution effect associated with the classical finite element method at high frequency. The higher-order finite element method takes advantage of the superior interpolation properties of the Lobatto functions compared to the conventional low-order Lagrange basis. It is also well established in the numerical acoustic community and has been recently implemented in a commercial code [21]. The wave-based method is part of the physics-based methods, in which the idea is to include some *a priori* knowledge about the local behaviour of the solution into the numerical model (for instance the wavenumber). Common belief is that compared to polynomial methods, physics-based methods can provide a significant improvement in performance, at the expense of a deterioration of the conditioning [68]. Yet, to our knowledge, no detailed comparison of a high-order polynomial method and a physics-based method can be found in the literature. Therefore, in the first part of the present work, the two high-order methods are benchmarked on four different two-dimensional Helmholtz problems and their performances are compared. The objective is to identify the ranges of frequency over which the high-order finite element method outperforms the wave-based method, and vice-versa.

The performance of a method relies heavily on efficient solving procedures for the resulting large, sparse, complex linear systems. Direct solving procedures are robust and predictable but do not scale well with respect to the problem size and frequency. This is an issue as the computational demand for the solution of large-scale problems is always ahead of computer capabilities. Continued efforts are deployed to develop efficient iterative solvers. These methods scale well with the problem size, are highly parallelisable and generally require a limited amount of computational resources. However, due to the properties of the underlying operator, the design of robust iterative solvers for Helmholtz problems remains a challenge. An alternative to purely iterative or direct solving procedures is to resort to domain decomposition methods. The idea is to split the numerical domain into smaller sub-domains and to solve the sub-problems separately. The sub-domains are related to each other only via the interface variables. So, if each sub-problem is assigned to one processor, the data to be exchanged is limited to the interface between the sub-domains. Domain decomposition can therefore ease the inter-processor communications which are the most penalising steps in parallel computing. The domain decomposition framework is naturally adapted for building efficient algorithms running on parallel computers but the methods are still profitable when used on single processor computers (solution of smaller problems). A non-overlapping domain decomposition method called the Finite Element Tearing and Interconnecting method

[60, 58, 55] employs two Lagrange multipliers to recover the connections between the sub-domains. This method is referred to as the FETI-2LM method [41, 145]. It adopts an iterative procedure to solve an interface problem whereas a direct solver is used for the local problems. The performance of this method is well documented for conventional finite elements. While the FETI approach has been used extensively for Helmholtz problems, to the author's knowledge, it has not been used for solving aero-acoustics problems. Therefore, in this work, we choose to focus on the performance of the FETI-2LM for aero-acoustic problems combined with a high-order FEM discretisation.

Outline of the thesis

The manuscript is organised as follows.

Chapter 2 introduces the hypothesis and model equations used to describe the propagation of time-harmonic acoustic waves in moving media. A literature survey of high-order polynomial and physics-based methods is also presented, justifying the choice of the two discretisation methods hereafter compared.

Chapter 3 is dedicated to the comparison of the high-order finite element method and the wave-based discontinuous Galerkin method. Four different two-dimensional Helmholtz problems including propagating and evanescent waves and smooth and non-smooth solutions are used such that the ability of the methods to tackle a large variety of problem is assessed. The chapter includes the presentation of the test cases, the assessment of the interpolation properties of the bases, anisotropy, convergence, conditioning and performance of the methods. This first part of the study concludes that there is still a significant benefit in using polynomial methods, both in terms of performance and robustness, even at high frequency. For the remainder of the work, we therefore choose to investigate solving procedures using the matrix systems arising from the high-order finite element discretisation.

Chapter 4 aims at outlining the difficulties inherent to the solving of the discrete Helmholtz equation with classical direct and iterative solvers. The advantages of the domain decomposition methods are brought up and a brief review is given.

Chapter 5 is dedicated to the application of the FETI-2LM solver to Helmholtz problems discretised using the high-order FEM. A two-dimensional scalability analysis is conducted and the influence of the polynomial order on the convergence of the iterative technique is examined. The FETI-2LM solver is then benchmarked against a state-of-the-art parallel direct solver for a simple three-dimensional Helmholtz solution. The method is finally applied to a large scale acoustics application, representing the scattering of a plane wave by a submarine.

Chapter 6 focuses on the generalisation of the FETI-2LM methods to flow acoustics applications. A new formulation is proposed and verified numerically, and a scalability study is carried out in order to assess the impact of the Mach number on the convergence. The proposed method is then demonstrated on a realistic three-dimensional air intake noise radiation application.

This thesis ends with a conclusion and a presentation of the possible future work.

Chapter 2

Review of High-order Methods

In this chapter, the physical models describing the propagation of sound waves in non-uniform flows are introduced. The main assumptions leading to the governing equations for the problems considered in this work are outlined. The limitations of the standard Finite Element Method are recalled and two advanced numerical methods, namely the high-order finite element method (p -FEM) and the wave-based Discontinuous Galerkin Method (DGM) are introduced.

2.1 Physical Models

In this section, the governing equations used to model the convected acoustic waves are presented [67]. Starting from the conservation equations, the linearised Euler and linearised potential flow equations are derived. The particular case of uniform flow is also examined as it gives insights into the physical properties of acoustic waves in moving media. The boundary conditions used in this work, required to close the problems, are also given.

2.1.1 Conservation Equations

We consider a volume \mathcal{V} of fluid, with an external boundary $\partial\mathcal{V}$, a density $\tilde{\rho}(\mathbf{x}, t)$, a pressure $\tilde{u}(\mathbf{x}, t)$ and a velocity $\tilde{\mathbf{v}}(\mathbf{x}, t)$. Under classical continuum mechanics assumptions, the motion of fluid in any point of the volume is governed by the conservation of the mass, momentum and energy.

The mass conservation reads:

$$\frac{D\tilde{\rho}}{Dt} + \tilde{\rho}\nabla \cdot \tilde{\mathbf{v}} = 0, \quad (2.1)$$

where $D(\cdot)/Dt = \partial(\cdot)/\partial t + \tilde{\mathbf{v}} \cdot \nabla(\cdot)$ is the material derivative representing the rate of change when moving with a fluid element and $\nabla \cdot \tilde{\mathbf{v}}$ represents the rate of dilatation.

The conservation of momentum —equivalent to Newton's second law applied to a fluid element— is written:

$$\tilde{\rho} \frac{D\tilde{\mathbf{v}}}{Dt} = \nabla \cdot \boldsymbol{\sigma} ,$$

where $D\tilde{\mathbf{v}}/Dt$ is the acceleration of a fluid element and $\boldsymbol{\sigma}$ is the stress tensor. It can be decomposed into:

$$\boldsymbol{\sigma} = -\tilde{u}\mathbb{I} + \mathbb{T} ,$$

where \mathbb{I} is the identity tensor and \mathbb{T} is the viscous stress tensor. If the flow is assumed inviscid, the viscous tensor is zero. The equation for conservation of momentum becomes:

$$\tilde{\rho} \frac{D\tilde{\mathbf{v}}}{Dt} + \nabla \tilde{u} = \mathbf{0} . \quad (2.2)$$

A third equation is retrieved by combining the equation of state and the mass conservation equation. A fluid particle is assumed to be a closed thermodynamic system at equilibrium meaning that it does not exchange any matter with its surrounding. It has a density $\tilde{\rho}$ and an entropy \tilde{s} . For a fluid element, the pressure is: $\tilde{u} = \tilde{u}(\tilde{\rho}, \tilde{s})$. If the flow is assumed isentropic, which means that the entropy of a fluid element is constant in time, it yields:

$$\frac{D\tilde{u}}{Dt} = \tilde{c}^2 \frac{D\tilde{\rho}}{Dt} . \quad (2.3)$$

where $\tilde{c}^2 = \partial\tilde{u}/\partial\tilde{\rho}$ and \tilde{c} is the speed of sound. For an ideal gas, the homentropy assumption leads to $\tilde{u} = C\tilde{\rho}^\gamma$ where C is a constant and γ is the ratio of the specific heats. Therefore, using the definition of the speed of sound, we have:

$$\tilde{c}^2 = \gamma \frac{\tilde{u}}{\tilde{\rho}} . \quad (2.4)$$

Combining the mass equation (2.1), (2.3) and (2.4), an energy equation for pressure can be written:

$$\frac{D\tilde{u}}{Dt} + \gamma \tilde{u} \nabla \cdot \tilde{\mathbf{v}} = 0 . \quad (2.5)$$

Under the isentropic and homentropic assumptions, the inviscid motion of a Newtonian fluid obeying the ideal gas law is therefore governed by (2.1), (2.2) and (2.5).

2.1.2 Linearised Euler Equations

The assumption of small perturbations allows the flow field to be decomposed into a base flow and perturbations:

$$\tilde{u}(\mathbf{x}, t) = u_0(\mathbf{x}) + u(\mathbf{x}, t) , \quad \tilde{\rho}(\mathbf{x}, t) = \rho_0(\mathbf{x}) + \rho(\mathbf{x}, t) , \quad \tilde{\mathbf{v}}(\mathbf{x}, t) = \mathbf{v}_0(\mathbf{x}) + \mathbf{v}(\mathbf{x}, t) .$$

The mean flow quantities (discernible by the subscript 0) are steady and of arbitrary amplitude, unlike the perturbations (unscripted variables) which are of small amplitude

compared to the base flow and unsteady. The decomposition is substituted into the previous Euler equations. The non-linear terms (products of small terms) can be neglected. Two sets of equations are obtained.

A first set describing the base flow only:

$$\begin{aligned}\nabla \cdot (\rho_0 \mathbf{v}_0) &= 0 , \\ \rho_0 \frac{D_0 \mathbf{v}_0}{Dt} + \nabla u_0 &= \mathbf{0} , \\ \frac{D_0 u_0}{Dt} + \gamma u_0 \nabla \cdot \mathbf{v}_0 &= 0 ,\end{aligned}$$

where $D_0(\cdot)/Dt = \partial(\cdot)/\partial t + \mathbf{v}_0 \cdot \nabla(\cdot)$ is the material derivative in the base flow (time variation seen by an observer moving with the mean flow). It should be noted that the perturbations do not appear in this set of equations; the base flow is not affected by the perturbations. These equations are usually solved using CFD methods and the base flow solution is used as an input for the second system of equations, solved using CAA methods.

The second set of equations are called the linearised Euler equations (LEE):

$$\frac{D_0 \rho}{Dt} + \nabla \cdot (\rho_0 \mathbf{v}) + \rho \nabla \cdot \mathbf{v}_0 = 0 , \quad (2.6a)$$

$$\rho_0 \frac{D_0 \mathbf{v}}{Dt} + \rho_0 (\mathbf{v} \cdot \nabla) \mathbf{v}_0 + \rho (\mathbf{v}_0 \cdot \nabla) \mathbf{v} + \nabla u = \mathbf{0} , \quad (2.6b)$$

$$\frac{D_0 u}{Dt} + \mathbf{v} \cdot \nabla u_0 + \gamma u_0 \nabla \cdot \mathbf{v} + \gamma u \nabla \cdot \mathbf{v}_0 = 0 . \quad (2.6c)$$

The LEEs describe the propagation of small perturbations on an inviscid mean flow. Heat transfers, viscosity, non-linearity and effect of the perturbations on the mean flow field have been neglected. This model admits acoustic, entropic and vortical disturbances. When solved in the time-domain, the system supports hydrodynamic instabilities which may ruin the solution. One way to alleviate this issue, is to formulate and solve the equations in the frequency-domain [2]. However, solving the LEEs remains very computationally intensive because of the large number of unknowns involved [140, 83].

2.1.3 Linearised Potential Flow Theory

In order to further simplify the model, the flow is assumed to be irrotational. In that particular case, the flow may be non-uniform but the velocity is derived from a potential $\tilde{\phi}$: $\tilde{\mathbf{v}} = \nabla \tilde{\phi}$. For an ideal gas in a compressible, homentropic, inviscid and irrotational flow, and in the absence of external forces, Bernoulli's equation reads [141]:

$$\frac{\partial \tilde{\phi}}{\partial t} + \frac{\tilde{\mathbf{v}} \cdot \tilde{\mathbf{v}}}{2} + \frac{\tilde{c}^2}{\gamma - 1} = \text{constant} .$$

Applying the linearisation $\tilde{\phi} = \phi_0 + \phi$, one obtains:

$$\frac{\partial \phi}{\partial t} + \mathbf{v}_0 \cdot \nabla \phi + \frac{(c^2)}{\gamma - 1} = 0 ,$$

where $(c^2) = c_0^2(\gamma - 1)\rho/\rho_0$. Finally, Bernoulli's equation yields:

$$\rho = -\frac{\rho_0}{c_0^2} \frac{D_0 \phi}{Dt} . \quad (2.7)$$

By substituting (2.7) into the linearised mass equation (2.6a), the linearised potential equation is found:

$$\rho_0 \frac{D_0}{Dt} \left(\frac{1}{c_0^2} \frac{D_0 \phi}{Dt} \right) - \nabla \cdot (\rho_0 \nabla \phi) = 0. \quad (2.8)$$

This equation is less computationally expensive compared to the complete LEE, as it involves only one variable. However, it is less general because of the absence of vorticity in the base flow and perturbations. As mentioned in the introduction, the irrotational assumption is still valid if we are interested in describing the inlet of the engine [7].

2.1.4 Uniform Flow

If the flow is uniform, the linearised Euler equations become:

$$\frac{D_0 \rho}{Dt} + \rho_0 \nabla \cdot \mathbf{v} = 0, \quad (2.9a)$$

$$\rho_0 \frac{D_0 \mathbf{v}}{Dt} + \nabla u = \mathbf{0}, \quad (2.9b)$$

$$\frac{D_0 u}{Dt} + \gamma u_0 \nabla \cdot \mathbf{v} = 0. \quad (2.9c)$$

A scalar equation involving only u can be obtained by considering:

$$\frac{1}{\gamma u_0} \frac{D_0}{Dt} (2.9c) - \frac{1}{\rho_0} \nabla \cdot (2.9b)$$

The resulting convected wave equation is:

$$\frac{1}{c_0^2} \frac{D_0^2 u}{Dt^2} - \nabla^2 u = 0. \quad (2.10)$$

Using (2.8), an alternative equation in terms of velocity potential for uniform flows can be obtained. The convected wave equation describes the propagation of acoustic perturbations in a uniform mean flow. This model does not support hydrodynamic or entropic fluctuations.

2.1.5 Non-reflecting Boundary Conditions

The previous equations have to be supplemented with boundary conditions to close the problems. In this work, we mostly consider exterior sound radiation applications in infinite domains. These are handled by introducing non-reflecting boundary conditions at the boundaries of the bounded computational domain Ω . These fictitious boundaries are designed to absorb the outgoing waves with minimal spurious reflections. In this section, two non-reflecting boundary conditions are introduced, the characteristics method and the Perfectly Matched Layer approach.

2.1.5.1 Characteristics Method

The characteristics boundary condition is used when the solution outside the numerical domain is known. This method was first proposed by Thompson in 1987 for the solution of hyperbolic systems in the time domain [152]. The idea is to split the wave at the boundary into an incoming part, which is to be specified, and an outgoing part, which is determined by the system being solved.

Let us consider the boundary outward normal direction \mathbf{n} and the corresponding coordinate n . At the boundary, in the normal direction, the acoustic wave is written as the sum of an incoming wave and an outgoing wave of angular frequency ω :

$$\begin{aligned} u(n) &= \mathcal{A}e^{-ik^+n} + \mathcal{B}e^{-ik^-n}, \\ \nabla u \cdot \mathbf{n} &= -\mathcal{A}ik^+e^{-ik^+n} - \mathcal{B}ik^-e^{-ik^-n}, \end{aligned}$$

where \mathcal{A} and \mathcal{B} are the wave amplitudes corresponding to the outgoing and incoming waves respectively, and $k^\pm = \omega/(\mathbf{v}_0 \cdot \mathbf{n} \pm c_0)$. The incoming wave is isolated by expressing $\nabla u \cdot \mathbf{n} + ik^+u$:

$$\nabla u \cdot \mathbf{n} + ik^+u = \mathcal{B}i(k^+ - k^-)e^{-ik^-n}$$

This is equivalent to a generalised Robin boundary condition which reads:

$$\nabla u \cdot \mathbf{n} + ik^+u = g, \tag{2.11}$$

where g is a given source term. If the solution outside the numerical domain is u_{ex} , the source term is given by $g = \nabla u_{\text{ex}} \cdot \mathbf{n} + ik^+u_{\text{ex}}$.

2.1.5.2 Perfectly Matched Layer

To avoid spurious reflections when the solution outside the domain is unknown, another technique is —instead of looking for an absorbing boundary condition— to create a

boundary layer that absorbs the waves entering it. The Perfectly Matched Layer (PML) was originally devised by Béranger for electromagnetic problems [18]. The method has since been used to solve the Helmholtz equation [16] and the LEE [51].

The principle of the PML lies in the application a complex coordinate stretching. When a wave leaves the computational domain and enters the absorbing layer, it decays exponentially. Even though a reflected component may be created at the exterior boundary of the PML, after a trip in the absorbing region, the effect of the reflected wave on the solution becomes in practice negligible. Let us now illustrate the method on a two dimensional radiation problem with Cartesian coordinates.

Assuming an implicit time dependence $e^{i\omega t}$ where ω is the angular frequency and a quiescent medium, in the absence of volume sources, an exterior Helmholtz problem can be defined as follows:

$$\begin{cases} k_0^2 u + \nabla^2 u & = 0, & \text{in } \Omega, \\ \nabla u \cdot \mathbf{n} + ik_0 u & = g, & \text{on } \Gamma_r, \\ \lim_{r \rightarrow \infty} r^{\frac{d-1}{2}} (\nabla u \cdot \mathbf{n} + ik_0 u) & = 0, & \end{cases} \quad (2.12)$$

where $k_0 = \omega/c_0$, g is a given source term. Γ_r can be the surface of a object within the domain. In (2.12), the Robin boundary condition presented above is used on Γ_r . The third equation is referred to as the Sommerfeld radiation condition. It ensures the uniqueness of the exterior Helmholtz problem and expresses that at infinity, the scattered waves are only outgoing. To ensure that this condition holds, the computational domain Ω is extended with a PML region Ω_{PML} . The interface between Ω and Ω_{PML} is denoted Γ_{in} whereas the exterior boundary of Ω_{PML} is referred to as Γ_{out} .

In order to absorb waves that enter the regions defined by $|x| > x_{\text{in}}$ and $|y| > y_{\text{in}}$, the following complex transformation is applied:

$$\hat{x} = x - i \int_{x_{\text{in}}}^x \sigma_x(s) ds, \quad \hat{y} = y - i \int_{y_{\text{in}}}^y \sigma_y(s) ds, \quad (2.13)$$

where $\sigma_x(x)$ and $\sigma_y(y)$ are positive absorption functions varying in the x and y direction respectively.

In Ω_{PML} , the governing equation becomes:

$$k_0^2 \hat{u} + \frac{\partial^2 \hat{u}}{\partial \hat{x}^2} + \frac{\partial^2 \hat{u}}{\partial \hat{y}^2} = 0 \quad (2.14)$$

where the derivatives can be written as:

$$\frac{\partial(\cdot)}{\partial \hat{x}} = \frac{\partial(\cdot)}{\partial x} \frac{\partial x}{\partial \hat{x}} = \frac{1}{1 - i\sigma_x(x)} \frac{\partial(\cdot)}{\partial x}, \quad \frac{\partial(\cdot)}{\partial \hat{y}} = \frac{\partial(\cdot)}{\partial y} \frac{\partial y}{\partial \hat{y}} = \frac{1}{1 - i\sigma_y(y)} \frac{\partial(\cdot)}{\partial y}.$$

These new expressions of the derivatives can be substituted in (2.14) and an equation valid in $\Omega \cup \Omega_{\text{PML}}$ can be defined:

$$k_0^2 u + \frac{1}{\gamma_x(x)} \frac{\partial}{\partial x} \left(\frac{1}{\gamma_x(x)} \frac{\partial u}{\partial x} \right) + \frac{1}{\gamma_y(y)} \frac{\partial}{\partial y} \left(\frac{1}{\gamma_y(y)} \frac{\partial u}{\partial y} \right) = 0, \quad (2.15)$$

where u now refers to the solution in $\Omega \cup \Omega_{\text{PML}}$ and

$$\gamma_x(x) = \begin{cases} 1 & \text{for } |x| < x_{\text{in}}, \\ 1 - i\sigma_x(x) & \text{for } |x| \geq x_{\text{in}}, \end{cases}, \quad \gamma_y(y) = \begin{cases} 1 & \text{for } |y| < y_{\text{in}}, \\ 1 - i\sigma_y(y) & \text{for } |y| \geq y_{\text{in}}. \end{cases}$$

Some limitations of the PML are discussed in [98]. The continuous PML is reflectionless when solving the exact wave equation. However, non-physical reflections appear at the interface between the PML and the computational domain because of the discretisation. To ensure negligible reflections, a solution is to use slowly varying absorption function.

On the exterior boundary, various conditions can be imposed. Adding a absorbing boundary condition can significantly reduce the non-physical reflections at Γ_{out} [97]. In practice, a standard homogeneous Dirichlet condition performs rather well, with the advantage of reducing the computational cost.

A more thorough presentation of the recent PML formulations for time-harmonic acoustics is given in [22].

2.2 Classical FEM and its Limitations

The term 'Finite Element method' was first introduced to the public by R. Williams in a paper of 1960 but the initial idea to split the numerical domain into "finite elements" came independently from A. Hrennikoff and R. Courant. In fact, the concept was born in the forties as a response to elasticity and structure problems in civil engineering [82]. In the finite element method, the governing equations are first written into a weak variational formulation to lower the constraints on the regularity of the solution. Then, the continuous domain is divided into smaller non-overlapping elements. The discrete variational formulation is constructed. In each element, the solution is to be found in a particular functional space. For linear problems, the resulting set of equations can be written in the form of a matrix system which may be solved using a direct, iterative or hybrid solver.

The method is relied upon because of its robustness. Moreover, it is easy to adapt to complex geometries and non-uniform media. However, for wave problems, the FEM suffers from dispersion error (phase shift between the numerical and exact solutions). At high frequencies, acoustic perturbations can propagate over many wavelengths, which leads to an accumulation of the dispersion error across the computational domain. This

phenomenon is called the pollution effect [14]. For instance with linear elements the numerical error scales like k^2h where k refers to the wavenumber and h to the elements size [13]. As k increases, the number of elements per wavelength must also increase to maintain the same level of accuracy, which rapidly becomes prohibitively expensive.

To circumvent this issue, a number of different strategies have been explored [153, 7]. They may roughly be classified in two categories: spectral methods and high-order polynomial FEM on the one hand and physics-based methods on the other hand.

2.3 Spectral and High-order Polynomial Methods

2.3.1 Literature Review

The high-order finite element method and spectral methods are closely related. They are built on the same idea: considering the solution as a sum of basis functions. The expansion functions of the FEM are defined locally which makes the method well-suited to handle complex geometries. This type of discretisation transforms differential equations into sparse matrix systems as only a few number of basis functions are non-zero in a given element. On the contrary, the spectral methods exploit basis functions which are infinitely differentiable global functions giving the methods good convergence properties. However, the matrices involved are dense and only simple geometries can be handled. Typically, spectral methods employ Fourier series or Tchebychev polynomials [29].

High-order finite element methods replace the standard, low-order Lagrange polynomials with higher-order functions providing superior interpolation properties. Different families of polynomials have been considered such as Bernstein, Hermite or Lobatto polynomials. Bernstein polynomials were introduced around a hundred years ago [62]. Over the years, this family of function has been used amongst others as a finite element basis to find solutions of the Laplace operator [164] or to model facial surgery using FEM [143]. Hermite polynomials are of class C^1 which means that they are differentiable and their first derivative is continuous. This makes them well-suited for fourth-order problems. These polynomials have been used in [6] to solve aero-acoustic problems in the time-domain. The authors compared a hybrid discontinuous Galerkin–Hermite method to a finite difference method. They found the former to be more efficient in terms of CPU time. Lobatto functions are obtained by integrating the Legendre polynomials. For Helmholtz problems, using the Lobatto shape functions allows a drastic reduction of the dispersion error [3, 154]. Hence, it is considered a valid approach to address the pollution effect and tackle large-scale problems [158, 21].

In [137], two high-order polynomial FEM are compared to a hybrid method. The Spectral Element Method (SEM), first proposed by Patera in 1984, was developed to combine

the advantages of the FEM and of the spectral methods. The authors show that high-order polynomial shape functions effectively control the pollution and provide more accurate solutions. SEM is found to be more efficient when regular quadrilateral elements are used but its efficiency decreases drastically when distorted or triangular elements are used. In [139], several higher-order polynomial bases, namely Bernstein, Hermite and Lobatto polynomials have been studied. The methods have been compared to a polynomial Partition of Unity Finite Element Method (PUFEM). The initial idea of PUFEM was to multiply the standard finite element method with problem dependent functions, called enrichment functions. However, in [139], polynomial enrichment functions [84, 124] which are independent of the frequency were used. This allows for faster frequency sweeps: the elementary matrices are only computed once for a given mesh and can be recalled for every frequencies. The Hermite method is found to be the most efficient but the variable-order functions are not available. Compared to Bernstein method, the Lobatto method is better conditioned. Moreover, the hierarchic nature of the Lobatto shape functions leads to efficient algorithms for solving the same model over a range of frequencies [21]. The study showed that polynomial PUFEM suffers from severe ill-conditioning which renders the method ineffective in practice.

The p -FEM approach provides exponential convergence when increasing the polynomial order p . Furthermore, p -FEM easily allows local order refinement, which makes it suited for p -adaptive and hp -adaptive strategies. These allow for a given mesh, frequency and flow properties, to automatically adjust the order of an element based on some *a priori* knowledge or *a posteriori* knowledge of the solution. For example in [139], a higher-order adaptive scheme has been proposed. The order to be used is determined using an *a priori* error estimator, which is based on solutions obtained from a one-dimensional finite element problem. The optimal order in the higher dimension element is obtained using the mean element size, the frequency and the properties of the physical medium. The variable-order p -FEM has been employed to solve a three-dimension industrial problem and significant improvements in terms of computational efficiency were reported when compared to the standard quadratic finite element method. Other approaches such as automated h -adaptivity for which the polynomial degree is kept constant in space but the local element sizes are varied and hp -adaptivity (combining both p and h refinements) generally based on *a posteriori* error estimators are also investigated [12, 131].

Hierarchical p -FEM has recently gained attention in the numerical acoustics community [21]. Therefore, in this work, the framework of p -FEM is used to construct a continuous, high-order approximation with Lobatto shape functions for solutions of the convected Helmholtz equation [159, 99].

2.3.2 Higher-order FEM

2.3.2.1 Statement of the Problem

We consider the propagation of harmonic, linear sound waves in a uniform medium within a bounded domain Ω with a boundary $\partial\Omega$. Throughout this report, we assume an implicit time dependence $e^{+i\omega t}$ where ω is the angular frequency. In the absence of volume sources, the complex amplitude u of the acoustic pressure is governed by the homogeneous convected Helmholtz equation:

$$\frac{1}{c_0^2} \frac{D_0^2 u}{Dt^2} - \nabla^2 u = 0, \quad (2.16)$$

where $D_0(\cdot)/Dt = +i\omega(\cdot) + \mathbf{v}_0 \cdot \nabla(\cdot)$ is the material derivative.

An inhomogeneous Robin condition (2.11) can be applied on Γ_r to generate a given solution u_{ex} inside the computational domain. In this case, the condition writes:

$$\nabla u \cdot \mathbf{n} + ik^+ u = g, \quad (2.17)$$

where the source term is given by $g = \nabla u_{\text{ex}} \cdot \mathbf{n} + ik^+ u_{\text{ex}}$. On the remainder of the boundary $\Gamma_d = \partial\Omega \setminus \Gamma_r$, a Dirichlet boundary condition can be imposed:

$$u = u_{\text{ex}}. \quad (2.18)$$

2.3.2.2 Weak Variational Formulation

Using the Galerkin method, the convected Helmholtz equation can be translated into a weak formulation. The aim is to relax the regularity requirements on the solution. Equation (2.16) is multiplied by a test function w associated to u . The integration of the equation over the whole domain Ω gives:

$$\int_{\Omega} \bar{w} \left[\frac{1}{c_0^2} \frac{D_0^2 u}{Dt^2} - \nabla^2 u \right] d\Omega = 0, \quad (2.19)$$

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

Using the Green theorem, we find, for $u \in \mathcal{U}$:

$$\forall w \in \mathcal{V}, \int_{\Omega} \left[-\frac{1}{c_0^2} \frac{\overline{D_0 w}}{Dt} \frac{D_0 u}{Dt} + \nabla \bar{w} \cdot \nabla u \right] d\Omega = \int_{\Gamma_r} \bar{w} \left[(\nabla u \cdot \mathbf{n}) - \frac{1}{c_0^2} (\mathbf{v}_0 \cdot \mathbf{n}) \frac{D_0 u}{Dt} \right] d\Omega, \quad (2.20)$$

where $\mathcal{U} = \{u \in H^1(\Omega), u = u_{\text{ex}} \text{ on } \Gamma_d\}$ and $\mathcal{V} = \{w \in H^1(\Omega), w = 0 \text{ on } \Gamma_d\}$ with

$$H^1(\Omega) = \left\{ f \in L^2(\Omega), \frac{\partial f}{\partial x_i} \in L^2(\Omega) \right\},$$

L^2 refers to the square integrable function space. The integral on the domain boundary is modified according to the type of specified boundary conditions.

2.3.2.3 Discretisation

The continuous domain is partitioned into non-overlapping finite elements with an average spacing between nodes h . The resulting mesh may be of different nature. It can be structured or unstructured: structured meshes are characterised by a uniform and regular grid. They are generally used when simple geometries are involved. Otherwise, unstructured meshes are preferred as they are more adapted to complex geometries. In this work, unstructured meshes are considered. If the nodes, edges and faces are perfectly matched, the mesh is said conformal. Otherwise, the mesh is non-conformal: hanging nodes or overlapping zones may be present. Only conformal meshes are examined in this work.

2.3.2.4 High-order Shape Functions

The approximate solution u^h and the associated test function w^h are constructed using the classical H^1 -conforming hierarchical polynomial shape functions. In this section, we recall the main definitions of these shape functions in two dimensions. A more detailed discussion can be found in Šolín *et al.* [159]. The interpolation of the solution is constructed on the reference triangular element defined as $\{(\zeta, \eta) \in \mathbb{R}^2; -1 < \zeta, \eta; \zeta + \eta < 0\}$. This reference element is equipped with the following affine coordinates:

$$\nu_1(\zeta, \eta) = \frac{\eta + 1}{2}, \quad \nu_2(\zeta, \eta) = -\frac{\eta + \zeta}{2}, \quad \nu_3(\zeta, \eta) = \frac{\zeta + 1}{2}. \quad (2.21)$$

The construction of the higher order shape functions rely on the Lobatto polynomials defined as:

$$l_q(x) = \sqrt{\frac{2q-1}{2}} \int_{-1}^x \mathcal{L}_{q-1}(s) ds, \quad \text{for } -1 \leq x \leq 1, \quad (2.22)$$

with \mathcal{L}_q the Legendre polynomial of order q . These functions verify the following orthogonality property:

$$\int_{-1}^1 \frac{\partial l_i}{\partial \zeta} \frac{\partial l_j}{\partial \zeta} d\zeta = 0, \quad \text{for } i \neq j, \quad (2.23)$$

which ensures an optimal conditioning of the stiffness matrix for Helmholtz problems [164]. It is also necessary to introduce the so-called kernel functions:

$$\Psi_q(x) = l_{q+2}(x) / [l_0(x)l_1(x)].$$

In each physical triangular element Ω_e , the numerical solution u^h is sought as a sum of contributions from nodal, edge and bubble shape functions (see Fig. 2.1). Nodal shape

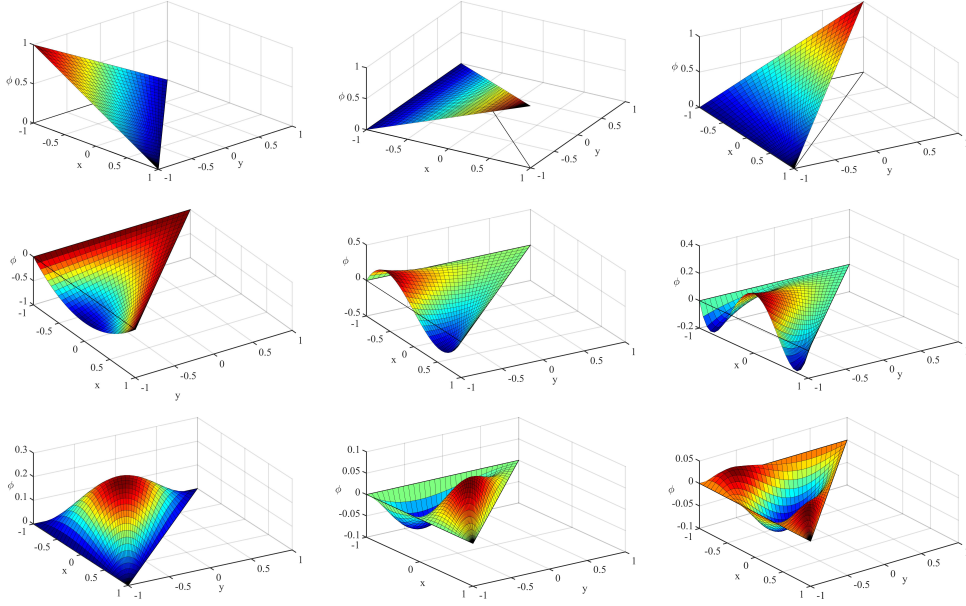


FIGURE 2.1: Example of triangular nodal, edge and bubble Lobatto shape functions.

functions are equal to unity at their node of definition and vanish at all other nodes:

$$\varphi^{v1} = \nu_2, \quad \varphi^{v2} = \nu_3, \quad \varphi^{v3} = \nu_1. \quad (2.24)$$

Edge functions are equal to the Lobatto polynomials on the edge they are associated to, and vanish on the other edges of the element:

$$\varphi_q^{e1} = \nu_2 \nu_3 \Psi_{q-2}(\nu_3 - \nu_2), \quad \varphi_q^{e2} = \nu_3 \nu_1 \Psi_{q-2}(\nu_1 - \nu_3), \quad \varphi_q^{e3} = \nu_1 \nu_2 \Psi_{q-2}(\nu_2 - \nu_1), \quad (2.25)$$

with $2 \leq q \leq p$.

Finally, the bubble shape functions vanish on the boundaries of the element but are non-zero inside the element:

$$\varphi_{q_1, q_2}^b = \nu_1 \nu_2 \nu_3 \Psi_{q_1-1}(\nu_3 - \nu_2) \Psi_{q_2-1}(\nu_2 - \nu_1), \quad (2.26)$$

with $q_1, q_2 \geq 1$ and $q_1 + q_2 \leq p - 1$.

On a given triangular element Ω_e , the discrete solution of order p is then expressed as:

$$u^h(\zeta, \eta) = \sum_{r=1}^3 \varphi^{v_r}(\zeta, \eta) u^{v_r} + \sum_{q=2}^p \sum_{r=1}^3 \varphi_q^{e_r}(\zeta, \eta) u_q^{e_r} + \sum_{q_1=2}^p \sum_{q_2=2}^{p-1-q_1} \varphi_{q_1, q_2}^b(\zeta, \eta) u_{q_1, q_2}^b, \quad (2.27)$$

where u^{v_r} , $u_q^{e_r}$ and u_{q_1, q_2}^b are the degrees of freedom (DOFs) associated to the nodal, edge and bubble functions respectively. On a given triangular element of order p , there are therefore 3 nodal functions, $3(p - 1)$ edge functions and $(p - 1)(p - 2)/2$ bubble functions.

An important property of the family of Lobatto shape functions is their hierarchic structure. The shape functions for order p are retained to construct the set of shape functions at order $p + 1$. As a consequence, the orders of the elements can easily be varied across the mesh without requiring any specific treatment which is particularly useful for p -adaptive and hp -adaptive finite-element methods.

All surface and line integrals in the p -FEM formulation (2.20) are evaluated using Gauss–Legendre quadrature [159]. For the quadratures used in the present study, all the quadratures are exact, except for the integrand involving g on Γ_r which may be non-polynomial. In this case, the order of the numerical quadrature is fixed to $2p + 2$. A large number of tests have been performed using quadrature orders up to $2p + 10$, without inducing any significant change in the results, thereby indicating that the resulting integration error can be considered negligible.

2.3.2.5 Condensation

The bubble functions defined in (2.26) vanish on the element boundaries and are therefore not coupled to the neighbouring elements. These internal degrees of freedom can be eliminated at the element level prior to assembly and do not appear in the global system matrix. This condensation procedure does not affect the final solution and is a very effective technique to reduce the size and improve the conditioning of the global system matrices [159]. The internal values of the solution are recovered during the post-processing stage by solving small linear systems for each element. The benefits of condensation tends to become more significant as the polynomial order p increases because the number of bubble functions scales like p^2 in 2D whereas the number of edge functions scales like p . In 3D, the benefit still holds as the number of bubble functions scales like p^3 . Note however that linear and quadratic triangular elements do not involve bubble functions and cannot benefit from condensation.

2.3.2.6 Error Estimates

For the hp -version of the finite element method for smooth solutions of the Helmholtz equation, Ihlenburg and Babuška [94] have obtained upper bounds for the global relative error in H^1 semi-norm

$$E \leq C_1 \left(\frac{kh}{2p} \right)^p + C_2 k \left(\frac{kh}{2p} \right)^{2p},$$

where C_1 and C_2 are independent of k and h and are weak functions of the order p . The first term in this expression represents the interpolation error and has a dependency of order p with kh . This error can be controlled by using a sufficient number of elements per wavelength, given by $2\pi/(kh)$. The second term is associated with the dispersion

error and the pollution effect. The dispersion error is the difference between the theoretical wavenumber and the wavenumber actually observed in the numerical model. From this error estimate, the advantages of resorting to higher order shape functions appears clearly for smooth solutions. Exponential convergence rates are obtained with p -refinement, while convergence with h -refinement is only algebraic.

At high resolution (*i.e.* at low values of kh), the interpolation error dominates and the H^1 -norm error converges asymptotically as $\mathcal{O}(kh/2p)^p$. In the L^2 -norm, the actual rate of convergence of the interpolation error is increased by one level and follows $\mathcal{O}(kh/2p)^{p+1}$ [10]. These theoretical error estimates are found to match closely with numerical experiments, see for instance [20].

2.4 Physics-based Methods

2.4.1 Literature Review

Physics-based methods, also referred to as Trefftz methods, are another category of method developed to remedy the pollution problem (a detailed review of these methods can be found in [87]). The idea is to incorporate *a priori* knowledge about the local behaviour of the solution into the numerical model. This is generally achieved by using local canonical solutions of the governing equations to build an approximation basis.

For the Helmholtz equation, the local solutions are generally plane waves, but other alternatives can be used, such as Bessel functions [110] or Green's functions [33]. In [110], the authors compared the accuracy and costs of using plane waves or Bessel functions in the basis. For the propagating and evanescent plane wave in free field problem, the results showed that to achieve high accuracy, the Bessel basis is more robust. When the accuracy needed is around 1%, the use of a plane-wave basis is preferred especially because of the expensive quadratures in the case of a Bessel basis. For a singular problem, they found that both bases performed similarly for uniform meshes. To obtain higher accuracies, the plane waves and Bessel basis can be coupled or the plane wave basis can be adopted on a non-uniform mesh (refined around the singularity points).

It is generally assumed that a plane-wave basis naturally provides better accuracy compared to polynomials since the latter bear no relation with the governing equations. A number of wave-based methods rely on discontinuous formulations where the solution is approximated with plane waves in each elements and the continuity between elements is imposed weakly. This includes the Ultra Weak Variational Formulation (UWVF) [34, 93, 35], the least square method [125], and the wave-based discontinuous Galerkin method [66]. They differ in the way the inter-element continuity is formulated but it is now recognised that these methods are all variants of the same discontinuous Galerkin method using different numerical fluxes to ensure continuity between elements [71]. The

DGM framework is particularly well suited for adaptivity since it is straightforward to change the number of plane waves in each element. Applied to the Helmholtz equation, the methods are found to be more efficient than the standard FEM [125, 35].

Another method using discontinuous interpolation is the Discontinuous Enrichment Method (DEM) where enrichment functions are added to the set of conventional shape functions to construct the approximate solution in each element [53]. The standard FE basis represents the coarse scales of the solution whereas the enrichment functions represent the fine scales. The resulting basis is not continuous across the elements interfaces and the continuity of the solution is weakly enforced by the mean of Lagrange multipliers on the edges between elements. The enrichment functions only contribute to the solution locally and have no connectivity with neighbouring elements. Static condensation can therefore be applied. This method has been used to solve the Helmholtz equation [54] and extended to other fields such as fluid-structure interactions [120]. The DEM leads to a significant reduction of the number of degrees of freedom when compared to linear, quadratic and quartic Lagrange finite elements for a fixed accuracy in 2D [54, 61] and 3D [150]. Removing the conventional shape functions from the DEM leads to a wave-based discontinuous Galerkin method with Lagrange multipliers [54]. It was indeed shown that the polynomial field is not crucial to efficiently solve the Helmholtz equation.

Other Trefftz methods relying on discontinuous interpolations include the wave-based method [43] and the variational theory of complex rays (VTCR) [103], but these two methods cannot be formulated in the DG framework [87]. Results in [90] indicate that the UWVF tends to outperform the VTCR for simple Helmholtz problems.

The PUFEM constructs an approximation of the solution by multiplying enrichments functions with the standard FEM shape functions [122]. The benefit of this approach is that the resulting approximation space is naturally continuous. The downside is the calculation of the element matrices which can be costly because it involves highly-oscillatory integrals with products of exponentials and polynomials. Astley and Gamallo were the first to apply the method in the aero-acoustic context [8] to model the propagation of acoustic waves on potential mean flows. The standard basis functions were enriched with plane waves solutions of the convected Helmholtz equation. They show that even though the conditioning of the method deteriorates when the number of wave directions increases, accurate solutions can be obtained with reasonable condition numbers. Compared to conventional FEM, PUFEM yields very large improvements. Elements of comparisons of PUFEM with the UWVF can be found in [92, 160] and the UWVF is relatively better conditioned and generally more efficient.

2.4.2 Wave-based Discontinuous Galerkin Method

This section describes the formulation of the wave-based DGM [66] for solving the convected Helmholtz equation.

2.4.2.1 Statement of the Problem

The two dimensional convected Helmholtz equation (2.16) is first reformulated as a set of linear, first-order equations of the form:

$$+ i\omega \mathbf{u} + \frac{\partial}{\partial x}(\mathbf{A}\mathbf{u}) + \frac{\partial}{\partial y}(\mathbf{B}\mathbf{u}) = \mathbf{0} , \quad (2.28)$$

by defining

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho_0 v_x \\ \rho_0 v_y \end{bmatrix} , \quad \mathbf{A} = \begin{bmatrix} 0 & 1 & 0 \\ c_0^2 - v_{x,0}^2 & 2v_{x,0} & 0 \\ -v_{x,0}v_{y,0} & v_{y,0} & v_{x,0} \end{bmatrix} , \quad \mathbf{B} = \begin{bmatrix} 0 & 0 & 1 \\ -v_{x,0}v_{y,0} & v_{y,0} & v_{x,0} \\ c_0^2 - v_{x,0}^2 & 0 & 2v_{x,0} \end{bmatrix} , \quad (2.29)$$

where ρ_0 is the mean density of the fluid. The unknowns ρ , $\rho_0 v_x$ and $\rho_0 v_y$ are the linear perturbations of density, x -momentum and y -momentum respectively. The numerical method is described here for constant coefficient matrices but it could be extended to the case where the matrices are functions of x and y [66] by approximating them as piecewise constant functions.

2.4.2.2 Variational Formulation

As for the p -FEM presented above, the computational domain Ω is divided into a collection of elements Ω_e but unlike p -FEM the solution is allowed to be discontinuous across the element edges. After integration by parts in each element the variational formulation for equation (2.28) reads:

$$\sum_e \int_{\Omega_e} \left(i\omega \mathbf{w}^T \mathbf{u} - \frac{\partial \mathbf{w}^T}{\partial x} \mathbf{A} \mathbf{u} - \frac{\partial \mathbf{w}^T}{\partial y} \mathbf{B} \mathbf{u} \right) d\Omega + \sum_e \int_{\partial\Omega_e} \left(\mathbf{w}^T \mathbf{A} \mathbf{u} n_x + \mathbf{w}^T \mathbf{B} \mathbf{u} n_y \right) d\Gamma = 0 , \quad (2.30)$$

where T denotes the conjugate transpose, \mathbf{w} is the test function associated to the solution \mathbf{u} and $\mathbf{n}_e = (n_x, n_y)$ is the outward unit normal to $\partial\Omega_e$.

After introducing the flux matrix $\mathbf{F} = \mathbf{A}n_x + \mathbf{B}n_y$ and rearranging terms in the first integral, equation (2.30) becomes:

$$\begin{aligned} - \sum_e \int_{\Omega_e} \left(i\omega \mathbf{w} + \mathbf{A}^T \frac{\partial \mathbf{w}}{\partial x} + \mathbf{B}^T \frac{\partial \mathbf{w}}{\partial y} \right)^T \mathbf{u} \, d\Omega + \int_{\partial\Omega} \mathbf{w}^T \mathbf{F} \mathbf{u} \, d\Gamma \\ + \sum_e \sum_{e' < e} \int_{\Gamma_{e,e'}} (\mathbf{w}^T \mathbf{F} \mathbf{u})_e + (\mathbf{w}^T \mathbf{F} \mathbf{u})_{e'} \, d\Gamma = 0, \end{aligned} \quad (2.31)$$

where $\Gamma_{e,e'}$ is the edge between the element e and e' . The continuity of the normal flux across the interface between the two elements is directly enforced by writing:

$$\mathbf{F}_e \mathbf{u}_e = -\mathbf{F}_{e'} \mathbf{u}_{e'} = \mathbf{f}_{e,e'}(\mathbf{u}_e, \mathbf{u}_{e'}) , \quad (2.32)$$

where $\mathbf{f}_{e,e'}$ is the so-called numerical flux which is discussed below. In the wave-based DGM, the solution in each element is taken to be an exact solution of the homogeneous equation:

$$i\omega \mathbf{u} + \mathbf{A} \frac{\partial \mathbf{u}}{\partial x} + \mathbf{B} \frac{\partial \mathbf{u}}{\partial y} = \mathbf{0} . \quad (2.33)$$

In addition, the test function \mathbf{w} in each element is a solution of the adjoint equation which can be easily identified from (2.31):

$$-i\omega \mathbf{w} - \mathbf{A}^T \frac{\partial \mathbf{w}}{\partial x} - \mathbf{B}^T \frac{\partial \mathbf{w}}{\partial y} = \mathbf{0} . \quad (2.34)$$

This particular choice of test function is crucial as it allows to remove the element integrals from the formulation and the variational formulation finally reads:

$$\sum_e \sum_{e' < e} \int_{\Gamma_{e,e'}} (\mathbf{w}_e - \mathbf{w}_{e'})^T \mathbf{f}_{e,e'}(\mathbf{u}_e, \mathbf{u}_{e'}) \, d\Gamma + \int_{\partial\Omega} \mathbf{w}^T \mathbf{F} \mathbf{u} \, d\Gamma = 0 . \quad (2.35)$$

2.4.2.3 Numerical Flux

The efficiency of a discontinuous Galerkin method depends to a large extent on the choice of the numerical flux. In the current formulation, the upwind flux-vector splitting method is used. The flux on the edge between two elements is expressed in terms of characteristic waves in the direction normal to the edge [161]. These characteristic waves can be obtained by diagonalising the flux matrix $\mathbf{F} = \mathbf{W} \mathbf{\Lambda} \mathbf{W}^{-1}$. The matrix of eigenvalues $\mathbf{\Lambda}$ contains the phase velocity of the characteristic waves and \mathbf{W} is the matrix of eigenvectors. They are used to split the flux matrix \mathbf{F} into the incoming and outgoing flux terms by writing

$$\mathbf{F} = \mathbf{F}^+ + \mathbf{F}^- , \quad \text{with } \mathbf{F}^\pm = \mathbf{W} \mathbf{\Lambda}^\pm \mathbf{W}^{-1}$$

where $\mathbf{\Lambda}^\pm$ is the diagonal matrix of the positive or negative eigenvalues. For a given edge, the incoming flux is calculated using the inner solution while the outgoing flux is

calculated using the outer solution, resulting in the following expression for the numerical:

$$\mathbf{f}_{e,e'}(\mathbf{u}_e, \mathbf{u}_{e'}) = \mathbf{F}^+ \mathbf{u}_e + \mathbf{F}^- \mathbf{u}_{e'} , \quad (2.36)$$

where the first term correspond to the flux associated to waves travelling from e to e' and the second term to the flux associated to waves travelling the other way. The variational formulation reads:

$$\sum_e \sum_{e' < e} \int_{\Gamma_{e,e'}} (\mathbf{w}_e - \mathbf{w}_{e'})^T (\mathbf{F}^+ \mathbf{u}_e + \mathbf{F}^- \mathbf{u}_{e'}) d\Gamma + \int_{\partial\Omega} \mathbf{w}^T \mathbf{F} \mathbf{u} d\Gamma = 0 . \quad (2.37)$$

In the case of the convected Helmholtz equation formulated with (2.28–2.29) the computation of $\mathbf{\Lambda}$ and \mathbf{W} gives:

$$\mathbf{\Lambda} = \begin{bmatrix} \mathbf{v}_0 \cdot \mathbf{n} & 0 & 0 \\ 0 & \mathbf{v}_0 \cdot \mathbf{n} - c_0 & 0 \\ 0 & 0 & \mathbf{v}_0 \cdot \mathbf{n} + c_0 \end{bmatrix} , \quad (2.38)$$

$$\mathbf{W} = \begin{bmatrix} 0 & 1 & 1 \\ -c_0 n_y & v_{x,0} - c_0 n_x & v_{x,0} + c_0 n_x \\ c_0 n_x & v_{y,0} - c_0 n_y & v_{y,0} + c_0 n_y \end{bmatrix} , \mathbf{W}^{-1} = \frac{1}{2c_0} \begin{bmatrix} 2(v_{x,0} n_y - v_{y,0} n_x) & -2n_y & 2n_x \\ c_0 + \mathbf{v}_0 \cdot \mathbf{n} & -n_x & -n_y \\ c_0 - \mathbf{v}_0 \cdot \mathbf{n} & n_x & n_y \end{bmatrix} . \quad (2.39)$$

2.4.2.4 Boundary Conditions

A variety of boundary conditions can be formulated within the wave-based DGM and a systematic approach for this is presented in [70]. In the present work, ghost cells and Dirichlet boundary conditions are used.

Ghost cells can be used to generate a given solution within the computational domain. This technique follows naturally from the numerical flux presented above and the idea is to treat the outer boundary of the computational domain as an internal edge between two elements. The same flux splitting as in equation (2.36) is used to write

$$\mathbf{F} \mathbf{u} = \mathbf{F}^+ \mathbf{u} + \mathbf{F}^- \mathbf{u} = \mathbf{F}^+ \mathbf{u} + \mathbf{F}^- \mathbf{g}_r .$$

The key difference with an internal edge is that the outer solution is known and will be denoted \mathbf{g}_r . The incoming flux is therefore expressed in terms of this prescribed solution by writing $\mathbf{F}^- \mathbf{u} = \mathbf{F}^- \mathbf{g}_r$. It should be noted that this approach is equivalent to the Robin boundary condition (2.11) used for the p -FEM. The forcing term \mathbf{g}_r for the wave-based DGM can in fact be expressed in terms of the target solution u_{ex} for the

p -FEM:

$$\mathbf{g}_r = -\frac{1}{i\omega} \begin{bmatrix} -i\omega u_{\text{ex}}/c_0^2 \\ \partial u_{\text{ex}}/\partial x \\ \partial u_{\text{ex}}/\partial y \end{bmatrix}. \quad (2.40)$$

For the Dirichlet boundary condition (2.18), the method outlined in [70] is used to modify the variational formulation accordingly. Imposing the pressure on the boundary amounts to specifying a relation between the outgoing and incoming characteristic waves. The flux term on the boundary can then be written $\mathbf{F}\mathbf{u} = \mathbf{F}_d\mathbf{u} + \mathbf{g}_d$. For conciseness, the definitions are given in the no flow case [70]:

$$\mathbf{F}_d = \begin{bmatrix} c_0 & n_x & n_y \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{g}_d = \begin{bmatrix} -u_{\text{ex}}/c_0 \\ u_{\text{ex}}n_x \\ u_{\text{ex}}n_y \end{bmatrix}. \quad (2.41)$$

The variational formulation finally becomes:

$$\begin{aligned} \sum_e \sum_{e' < e} \int_{\Gamma_{e,e'}} (\mathbf{w}_e - \mathbf{w}_{e'})^T (\mathbf{F}^+ \mathbf{u}_e + \mathbf{F}^- \mathbf{u}_{e'}) d\Gamma + \int_{\Gamma_r} \mathbf{w}^T \mathbf{F}^+ \mathbf{u} d\Gamma \\ + \int_{\Gamma_d} \mathbf{w}^T \mathbf{F}_d \mathbf{u} d\Gamma = - \int_{\Gamma_r} \mathbf{w}^T \mathbf{F}^- \mathbf{g}_r d\Gamma - \int_{\Gamma_d} \mathbf{w}^T \mathbf{g}_d d\Gamma. \end{aligned} \quad (2.42)$$

2.4.2.5 Interpolation Basis

With Trefftz methods the solution is approximated using local, canonical solutions of the governing equations. Various choices of approximating functions are available: propagating plane waves [122, 34, 125, 53], evanescent plane waves [120, 151, 111], Bessel functions [110], Green's functions [33]. Some elements of comparisons of these methods can be found in [73, 92, 68, 71, 160].

In the wave-based DGM the approximate solution \mathbf{u}_e^h and test function \mathbf{w}_e^h in each element Ω_e are constructed as sums of propagating plane waves:

$$\mathbf{u}_e^h(\mathbf{x}) = \sum_{n=1}^{N_w} a_n^e \mathbf{r}_n \exp(-ik_n \boldsymbol{\theta}_n \cdot \mathbf{x}), \quad \mathbf{w}_e^h(\mathbf{x}) = \sum_{m=1}^{N_w} b_m^e \mathbf{l}_m \exp(-ik_m \boldsymbol{\theta}_m \cdot \mathbf{x}). \quad (2.43)$$

With this approximation the degrees of freedom in each element are the wave amplitudes a_n^e and b_m^e as all the other quantities in (2.43) are fully defined.

The unit vectors $\boldsymbol{\theta}_n$ represent the directions of the plane waves. In the framework of the DGM, the number and the orientations of the plane waves can easily be chosen independently in each element. For simplicity, in this study the same number N_w of plane waves is used in every element. Unless otherwise stated, the wave directions are

evenly distributed on the interval $[0, 2\pi[$. In the literature, both theoretical results and numerical evidence indicate that $N_w = 2n$ and $N_w = 2n - 1$ (where $n \in \mathbb{N}^*$) provide the same order of convergence, see for instance [79, 78]. This implies that using an odd number of plane waves offers a gain in performance [93, 85]. Therefore, for the sake of conciseness, only odd numbers of plane waves will be shown in this study, with N_w varying from 3 to 25. Results for even numbers of plane waves have also been calculated and the performance and conditioning of the wave-based DGM was indeed found to be slightly below those obtained with odd values of N_w .

The formulation of the wave-based DGM requires that the solution in each element is an exact solution of equation (2.33). It follows that the wavenumbers k_n and vector amplitudes \mathbf{r}_n introduced in (2.43) have to be solution of the following eigenvalue problem:

$$(\mathbf{A} \cos \theta_n + \mathbf{B} \sin \theta_n) \mathbf{r}_n = \frac{\omega}{k_n} \mathbf{r}_n. \quad (2.44)$$

This equation represents the dispersion relation of the waves propagating in the θ_n direction.

Similarly, the test function \mathbf{w}^h has to satisfy the adjoint equation (2.34) and it follows that the wavenumbers k_n and vector amplitudes \mathbf{l}_n are solution of the adjoint dispersion relation:

$$(\mathbf{A} \cos \theta_n + \mathbf{B} \sin \theta_n)^T \mathbf{l}_n = \frac{\omega}{k_n} \mathbf{l}_n. \quad (2.45)$$

In the case of the convected Helmholtz equation formulated with (2.28–2.29), for a given direction θ_n , one finds:

$$k_{n,1} = \frac{\omega}{\mathbf{v}_0 \cdot \boldsymbol{\theta}_n}, \quad k_{n,2} = \frac{\omega}{\mathbf{v}_0 \cdot \boldsymbol{\theta}_n - c_0}, \quad k_{n,3} = \frac{\omega}{\mathbf{v}_0 \cdot \boldsymbol{\theta}_n + c_0}, \quad (2.46)$$

and the corresponding wave amplitudes:

$$\mathbf{r}_1 = \begin{bmatrix} 0 \\ -c_0 \sin \theta_n \\ c_0 \cos \theta_n \end{bmatrix}, \quad \mathbf{r}_2 = \begin{bmatrix} 1 \\ v_{x,0} - c_0 \cos \theta_n \\ v_{y,0} - c_0 \sin \theta_n \end{bmatrix}, \quad \mathbf{r}_3 = \begin{bmatrix} 1 \\ v_{x,0} + c_0 \cos \theta_n \\ v_{y,0} + c_0 \sin \theta_n \end{bmatrix}, \quad (2.47)$$

$$\mathbf{l}_1 = \begin{bmatrix} v_{x,0} \sin \theta_n - v_{y,0} \cos \theta_n \\ -\sin \theta_n \\ \cos \theta_n \end{bmatrix}, \quad \mathbf{l}_2 = \begin{bmatrix} -c_0 - \mathbf{v}_0 \cdot \boldsymbol{\theta}_n \\ \cos \theta_n \\ \sin \theta_n \end{bmatrix}, \quad \mathbf{l}_3 = \begin{bmatrix} c_0 - \mathbf{v}_0 \cdot \boldsymbol{\theta}_n \\ \cos \theta_n \\ \sin \theta_n \end{bmatrix}. \quad (2.48)$$

The first eigenvectors correspond to vorticity waves convected with the mean flow. They do not affect the pressure perturbation. The two other vectors correspond to acoustic waves propagating in opposite directions and produce a potential velocity field. As we are interested in potential problems only, where no vorticity waves are involved, the number of vorticity waves in our basis will be zero.

The wavenumbers k_n and the vector amplitudes \mathbf{r}_n and \mathbf{l}_n are exact solutions of the dispersion relations (2.44–2.45). By construction the numerical model will directly include these properties.

It can be noted that using this particular wave basis, the variational formulation (2.42) only involves integrals of exponentials on the edges between elements which are relatively inexpensive to compute. This is in contrast with the PUFEM, for instance, which involves costly integrals of polynomial-exponential products [122]. The integrals on the boundaries Γ_d and Γ_r involving the source terms are evaluated in closed form if possible, or using high-order Gauss-Legendre quadrature.




FIGURE 2.2: Plane waves discretisation in a triangle.

2.4.2.6 Error Estimates

Theoretical error estimates for the h -version of the UWVF are given in [30], for the Helmholtz equation. It is shown that for smooth solutions, the global discretization error measured in the L^2 -norm behaves asymptotically like $h^{(N_w-1)/2-1}$. However numerical tests carried out for different frequencies and number of plane waves indicate that this expression under estimates the actual convergence rates. The convergence rates obtained from numerical experiments reported in [30] are found to be close to $h^{[(N_w+1)/2]}$, where $[x]$ denotes the integer part of x . Similar error estimates are also derived for h -refinement by Cessenat *et al.* in [34], although for a slightly different error norm.

In a series of papers, Hiptmair *et al.* [85, 86, 88] discussed the behavior of the p and the hp -version of the wave-based DGM. The main result that, for a fixed mesh, the L^2 -norm of the discretisation error converges exponentially in the square root of the number of degrees of freedom, including in the presence of strong mesh refinements [88]. Concerning the dispersion error, likely to dominate in the lower resolution regime, it is analyzed for both triangular and quadrilateral meshes in 2D in [79]. The relative dispersion error is shown to behave asymptotically like $(kh)^{(N_w-1)}$ for large values of kh , hence exhibiting an exponential decrease as the number of plane waves is increased.

Chapter 3

Performance Study of p -FEM and Wave-Based DGM

In this chapter, the performance of the p -FEM and the wave-based DGM are benchmarked for the two-dimensional Helmholtz and convected Helmholtz equations. Unless otherwise stated, the problems considered are no-flow situations. The methodology put in place for the comparison is first introduced. This includes the definition of the used cases, as well as the choice of the measures for the relative cost and the accuracy. The comparison begins by examining the interpolation properties of the two approximation bases. The convergence, conditioning and computational costs of the two methods are then investigated in details. The analysis ends with some practical considerations on the strengths and weaknesses of each approach, and outlines possible improvement paths.

3.1 Description of the Test Cases

The performance of high-order methods is known to be highly problem dependent [87]. Therefore, a total of four cases are introduced, in order to assess the ability of the two methods to tackle a large variety of problems. In all computations, all quantities are non-dimensionalised, the speed of sound c_0 and the density ρ_0 are set to unity. Examples of solutions for each of the four problems are shown in figure 3.1.

3.1.1 Plane Wave in Free Field Problem

The first test case is a simple plane wave propagating on a square domain (see Fig. 3.1(a)). This rudimentary scenario allows for a detailed analysis of the dispersion properties of the numerical models such as the anisotropy. It has been used in previous

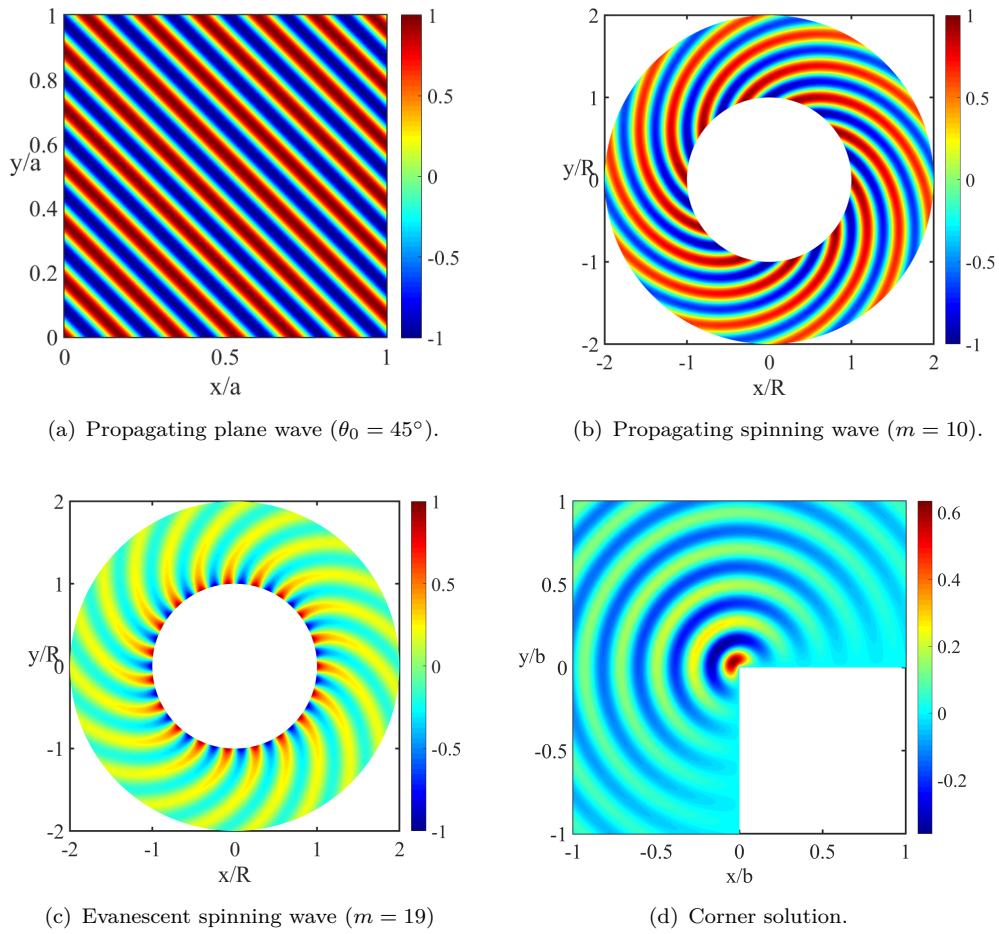


FIGURE 3.1: Examples of analytical solutions of the four benchmark problems ($kL = 50$).




FIGURE 3.2: Examples of unstructured meshes used for the test cases.

studies to analyse the performance of the wave-based DGM [66], the Partition of Unity method [122] and the Discontinuous Galerkin method with Lagrange multipliers [54].

The computational domain is a square of size a and is represented by an unstructured mesh of triangular elements of typical size h (see Fig.3.2(a)). A single propagating plane wave with direction θ_0 is generated inside the domain:

$$u_{\text{ex}} = e^{-ik\theta_0 \cdot \mathbf{x}} ,$$

where $\theta_0 = (\cos \theta_0, \sin \theta_0)$. For the p -FEM this is achieved using the Robin condition (2.11) and for the wave-based DGM the ghost cells are used with (2.40).

While providing insight into the accuracy of the numerical schemes, this plane wave solution is not representative of realistic applications where the sound fields typically contain a wide range of wave directions. Realistic problems can also involve evanescent waves in addition to propagating waves, and these can be difficult to capture by some numerical models [92, 110]. This motivates the introduction of a second and third test cases.

3.1.2 Spinning Wave Problems

The second and third test cases consist in the simulation of a single cylindrical harmonic of order m which will be referred to as the spinning wave problems. Examples of such solutions are shown in figures 3.1(b) and 3.1(c) and it can be seen that all wave directions are equally represented in these problems. The analytical solution is given in cylindrical coordinates:

$$u_{\text{ex}}(r, \theta) = H_m^{(2)}(kr)e^{-im\theta} , \quad (3.1)$$

where $H_m^{(2)}$ is the Hankel function of the second kind of order m which determines the number of spiral waves around the circumference.

These test cases are solved on an annular computational domain with inner and outer radii R and $2R$, respectively. The exact spinning waves (3.1) are enforced on the inner and on the outer boundary using the Robin condition (2.20) for p -FEM or the ghost cell formulation with (2.40) for the wave-based DGM. An example of unstructured triangular mesh used is shown in figure 3.2(b). In order to avoid introducing any additional source of error in the model, the circular geometry is represented exactly. For p -FEM, this is achieved by using a linear coordinate mapping between the reference element and the cylindrical coordinates r and θ (instead of Cartesian coordinates x and y). For the wave-based method the edges of the elements on the boundaries are defined as arcs rather than straight lines. For both methods, this leads to an exact representation of the cylindrical boundaries of the domain.

It is worth noting that this situation is closely related to the well-known problem of a plane wave scattered by a cylinder since the solution for the latter is a sum of cylindrical harmonics, see for instance [160, 70]. Although this problem has been investigated, it is preferable to consider different cylindrical harmonics individually (the reader can find detailed justifications in Annexe A). For $m < kR$ the sound field decays slowly like $1/\sqrt{r}$ away from the inner boundary. However for $m > kR$ there is a region near the inner boundary where the sound field decays at a much faster rate $1/r^m$. These harmonics represent creeping waves that propagate along smooth convex surfaces (see section 9-5 in [138]). While creeping waves do radiate sound to the far field, they play the same role as evanescent waves radiating from a flat surface. As will be shown below, the large gradients present in the solutions for $m > kR$ lead to a change in the performance of the numerical models. For this reason both propagating (Fig. 3.1(b)) and evanescent (Fig. 3.1(c)) spinning wave harmonics will be considered.

3.1.3 Corner Problem

All the benchmark problems described above involve smooth solutions, but the performance of high-order schemes is known to deteriorate significantly when dealing with non-smooth solutions, including the p -FEM [11] and the wave-based DGM [92]. Therefore we also consider the simulation of the corner solution, shown in figure 3.1(d), which was also used in [92, 158]. The analytical expression for the pressure field is given by

$$u_{\text{ex}}(r, \theta) = J_{2/3}(kr) \sin\left(\frac{2}{3}\theta\right), \quad \text{for } 0 < \theta < \frac{3}{2}\pi, \quad (3.2)$$

where $J_{2/3}$ is the Bessel function of the first kind. While u_{ex} itself remains bounded, its gradient is singular at the corner. This problem will allow to illustrate how the two methods cope with non-smooth solutions.

As shown in figure 3.1(d) an L-shaped computational domain is used for this test case. On the two edges of length b which meet at the origin, an homogeneous Dirichlet condition is enforced: $u = 0$. This is done using the Dirichlet boundary condition (2.18) for p -FEM and the formulation (2.42) with (2.41) for wave-based DGM.

For all the other boundaries the Robin condition (2.11) is used for p -FEM and the ghost cells (2.40) are used for the wave-based DGM. Unstructured triangular meshes will be used and an example is shown in figure 3.2(c). For p -FEM it is recommended to use local h -refinement to improve the accuracy of the model near the singular part of the solution, so both a uniform and refined grid will be considered here. For the wave-based DGM incorporating *ad-hoc* function in the approximation basis is generally required in such situations. These issues will be addressed in more details when discussing the results for this test case in sections 3.3 and 3.4.

3.2 Measures of Accuracy and Costs

For all the test cases described above a number of parameters can be varied: the angular frequency ω , the element size h and the approximation order (that is the polynomial order p for the p -FEM and the number of plane waves N_w for the wave-based DGM). Since the sizes and shapes of the computational domains vary between the different test cases it is more representative to define the frequency in terms of the Helmholtz number kL . The length $L = \sqrt{\text{area}(\Omega)}$ is indicative of the size of the domain and different test cases solved with the same Helmholtz number kL can be expected to be as demanding.

To provide a measure of the numerical resolution, we define the number of degrees of freedom per wavelength as follows [66]:

$$D_\lambda = \frac{2\pi}{kL} \left(\sqrt{N_{\text{DOF}}} - 1 \right) . \quad (3.3)$$

The number of elements per wavelength could also be used to describe the resulting problem but when dealing with high-order methods, it is more relevant to use D_λ which can be understood as the density of information describing one wavelength. Therefore, unless stated otherwise, the mesh resolution will be measured using D_λ .

The accuracy of the numerical solutions is assessed by calculating the relative L^2 -error

$$E_{L^2(\Omega)} = \frac{\|u^h - u_{\text{ex}}\|_{L^2(\Omega)}}{\|u_{\text{ex}}\|_{L^2(\Omega)}} , \quad \text{with } \|\cdot\|_{L^2(\Omega)} = \left(\int_{\Omega} |\cdot|^2 d\Omega \right)^{1/2} . \quad (3.4)$$

This integral is evaluated using high-order Gauss-Legendre quadrature rules on each element with $2p + 10$ points for p -FEM and $N_w + 10$ for the wave-based DGM.

Once assembled, both p -FEM and wave-based DGM lead to a sparse global system of equations with complex-valued coefficients of the form:

$$\mathbb{K}\mathbf{u} = \mathbf{b} . \quad (3.5)$$

The latter can be solved either using a direct or an iterative approach.

In this first part of the study, only direct solution procedures are considered. The development of robust iterative solvers for Helmholtz problems discretised by the standard FEM is undergoing active research (for a review see [133]). However, the indefiniteness of the operator tends to hinder their performance at higher frequencies. Furthermore, the system matrices arising from a high-order discretisation typically exhibit higher condition numbers than that of the linear FEM ones, which renders the problem even more challenging. For these reasons, iterative solutions will not be considered here. Instead, all systems are solved using either the Cholesky (in the no-flow case, when \mathbb{K} is symmetric) or LU factorisation algorithm of the multi-frontal sparse direct solver MUMPS (version 4.10.0) [1].

This operation represents the most expensive part of the solution procedure. The cost of solving the global system is monitored in a number of ways:

- The total number of degrees of freedom N_{DOF} is a first indicator, but in practice it does not necessarily correlate directly with the cost of a numerical model, see for instance [21].
- The number of non-zero entries N_{NZ} in the sparse global matrix \mathbb{K} is another measure of cost. In particular the memory required to store the matrix before the factorisation is directly proportional to N_{NZ} .
- The factorisation memory, as reported by the solver MUMPS, is a crucial parameter because it is generally the bottleneck in solution procedures based on direct solvers. In real-world applications it is often the amount of memory available that dictates the size of the problems that can be solved on a given platform.
- Although this is more dependent on the platform used for the calculations, the time taken by the factorisation will also be reported.
- Finally, it is important to monitor the condition number κ of the system:

$$\kappa = \|\mathbb{K}\|_1 \left\| \mathbb{K}^{-1} \right\|_1$$

where $\|(\cdot)\|_1$ refers to the 1-norm. The conditioning of a system quantifies the sensitivity of the solution vector \mathbf{u} to small changes in the inputs. In other words, it is a measure of the effect of a small change in the problem parameters on the resulting solution. These little changes may correspond for example to roundoff errors during the numerical process. This quantity is important as the accuracy, or even success, of the direct solver can be affected by poor conditioning. In addition it is often reported in the literature that wave-based methods are poorly conditioned and this aspect will be carefully examined in the next sections.

Note that for p -FEM, when condensation is applied, N_{DOF} and N_{NZ} do not take into account the degrees of freedom corresponding to the internal bubble functions, which are eliminated before the global system is assembled. Unless otherwise stated, the cost of solving the linear system will be reported with condensation. However, its effects on computational cost and conditioning will also be addressed.

3.3 Comparison of the Interpolation Properties

Before trying to solve the test cases previously defined using the two methods, a first analysis is performed to assess the interpolation properties of their respective approximation bases. Independently of the numerical formulations, we compare in this section the ability of the continuous polynomial basis on the one hand, and the discontinuous plane-wave basis on the other hand to approximate the solutions of the four different test cases.

This is done by computing the best interpolation $u^{h,\text{opt}}$ that minimises the L^2 -error relative to a known exact solution u_{ex} where the best interpolation is given by:

$$u^{h,\text{opt}} = \sum_n a_n \varphi_n$$

where φ_n the n^{th} shape functions. Finding the degrees of freedom a_n which minimise the L^2 error involves the following cost function \mathfrak{J} :

$$\begin{aligned} \mathfrak{J} &= \int_{\Omega} |u_{\text{ex}} - u^{h,\text{opt}}|^2 d\Omega \\ &= \int_{\Omega} (u_{\text{ex}} - \sum_n a_n \varphi_n)^T (u_{\text{ex}} - \sum_n a_n \varphi_n) d\Omega \end{aligned} \quad (3.6)$$

The expression can be further developed:

$$\mathfrak{J} = \mathbf{a}^T \mathbb{M} \mathbf{a} + \mathbf{a}^T \mathbf{y} + \mathbf{a} \mathbf{y}^T + z, \quad (3.7)$$

where

$$(\mathbb{M})_{ij} = \int_{\Omega} \overline{\varphi_i} \varphi_j d\Omega, \quad y_i = - \int_{\Omega} \overline{\varphi_i} u_{\text{ex}} d\Omega, \quad z = \int_{\Omega} |u_{\text{ex}}|^2 d\Omega,$$

and where \mathbf{a} refers to the vector of DOFs a_n . Decomposing every terms into their real and imaginary parts, $(\cdot)_r$ and $(\cdot)_i$ respectively, leads to:

$$\mathfrak{J} = \mathbf{a}_r^T \mathbb{M}_r \mathbf{a}_r - 2\mathbf{a}_r \mathbb{M}_i \mathbf{a}_i + \mathbf{a}_i^T \mathbb{M}_r \mathbf{a}_i + 2\mathbf{a}_r^T y_r + 2\mathbf{a}_i^T y_i + z. \quad (3.8)$$

Minimising the cost function \mathfrak{J} is equivalent to solving the system:

$$\begin{cases} \mathbb{M}_r \mathbf{a}_r - \mathbb{M}_i \mathbf{a}_i = -\mathbf{y}_r \\ \mathbb{M}_i \mathbf{a}_r + \mathbb{M}_r \mathbf{a}_i = -\mathbf{y}_i \end{cases} \quad (3.9)$$

Finding the optimal interpolation involves a simple least-square problem which amounts to solving the linear system $\mathbb{M} \mathbf{a} = -\mathbf{y}$.

The relative error on the optimal interpolation is denoted by

$$E_{L^2(\Omega)}^{\text{opt}} = \frac{\|u^{h,\text{opt}} - u_{\text{ex}}\|_{L^2(\Omega)}}{\|u_{\text{ex}}\|_{L^2(\Omega)}}. \quad (3.10)$$

It is worth noting that the results for the wave-based DGM presented in this section also apply to a number of other wave-based methods: UWVF, least-square method and wave-based DGM with Lagrange multipliers since all of these methods interpolate the solution with the same plane-wave basis defined in section 2.4.2.5.




FIGURE 3.3: Relative L^2 error associated with the best interpolation plotted against the direction of the incident plane wave with $h/a = 0.1$, $kL = 20$. Black solid lines: p -FEM, red dashed lines: wave-based DGM. Numbers indicates the polynomial order or the number of plane waves.

3.3.1 Anisotropy

The anisotropy of the two different approximating bases is best shown using the test case of the single plane wave by varying the incident wave direction θ_0 from 0 to 2π . The relative L^2 -error of the optimal interpolation is plotted against θ_0 in Figure 3.3 for various approximation orders and with element size $h/a = 0.1$ and Helmholtz number $kL = 20$.

As expected, both methods exhibit exponential convergence with respect to the order p or the number of plane waves N_w . However, the behaviour of the two bases with respect to the propagation angle θ_0 is very different.

The polynomial basis is rather insensitive to the plane wave direction and, for a given order p , only small variations of accuracy are observed when varying θ_0 . This is due to the use of an unstructured mesh (a structured mesh would result in a more pronounced anisotropy governed by the topology of the mesh [20]).

In contrast, the plane-wave basis is very sensitive to the orientation of the incident plane wave. When it is aligned with one of the waves in the basis, the error drops to zero since the exact solution belongs to the interpolation basis. The error reaches its maximum precisely half-way between two wave directions in the basis [66]. In realistic cases, sound fields are composed of waves propagating in multiple directions and therefore, the global error will most likely be driven by the wave direction resulting in the largest error. For the remainder of the report, for this particular test case, the incident wave direction θ_0 will be taken half-way between two plane waves of the basis. This strong anisotropy of the wave-based DGM also justifies the introduction of the spinning wave test case (Figure 3.1(a)) where all wave directions are equally represented.

3.3.2 Convergence Properties

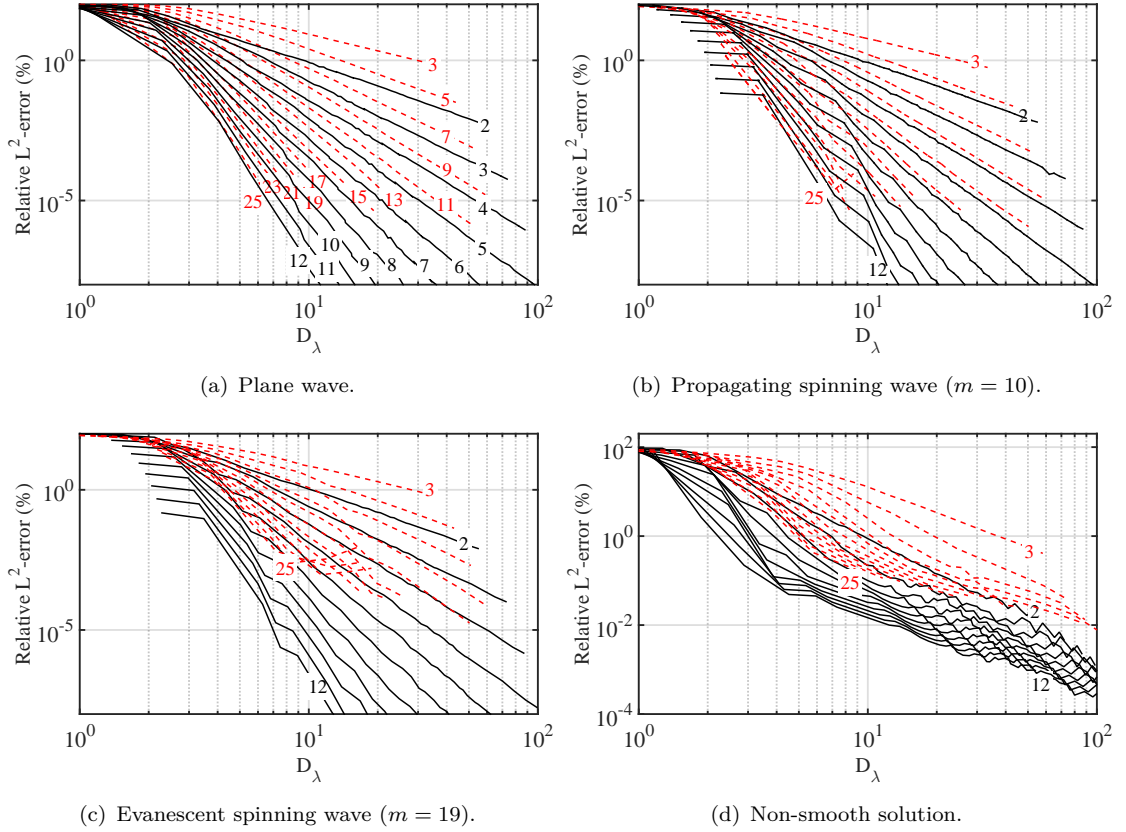


FIGURE 3.4: Relative best L^2 interpolation error (%) against D_λ at $kL = 50$; red dashed lines: wave-based method; black solid lines: p -FEM (with static condensation).

We now examine the interpolation properties of the two approximation bases with respect to order and mesh resolution. In Figure 3.4, the relative error of the best interpolation is plotted against the number of degrees of freedom per wavelength D_λ . This is done for each value of p or N_w by fixing the frequency $kL = 50$ and varying the mesh size.

For the single plane wave and the propagating spinning wave test cases, the expected behaviour is observed, with an algebraic convergence with the mesh size and an exponential convergence with p or N_w . For p -FEM, the L^2 error scales asymptotically like $D_\lambda^{-(p+1)}$ as expected [10, 20]. For the wave-based DGM the error scales like $D_\lambda^{-[(N_w+1)/2]}$ where $[x]$ denotes the integer part of x which is consistent with previous work [34, 30, 66]. These are rates of h -convergence as they have been identified by varying the element size while keeping the order or the number of plane waves constant.

For the wave-based method the interpolation error is not able to go beyond 10^{-5} . This is due to the high-conditioning of the problem: the condition number of the matrix \mathbf{X} reaches 10^{12} when the relative L^2 error is below 10^{-4} .

The conclusion from Figures 3.4(a) and 3.4(b) is that, for smooth solutions, there is no significant gap between the two bases in terms of interpolation error. Using reasonably high polynomial orders or numbers of plane waves, it is possible to achieve similar levels of accuracy with similar number of degrees of freedom per wavelength. For instance for the plane wave test case, the polynomial basis with order $p = 9$ matches closely the accuracy of the plane wave basis approximation with $N_w = 19$.

Results for the evanescent spinning wave test case are shown in Figure 3.4(c). For the p -FEM basis, the same convergence behaviour as in Figure 3.4(b) is observed, with only a slight increase in interpolation error. However for the wave basis, the interpolation error does not converge fully at the expected asymptotic rate. Instead, a region of slow convergence is observed. An inspection of the interpolation error in the domain reveals that the error is concentrated in the region close to the inner cylinder. As explained previously, the solution in this region decays very rapidly with r (as r^{-19} in this case) and this decay is difficult to represent with a plane-wave basis.

Finally figure 3.4(d) shows the results for the corner solution obtained with a uniform mesh. The exponential convergence with p and N_w is lost for both methods. The interpolation error still decreases but at the same slow rate for all orders. This behaviour is expected when dealing with solutions that are not sufficiently regular (in this case the derivative of the solution is singular). The use of local h -refinement is a valid option to improve the FEM solution and will be used in the next section.

For problems with singular solutions or evanescent waves, Trefftz methods can benefit from the introduction of additional *ad hoc* functions in their approximation basis. The use of evanescent plane waves in the basis is for instance examined to represent the transmission and reflection of sound at an interface between two media in [111] for the UWVF and in [151, 120] for the DEM. However, this approach obviously requires a detailed knowledge of the solution behaviour before the numerical model is even formulated. It may hence be difficult to generalise it to all the scenarios, particularly in the presence of complex geometries. Such practical considerations will be discussed in more detail in Section 3.6.

3.4 Performance of the Numerical Methods

In this section, the efficiency of the two numerical methods is compared. The convergence with respect to p - and h -refinements is firstly considered, followed by the study of the conditioning properties of the underlying system of equations. Finally, the computational cost at fixed accuracy is assessed.

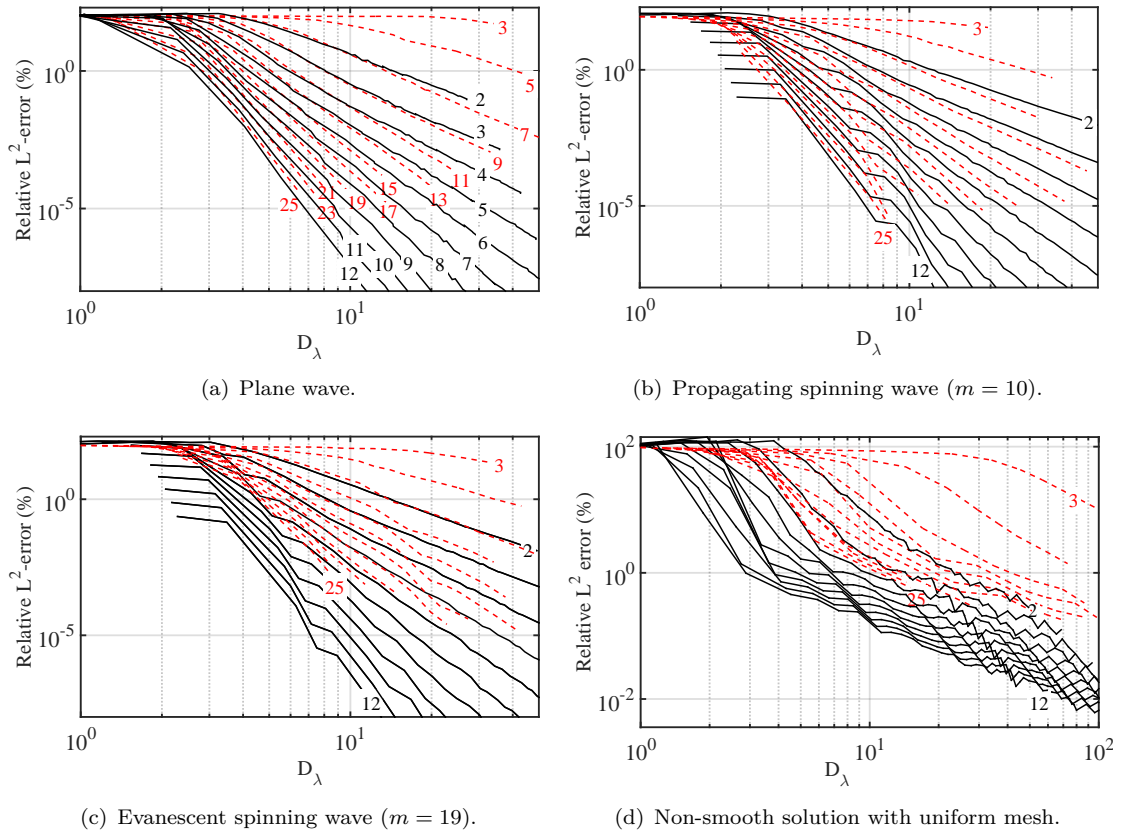


FIGURE 3.5: Relative L^2 -error (%) against D_λ at $kL = 50$; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation.

3.4.1 Convergence

Figure 3.5 shows the relative L^2 -error defined in (3.4) plotted against the number of degrees of freedom per wavelength for the four benchmark problems at $kL = 50$.

For the plane wave and the spinning wave test cases, these convergence results follow, at least qualitatively, the results in Figure 3.4 for the best interpolation error. The usual algebraic and exponential convergence rates are also observed for h -refinement and p -refinement, respectively. For the wave-based method, it is not possible to reach levels of error below 10^{-6} due to poor conditioning (this is discussed in Section 3.4.2). Again, the overall conclusion is that the two methods are able to achieve similar levels of accuracy with the same number of degrees of freedom per wavelength.

In order to provide a more quantitative comparison with the best interpolation error, it is useful to calculate the optimality coefficient C_{opt} which is defined as the ratio between the numerical solution accuracy and the best interpolation:

$$C_{\text{opt}} = \frac{E_{L^2(\Omega)}}{E_{L^2(\Omega)}^{\text{opt}}} . \quad (3.11)$$

This parameter indicates the efficiency with which the numerical formulation can exploit the interpolation properties of the underlying approximation basis. An example for the propagating spinning wave test case is provided in Figure 3.6. A large increase in C_{opt} is observed for both methods in the pre-asymptotic regime, corresponding to low values of D_λ . This increase can be attributed to the pollution effect, corresponding to the presence of large dispersion errors in the model, which typically dominate for poor resolutions. It is well known that high order methods allow to mitigate efficiently the pollution effect. This is confirmed by the current analysis. For large values of p or N_w , the increase in C_{opt} is hardly observed. It can be also mentioned that, overall, the p -FEM




FIGURE 3.6: Optimality coefficient against D_λ for the propagating spinning wave problem at $kL = 50$; red solid lines: wave-based DGM; black solid lines: p -FEM with condensation.

formulation exhibits slightly superior optimality properties. In the asymptotic regime where the interpolation error dominates, the p -FEM has close to optimal coefficients with C_{opt} varying from 1.2 to 2. By contrast, the range of variation of the wave-based C_{opt} is between 2 and 10. The difference is particularly visible for low orders: the DGM formulation is far from optimal. This indicates that there is still room for improvement in the DGM formulation, even for regular problems.

The results obtained for the evanescent spinning wave and the corner solution, depicted respectively in Figures 3.5(c) and 3.5(d), follow qualitatively those obtained for the optimal L^2 interpolation error given in Figure 3.4. For the corner solution the loss of exponential convergence with p or N_w is again clearly visible. An additional observation is that for a relative error below 1% the same convergence rate is observed for all orders p . Quantitatively, the coefficient of optimality varies between 5 and 25 for both methods, which is much higher than for the smooth solution. This indicates that the numerical schemes do not fully exploit the potentialities of their respective basis. This is an interesting finding which suggests that, as an alternative or perhaps in complement to the usual remedies (h -refinement or the introduction of special functions), modifications of the formulation could also be used to improve the performance for non-smooth solutions. This indicates that not only is the corner solution an issue in terms of interpolation but

the expected convergence properties of high-order methods are also lost. A summary of these aspects can be found in [73].

The same corner solution is examined in [73, 92] for Trefftz methods and in [158] for the p -FEM. The main outcome of the publications is that h -refinement is required to obtain accurate solutions. The study presented below aims at comparing the relative impact of this grid refinement on the accuracy of the two methods. To this end, the following local grid modification is introduced: adopting a mesh size h_{\max} at the boundaries $\{(x, y) \setminus y = \pm 1\}$ and $\{(x, y) \setminus x = \pm 1\}$, a refinement coefficient C_r is varied from 1 to 400 such that at the origin, the typical mesh size is $h = h_{\max}/C_r$. Three examples of locally refined meshes are given in Figure 3.7.

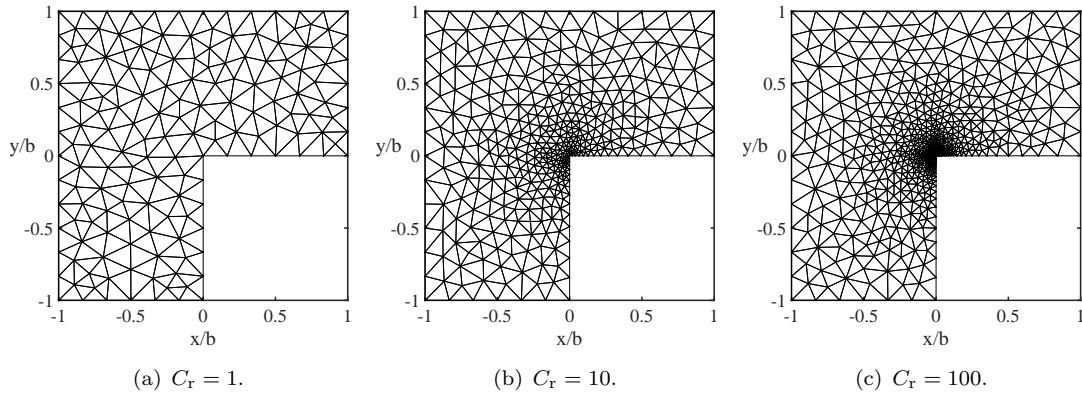


FIGURE 3.7: Examples of unstructured non-uniform meshes used for the non-smooth test case with $h_{\max}/b = 0.17$.

In Figure 3.8, the convergence of the solution is plotted against the value of C_r for a fixed mesh size $h_{\max}/b = 0.17$ and a fixed frequency ($kL = 50$). Two different regimes can be observed. For moderate refinements (i.e. low values of C_r) the error due to the singularity dominates the global numerical error. Hence, increasing the refinement allows an improvement of the relative L^2 error. It is worth noting that the C_r -convergence rate in this regime is nearly independent of the approximation order p and N_w , and is similar for both methods. For strong refinement (i.e. large values of C_r) a plateau is reached and the error stagnates as the numerical error is dominated by the approximation of the solution in the rest of the domain.

Figure. 3.8 shows the impact of the mesh refinement on the condition number. For the p -FEM, modifying C_r has little effect on the conditioning of the systems. By contrast, it has a significant impact on the behaviour of the wave-based DGM. The combination of small elements in the neighbourhood of the singularity and high number of plane waves in the basis leads to a drastic growth of the condition number. As a consequence, the number of plane waves cannot be arbitrarily increased and in this example, values above $N_w = 16$ could not be considered.




FIGURE 3.8: Relative L^2 -error (left) and condition number (right) against C_r at $kL = 50$ with $h_{\max}/b = 0.17$ for the non-smooth problem with locally refined meshes; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation.

This analysis reveals that h -refinement has a significantly different impact on the behaviour of the two methods. While it consistently enhances p -FEM results when the singularity error dominates, it should be manipulated cautiously with wave-based DGM, where it can lead to stability issues. For this method, the introduction of local *ad hoc* functions appears more adequate. In particular, an alternative is to couple the plane wave basis with Bessel functions [110].

3.4.2 Conditioning




FIGURE 3.9: Propagating spinning wave case with $kL = 50$ and $m = 10$. Left: Condition number as a function of D_λ . Right: relative L^2 error as a function of the condition number (left). Black solid lines: p -FEM with condensation, blue dotted lines: p -FEM without condensation, red dashed lines: wave-based DGM.

In this section, the conditioning properties of the two methods for solving smooth propagating problems is further examined. In Figure 3.9 the condition numbers are plotted against D_λ for the propagating spinning wave test case. The results obtained for the plane wave test case are very similar and are not repeated here.

For well resolved models (i.e. large values of D_λ), the conditioning of p -FEM tends to increase slowly with D_λ . In this regime, the polynomial order has a loose effect on

the conditioning. Conversely, for poorly resolved models (typically $D_\lambda < 10$) there is a visible effect of the order, with a progressive deterioration of the conditioning when p increases. Also shown in Figure 3.9 is the conditioning of the p -FEM model without condensation. In this case, the global system includes the degrees of freedom corresponding to the bubble shape functions in each elements. The impact of condensation on the conditioning is very significant. Without condensation, the conditioning of the p -FEM model deteriorates very rapidly when the polynomial order is increased. For this test case, each increment of p results in a growth of one order of magnitude for the condition number. For instance for $p = 10$, there is a difference of 5 or 6 orders of magnitude between the models with and without condensation. The effect of the resolution rate D_λ remains comparatively limited. This explains why the use of condensation is generally recommended to improve the performance of direct [25] as well as iterative [134] solving procedures.

In comparison with p -FEM, the conditioning of the wave-based method varies more significantly with the mesh resolution and the order N_w [34, 104]. In particular, for a large number of plane waves, the conditioning number grows very steeply with the mesh resolution. When plotting the relative error against the accuracy (see Figure 3.9), the two appear to be closely related. Independently of N_w , the results collapse on a single curve indicating that the condition number scales like $E_{L^2}^{-2}$. Hence, for reasonable levels of accuracy, say for instance 1%, the conditioning of the wave-based method is not problematic, and is in fact equivalent to the p -FEM with condensation. The conditioning of the wave-based method only becomes an issue when targeting very high levels of accuracy.

For both methods different techniques can be used to improve the conditioning, in particular by selecting other approximation bases. For the wave-based method, the use of Bessel functions improves the conditioning [110], while Bernstein polynomials have a similar impact on p -FEM [137].

3.4.3 Computational Costs

Given the large number of parameters involved, it is convenient to compare the two methods for a fixed accuracy and frequency. In this way we can assess the requirements needed for the two methods to solve a particular problem with a given error threshold. In a second stage, we will discuss to which extent the conclusions drawn are applicable to other frequencies and other levels of accuracy. For a given test case and for every order p or number of plane waves N_w , the mesh size h is varied until the numerical error reaches the target precision, then the costs and conditioning of this model are recorded. The test case of the spinning wave is used here but the conclusions are equally applicable to the plane-wave test case.

3.4.3.1 Fixed Accuracy and Frequency

The frequency is fixed to $kL = 50$ and a target accuracy of 1% is considered, which is regularly suggested as relevant for industrial applications.

Figure 3.10 shows the link between the mesh size and the factorisation time. The timing was measured on the Iridis4 cluster with Intel Xeon E5-2670 processor using one core (2.6 GHz, 4 GB of memory). Every data point in Figure 3.10 achieves the same relative error of 1% but the mesh resolution required to do so varies with p or N_w . As one would expect, low-order models (e.g. $p = 2$ or $N_w = 5$) require very fine meshes to meet the target accuracy and result in longer factorisation times. When using higher-order models it is possible to maintain the same error level with larger elements, and the factorisation time is also reduced significantly. The benefit of increasing the polynomial order or the number of plane waves tends to level off for large values of p and N_w .

Also shown in Figure 3.10 is the reduction in factorisation time obtained from condensation (the mesh resolution remains the same with and without condensation). This benefit is however balanced by the fact that, when using static condensation, the degrees of freedom associated with the bubble functions have to be calculated after the factorisation. From experience, the use of static condensation does not significantly reduce the overall runtime of the solution procedure unless the recovery of the internal degree of freedom in each element is parallelized.

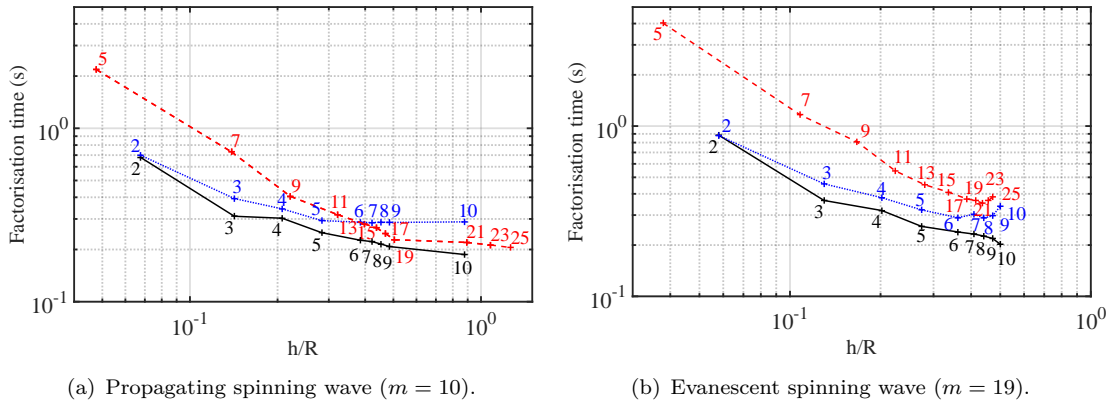


FIGURE 3.10: Factorisation time (s) against the element size required to obtain a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.

Figure 3.11 shows the evolution of the condition number and the factorisation memory when p and N_w are varied (but the mesh resolution is adjusted to maintain the error at 1%).

For the propagating spinning wave, see Figure 3.11(a), increasing N_w for the wave-based method leads to a consistent reduction in the memory required to solve the problem but the conditioning does not vary significantly. This is consistent with the results in Figure

3.9 showing that for a fixed accuracy the conditioning is only weakly influenced by the number of plane waves. For p -FEM with condensation, increasing the polynomial order also reduces the factorisation memory and induces a moderate increase in conditioning, although the condition number remains comparable to that of the wave-based method. We can see that the two methods are able to achieve the same level of accuracy with similar cost and conditioning. Without condensation, the cost of p -FEM does not greatly decrease with p (except between $p = 2$ and 3) and the conditioning is significantly higher.

For the evanescent spinning wave, Figure 3.11(b), the results for p -FEM are very similar to that for the propagating wave. However for the wave-based method, the condition number and the memory requirements have both increased. The benefit of using more plane waves therefore tends to level off. Again, this is due to the strong gradient in the solution close to the inner boundary of the computational domain.

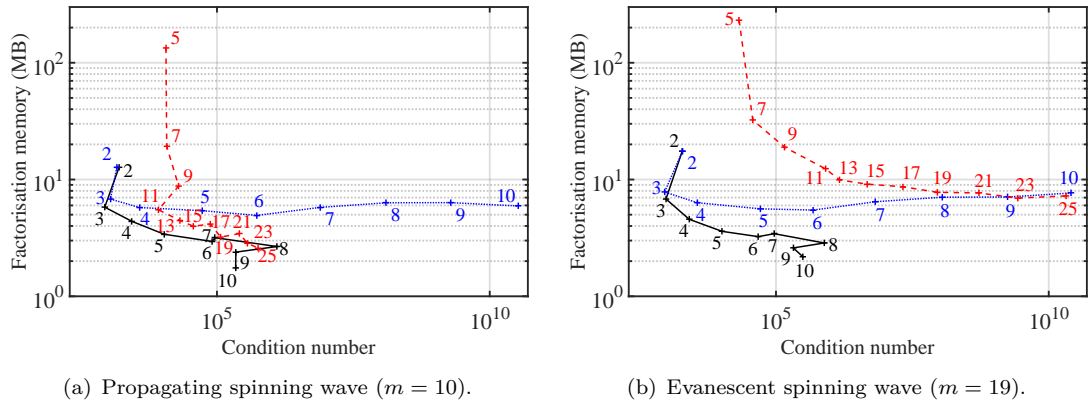


FIGURE 3.11: Factorisation memory against condition number for a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.

The behaviour of the factorisation memory can be related to the properties of the sparse, global system matrices generated by the two methods. Figure 3.12 shows the evolution of the number of non-zero entries (N_{NZ}) against the number of degrees of freedom (N_{DOF}). Note that the ratio $N_{\text{NZ}}/N_{\text{DOF}}$ provides an indication of the average number of terms in each row of the matrix. For p -FEM with condensation, increasing the polynomial order leads to a rapid reduction in the number of degrees of freedom. The same applies to the wave-based method when increasing N_w . For both methods the sparsity of the global matrices is reduced when increasing p or N_w , but the wave-based method generates denser matrices (i.e. with more non-zero entries). This difference is further accentuated when considering the evanescent spinning wave.

3.4.3.2 Influence of the Target Accuracy

So far, a fixed accuracy of 1% on the relative L^2 error was considered, but it is also interesting to examine other levels of target accuracy. Figure 3.13 shows the factorisation

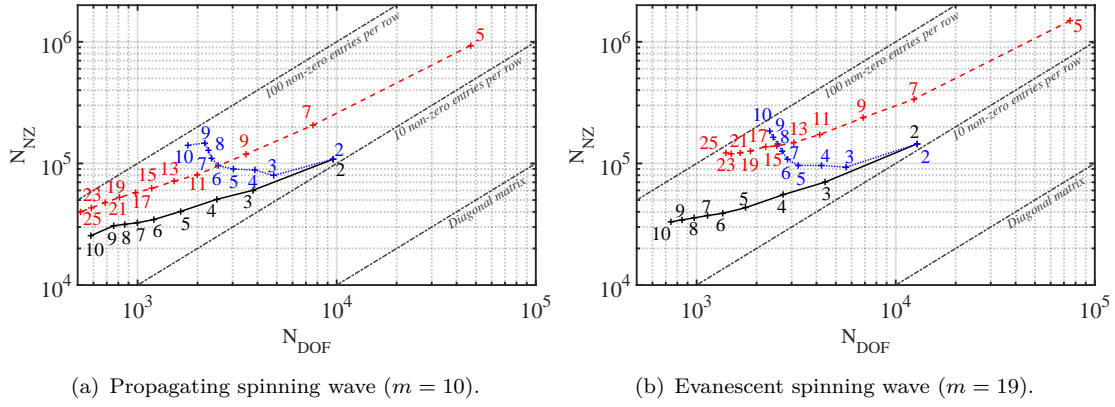


FIGURE 3.12: Number of non-zero entries against the number of degrees of freedom for a numerical error of 1% at $kL = 50$. Red dashed lines: wave-based DGM; Blue dotted lines: p -FEM without condensation; Black solid lines: p -FEM with condensation.

memory and the conditioning as a function of the number of degrees of freedom per wavelength D_λ . The target error is varied from 10% to $10^{-4}\%$. The test case of the propagating spinning wave is used, but the results for the plane wave and the evanescent spinning wave test cases are similar.

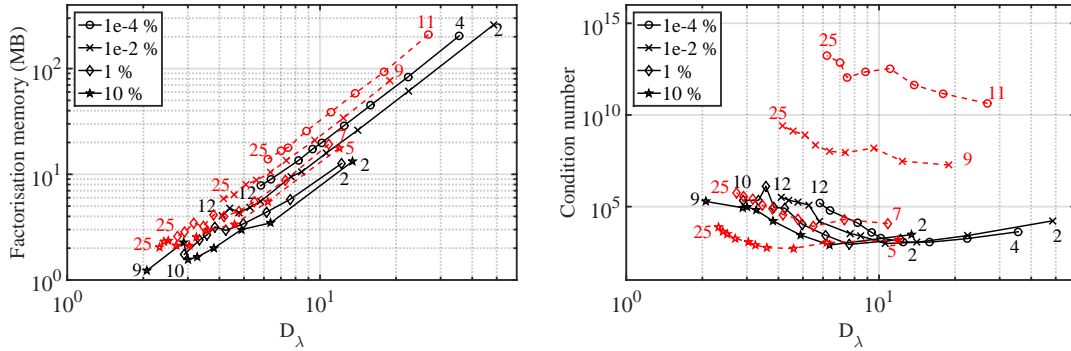


FIGURE 3.13: Propagating spinning wave test case ($kL = 50, m = 10$); factorisation memory (left) and condition number (right) against D_λ to achieve different levels of accuracy; red dashed lines: wave-based DGM, solid black lines: p -FEM with static condensation.

For the factorisation memory, the two methods follow the same trends: the memory increases when the target error level is reduced and when D_λ increases. The fact that the wave-based method requires slightly more memory than p -FEM is observed at all error levels, confirming that this observation is not specific to the 1% target accuracy used above.

The conditioning of the p -FEM model tends to increase slightly when the target error level is lowered, mostly for lower resolutions (i.e. $D_\lambda < 10$). However, for the wave-based method, the condition number increases very rapidly to reach levels that are problematic. Again, this is consistent with the results in Figure 3.9 which shows that it is only for

very accurate solutions that the conditioning of the wave-based method becomes a real issue.

3.4.3.3 Influence of the Frequency

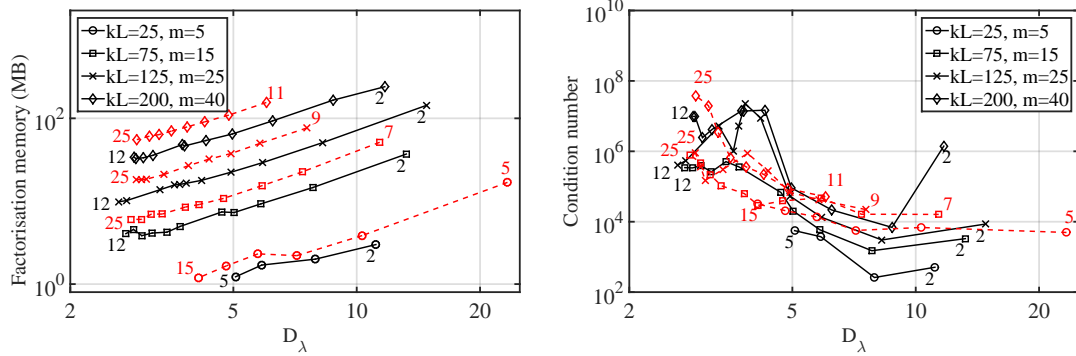


FIGURE 3.14: Propagating spinning wave problem. Factorisation memory (left) and condition number (right) against D_λ to achieve 1% of accuracy at different frequencies; red dashed lines: wave-based DGM, solid black lines: p -FEM with static condensation.

Finally, it is important to assess whether the trends described above are also observed for other frequencies. To this end, Figure 3.14 shows the factorisation memory and the conditioning for Helmholtz numbers ranging from $kL = 25$ to 200 for the propagative spinning wave test case (the azimuthal order m is adjusted at each frequency to ensure the wave is propagating). A target error level of 1% is used.


As expected, the factorisation memory is found to increase regularly with frequency and resolution. As the polynomial order or the number of plane waves is increased, the two methods are able to achieve 1% error with less memory. For every frequency, the wave-based method requires a larger amount of memory than p -FEM, although the difference is relatively small for the lowest frequency ($kL = 25$).

The evolution of the condition number with respect to the frequency is more complex and does not follow a simple trend. For $D_\lambda > 5$, p -FEM appears to be better conditioned, but the opposite is generally true for $D_\lambda < 5$.

3.5 Additional Study Involving Flow

In this section, the performance of the p -FEM and the wave-based DGM is now studied in the presence of a strong uniform background flow.

For the wave-based DGM, the basis is generally built by evenly spacing the N_w wave directions in the interval $[0, 2\pi]$ [66, 92]. However, in the presence of mean flow, a uniform distribution of the waves directions may become suboptimal. The convective effect leads

FIGURE 3.15: Plane wave basis angles θ and α .

to a shortening of the wavelength λ when the wave is propagating against the flow direction (upstream direction) and it leads to a stretching of the wavelength when the wave is propagating in the flow direction (downstream direction). When the wavelength is small, waves can propagate over many wavelength within the computational domain Ω which allows large accumulations of the error on the phase speed. A way to alleviate this issue is to adjust the plane waves distribution in order to obtain a better discretisation in the upstream direction. The plane waves in the basis have a direction θ and a ray direction α that takes into account the mean flow (see Fig. 3.15). If the flow is propagating in the $+x$ direction then from trigonometry, the following expression is derived:

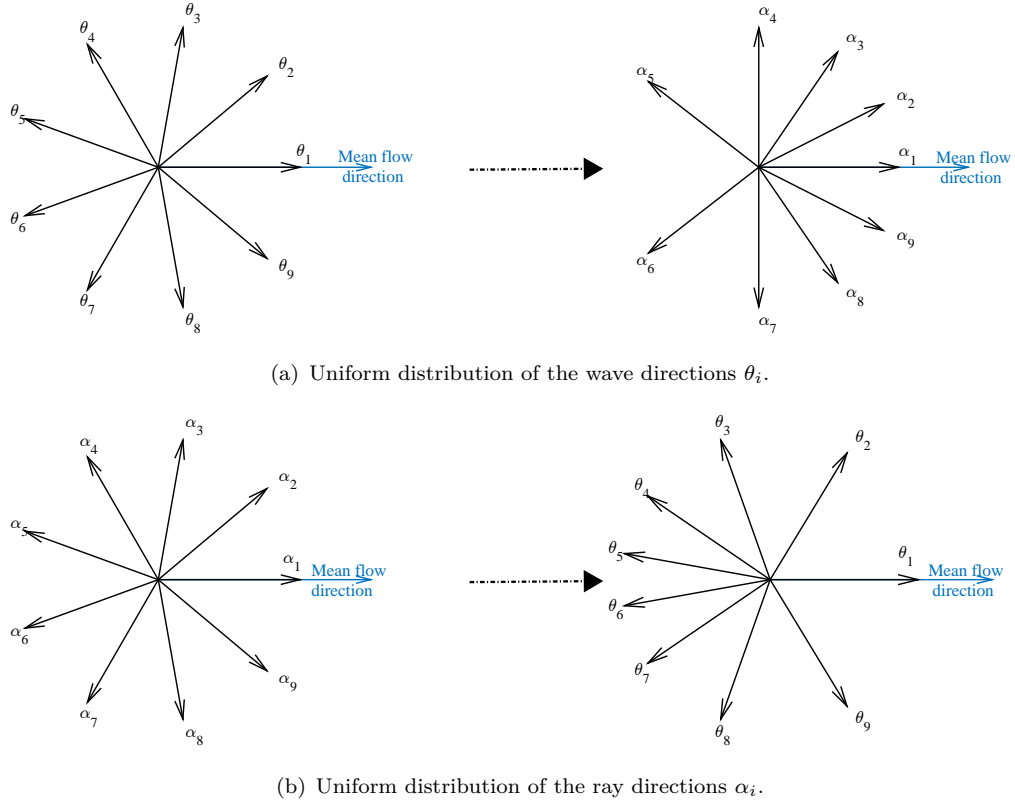

$$\theta = \alpha + \text{asin}(M \sin(\alpha)), \quad (3.12)$$

where $M = \|\mathbf{v}_0\|/c_0$ is the flow Mach number. The alternative approach is to evenly distribute the ray directions and infer the wave directions using Eq. 3.12. An illustration of the two approaches is shown in Fig. 3.16. It can be seen that, on the one hand, a uniform distribution of the θ_i results in a clustering of the ray direction in the downstream direction. On the other hand, the consequence of uniformly distribution the ray directions is a clustering of the plane waves in the upstream direction.

The two clustering approaches are used to solve a simple test case in the presence of flow and compared to the high-order polynomial method. The problem of plane wave propagation in free field exposed in Section 3.1.1 is chosen and a uniform mean flow aligned with the x axis is introduced.

3.5.1 Anisotropy

The anisotropy of the different bases is shown using the single plane wave propagation in free field test case. This problem is solved for the case of a Mach number $M = 0.5$, Helmholtz number $kL = 50$ and a typical mesh size $h/a = 0.1$. The incident wave angle is varied from 0 to 2π . The relative L^2 -error is measured and plotted against the plane wave direction θ_0 in Fig. 3.17.

FIGURE 3.16: Two approaches for the wave basis distribution with $N_w = 9$.FIGURE 3.17: Relative L^2 -error against the wave direction for $M = 0.5$, $h/a = 0.1$ and $kL = 50$; black lines: p -FEM with condensation, red lines: DGM with uniform wave distribution, magenta dashed lines: DGM with a non-uniform wave distribution. The number of plane wave or polynomial order is shown next to each curve.

The different methods exhibit exponential convergence with respect to the polynomial order and number of plane waves. However, their behaviour with respect to the propagation angle is again different.

In contrast to the no-flow case, the error obtained with the p -FEM now varies significantly with the incident wave angle. Due to the shortening of the wavelength in the upstream direction, the resolution rate D_λ is reduced when $\theta_0 = 180^\circ$ compared to $\theta_0 = 0^\circ$. Therefore, for a given polynomial order, the relative error reaches its minimum (resp. maximum) in the downstream (resp. upstream) direction. For angles close to 0° —wave propagating in the flow direction—, the errors are decreased by more than four orders of magnitude when the polynomial degree is increased from p to $p + 4$. Close to the upstream case, the benefit of increasing p is reduced but remains significant: indeed, the difference is of two orders of magnitude between the relative L^2 -error for p and $p + 4$.

For the wave-based DGM, on top of the effect of the relative alignment between the plane wave direction θ_0 and the basis wave directions (see Section 3.3.1), the convective effect is also noticed. Having a non-uniform distribution of the wave directions allows for a better resolution in the upstream direction. Compared to the uniform basis, the relative error is therefore increased in the downstream direction but reduced in the upstream direction. The result is a more uniform distribution of the error over the whole range of θ_0 .

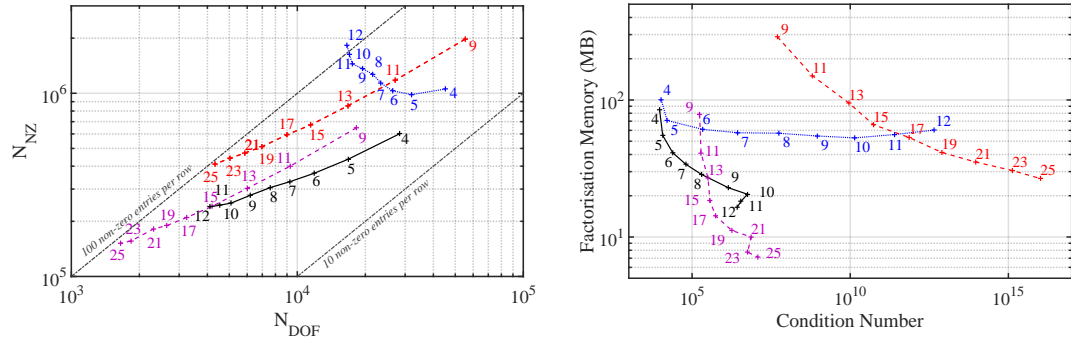
3.5.2 Performance of the Numerical Methods

The test case of the single plane wave is also used to assess the performance of the methods. For each model, the propagating wave direction which is the most penalising (meaning leading to the largest error) is considered. For the p -FEM, the upstream case is chosen. For the wave-based DGM, the incident wave direction θ_0 is taken:

- half-way between two plane waves of the basis,
- in the upstream (resp. downstream) direction in the case of the uniform (resp. non-uniform) distribution.

Figure 3.18(a) presents the number of non zeros against the total number of degrees of freedom used to achieve 1% of relative error. Increasing the order of the methods leads to a rapid reduction in the number of degrees of freedom and the sparsity of the global matrices is reduced. Similarly to the no-flow cases, we notice that the wave-based method generates denser matrices compared to the p -FEM with condensation.

Figure 3.18(b) shows the factorisation memory and condition number to achieve a defined accuracy. Static condensation for p -FEM reduces drastically not only the conditioning but also the storage and factorisation memory needed. The same effect is observed for the wave-based DGM when the basis is adjusted. The conditioning of the systems for



(a) Number of non-zero entries against the number of degree of freedom. (b) Factorisation memory against condition number.

FIGURE 3.18: Convected plane wave in free field problem with 1% of accuracy at $kL = 75$, $M = 0.5$; blue lines: p -FEM without condensation, black lines: p -FEM with condensation, red lines: DGM with uniform wave distribution, magenta dashed lines: DGM with non-uniform wave distribution. The numbers of plane wave or polynomial order are shown next to each point.

DGM with an adapted basis is not limiting: the condition numbers are similar to that of the p -FEM with condensation and they mainly depend on the accuracy, rather than the number of plane waves in the basis. The conditioning of DGM with a uniform basis and p -FEM without condensation systems to reach 1% of accuracy significantly increase with p and N_w . This is an issue if iterative solvers are to be used. For the range of orders and number of plane waves investigated, wave-based DGM with the non-uniform basis reaches the lowest demands in terms of memory.

In the presence of uniform mean flow, the malleability of the wave-based basis may be utilised to improve the method's performance. It has been observed (see Figure 3.18) that clustering the plane waves in the upstream direction leads to a drastic reduction in the computational requirements for the wave-based DGM. As a result, for the range of p and N_w investigated, to reach 1% of accuracy, the wave-based DGM with the adapted basis significantly outperforms the p -FEM with condensation. We should point out, however that, anisotropic-order bases have recently been derived for the p -FEM, and significant improvements have been reported for convected applications, including for triangular unstructured meshes [69].

3.6 Practical Considerations

So far, the discussion has focused on the intrinsic computational performance of the numerical schemes. To offer a broader perspective on the use of these two methods, it is also useful to consider more practical aspects.

The performance of the wave-based method, and other Trefftz methods in general, is obviously directly linked to the selection of canonical solutions used to construct the basis, and this has a number of practical consequences:

- As already mentioned above, when dealing with evanescent waves, wave-based methods require the introduction of *ad hoc* interpolating functions directly associated with these types of solutions. On the other-hand, the performance of p -FEM is not significantly altered by the presence of evanescent waves.
- Both methods lose p -convergence when dealing with singular solutions. For p -FEM, local h -refinement is sufficient to reach a good level of accuracy, but due to conditioning issues, the use of special functions should be preferred for wave-based DGM [110].
- For problems involving non-uniform coefficients, which have not been considered here, individual plane waves are not solutions of the governing equations. Defining a basis of simple local solutions for this class of problems is not trivial but the recent work of [95] provides a way forward. On the other hand, p -FEM can accommodate non-uniform coefficients without additional treatment.
- The presence of source terms in the domain also requires special techniques to address the fact that the plane waves are solutions of the homogeneous equations. The p -FEM can tackle such problems ‘out-of-the-box’ and a number of techniques have been proposed for wave-based methods [34, 4, 66].

It is therefore needed to adjust the Trefftz interpolation basis locally, on a case-by-case basis, which requires a detailed *a priori* knowledge of the solution. While this is acceptable for applications to academic test cases, it is difficult to envisage for real-world applications where robust, fully-automated solvers are required.

In typical applications, it is often needed to repeat a simulation over a range of frequencies. In such situations, the wave-based method requires calculating and assembling the element matrices again for each frequency due to the dependence of the shape functions on frequency, which represents a significant overhead. By contrast, with p -FEM the use of hierarchic shape functions allows to store and reuse the element matrices for other frequencies [21].

3.7 Conclusions

A detailed comparison of a high-order polynomial method (p -FEM) and a wave-based method (DGM) has been presented in terms of interpolation properties, performance and conditioning. The main objective was to put in perspective the relative merits of these two categories of discretisations.

The general argument for the use of wave-based methods is that plane waves are canonical solutions of the underlying equations and are therefore better suited than polynomials to construct an approximation basis. They incorporate key properties of the exact solution such as the wavenumber (and polarisation for electromagnetism for instance) and the expectation is that they provide a more accurate and efficient family of functions to interpolate the solution. The results presented here indicate that this intuitive expectation does not necessarily translate into a clear benefit in terms of interpolation accuracy (for a fixed resolution D_λ) or performance (for a fixed numerical error). Rather surprisingly, the high-order polynomial method is able to achieve comparable, and in some cases superior, results compared to a wave-based method. This observation was found to hold for a range of frequencies and error levels. It should be noted that the use of static condensation is crucial to obtain good performance and conditioning from p -FEM. The conditioning of the wave-based method was found to be acceptable and comparable to that of p -FEM for an error level of 1%. It is only when very high levels of accuracy are targeted that the conditioning deteriorates significantly.

The general outcome of this comparison is that there is still significant benefits in using polynomial methods, both in terms of performance and robustness, and that more work is required on wave-based methods to fully realise their potential.

It should be noted that the present work used the wave-based DGM as an example of Trefftz method and other wave-based methods should also be considered. The wave-based DGM is closely related to a number of other methods, in particular the UWVF [34] and the least-squares method [125]. Both methods can be recovered in the DG framework by using different numerical fluxes. Existing results indicate that the change of numerical flux is unlikely to change the performance of the methods sufficiently to alter the conclusions in this work [71]. But other categories of wave-based methods, such as DEM and PUFEM remain to be compared systematically with high-order polynomial methods.

For the remainder of this work, we therefore choose to discretise the problems using the hierarchical high-order FEM. In the next sections, the solver used for the solution of the system matrix arising from this discretisation is investigated.

Chapter 4

Review of Domain Decomposition Methods

In this chapter, we review some methods developed to solve the linear matrix system arising from the finite element discretisation of the Helmholtz equation:

$$\mathbb{K}\mathbf{x} = \mathbf{b}, \tag{4.1}$$

where \mathbb{K} is a large, square, sparse and complex matrix; \mathbf{x} is the unknown vector and \mathbf{b} is the right-hand side.

Difficulties are encountered when using classical direct and iterative solvers for the solution of large scale Helmholtz problems. Domain decomposition methods constitute an alternative approach to purely direct or iterative methods. In this chapter, Schwarz algorithms, from which the domain decomposition originated, are presented, followed by a review of more recent methods.

4.1 On the Difficulty to Solve the Helmholtz Equation with Classical Solvers

There are traditionally two classes of methods to solve a linear system: direct methods and iterative methods.

Until recently, for real life engineering applications, direct solvers were often preferred to iterative methods because of their robustness and predictable behaviour. Provided the matrix \mathbb{K} is not singular and in the case of exact arithmetic — meaning no rounding error — the solution computed is the exact solution of the linear system. The number of steps required to reach a solution is known beforehand. Most direct methods derive from Gaussian elimination. They transform matrix \mathbb{K} into a product of other matrices

which are easier to invert for a given right-hand side. For example, the LU factorisation [155] of \mathbb{K} generates lower and upper triangular matrices, \mathbb{L} and \mathbb{U} respectively, such that $\mathbb{K} = \mathbb{L}\mathbb{U}$. For a dense matrix of dimension N , this procedure requires $\mathcal{O}(N^3)$ arithmetic operations making the factorisation of very large matrices extremely computationally intensive. The initial system is then written as $\mathbb{L}\mathbb{U}\mathbf{x} = \mathbf{b}$ and is solved in two steps:

$$\mathbb{L}\mathbf{y} = \mathbf{b}, \quad \mathbb{U}\mathbf{x} = \mathbf{y}. \quad (4.2)$$

These two systems are triangular and can be solved very easily using forward and backward substitutions. These substitutions only require $\mathcal{O}(N^2)$ arithmetic operations per right-hand side.

The most resource intensive operation is therefore the matrix factorisation. For a banded matrix, the number of operations scales like $\mathcal{O}(b_w^2 N)$, where b_w refers to the bandwidth (any entry (i, j) of the matrix is zero whenever $|i - j| > b_w$) [80]. The algorithm cannot however take advantage of the zero entries inside the band, as these typically fill-in with non-zeros during the factorisation, yielding extra memory requirements. Sparse matrices arising from a FEM discretisation on an unstructured grid have a non regular non-zero pattern distribution. For these matrices, the cost of the factorisation depends to a large extent on the partitioning for distributed memory calculations and on the re-numbering strategy (which aims at limiting the typical bandwidth b_w). A large variety of direct solvers specialised in sparse matrices are available [47]. An example of sparse method is the frontal technique due to Irons [96], initially applied to symmetric positive-definite matrices and then extended to unsymmetric matrices by Hood [89]. The idea is to consider the sparse matrix as the sum of element matrices. The LU decomposition is progressively done on a subset of elements (called front) at a time. An improvement of the frontal solver is the multi-frontal method [48] using multiple fronts at a time. Parallel multi-frontal solvers like MUMPS [5, 1] used in the previous chapter and PARDISO [148, 102] stand among the most efficient sparse linear solvers. They proceed in three steps:

- Analysis: it begins with the fill-in minimisation through the reordering of rows and columns of the coefficient matrix. A symbolic factorisation is then performed. It aims at finding the exact positions of the new non-zero entries and allocating memory. At this stage, the numerical values of the entries of the matrix \mathbb{K} are not required.
- Factorisation: the factors \mathbb{L} and \mathbb{U} are computed and stored. The numerical values of the entries of \mathbb{K} are hence required, but not yet the values of the component of the right-hand sides (RHS).
- Solution: the two systems (4.2) are solved for all the RHS using the factors stored in-core (meaning in the main memory) or out-of-core (meaning in the auxiliary memory).

However, direct solving procedures do not scale well with respect to the problem size, the limitation lying in the storage of the factors. Large system matrices appear in our case, especially when high-frequency problems have to be tackled. When doubling the frequency ω , the usual meshing criterion implies to divide the grid length by a factor two. On a three-dimensional grid the factorisation hence scales roughly with $\mathcal{O}(\omega^3)$ neglecting the fact that the bandwidth also generally increases with the model size. Doubling the frequency generally yields a factor ten on the memory requirements. This third order dependency is also observed on high-order models with p -adaptivity [21]. Hence, for large scale applications at high frequency, the amount of computational resources available is in practice often not sufficient.

In contrast with direct solvers, iterative methods approach the solution gradually, rather than in one large computational step, without modifying the matrix \mathbb{K} . As a matter of fact, most iterative methods do not require the explicit representation of the coefficient matrix but only the product of the matrix with some vector. An initial guess for the solution is successively refined by repeating an algorithm in order to (hopefully) approach the exact solution of the system. The matrix-vector operations on which iterative methods rely on are relatively cheap processes for large scale sparse matrices. However, the exact number of steps needed to achieve a solution cannot be predicted. Several aspects influence the efficiency of the approach: the solution accuracy that is required, the initial guess which can be close or far from the solution of the given system, and the properties (clustering of the eigenvalues) of the matrix. In terms of computational complexity, iterative methods scale much better than direct methods because they only require the storage of a few additional vectors (there is no fill-in involved). In practice, the initial matrix equation has to be modified into a system that is easier to solve. This process is called preconditioning. It usually increases the cost of each iteration, however a proper choice of preconditioner is expected to improve the efficiency and robustness of the solver [146, 50].

Even though the early iterative techniques such as Gauss-Seidel method are still used, the Krylov methods are now considered the methods of choice for solving large sparse linear systems. From an initial guess \mathbf{x}^0 , an approximate solution \mathbf{x}^m is sought in a subspace $\mathbf{x}^0 + \mathcal{K}_m$ of dimension m under some optimality conditions (minimising the 2-norm of the m^{th} residual for example). This amounts to imposing the orthogonality of the residual \mathbf{r}^m with respect to another subspace \mathcal{L}_m [146]. \mathcal{K}_m is called the Krylov subspace and:

$$\mathcal{K}_m(\mathbb{K}, \mathbf{r}^0) = \text{span}\{\mathbf{r}^0, \mathbb{K}\mathbf{r}^0, \mathbb{K}^2\mathbf{r}^0, \dots, \mathbb{K}^{m-1}\mathbf{r}^0\}, \quad (4.3)$$

where $\mathbf{r}^i = \mathbf{b} - \mathbb{K}\mathbf{x}^i$. The efficiency of a Krylov method depends to a large extent on the choice of the constraints used to build the approximations *i.e* the choice of the subspace \mathcal{L}_m .

For the solution of complex symmetric, non-Hermitian and indefinite systems, methods such as the Generalised Minimal Residual (GMRES) [147], BiCGSTAB [157], and Quasi Minimal Residual [63] can be considered. A review can be found in [146].

In 1986, at a time when the use of iterative solvers was popular mainly for either diagonally dominant matrices or positive-definite matrices (Conjugate Gradient for symmetric matrices), Saad and Schultz proposed the GMRES [147] for general unsymmetric systems. The principle is to solve a least-squares problem at each iteration: at step m , the approximate solution is sought in the m^{th} Krylov space such that the L^2 norm of the residual \mathbf{r}^m is minimised. This is equivalent to choosing $\mathcal{L}_m = \mathbb{K}\mathcal{K}_m$. If we look for an orthonormal basis for \mathcal{K}_m formed by the columns of $\mathbf{Q}_m = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_m]$, the solution $\mathbf{x}^m \in \mathcal{K}_m$ can be expressed as:

$$\mathbf{x}^m = \mathbf{Q}_m \mathbf{y}^m ,$$

where \mathbf{y}^m is a coefficient vector of dimension m . The function to be minimised becomes:

$$J = \|\mathbb{K}\mathbf{x}^m - \mathbf{b}\|_{L^2} = \|\mathbb{K}\mathbf{x}^m - \mathbf{b}\|_{L^2} = \|\mathbb{K}\mathbf{Q}_m \mathbf{y}^m - \mathbf{b}\|_{L^2} .$$

The GMRES method converges in maximum N iterations, where N is the dimension of the linear system. However, for large systems, the number of iterations required to converge may be considerable and therefore, the storage may become a problem. To remedy this problem, a restarted version of GMRES can be used: after m iterations, the computation is restarted with \mathbf{x}^m as initial guess [147].

For the solution of large linear systems, iterative techniques have been gaining popularity in many areas of scientific computing. However, for wave problems, the oscillatory nature of the solution and the indefiniteness of the final matrix lead to severe convergence problems: slow convergence, stagnation or even divergence when one is close to a resonant frequency of the problem [146]. A comparison of the convergence of the GMRES algorithm for a Helmholtz and Laplace problem is given in [50]. Unless a preconditioner is employed, the authors show that the convergence for the Helmholtz problem stagnates.

4.2 Domain Decomposition Methods for Wave Problems

As mentioned above, the discretisation of wave problems at high-frequency yields very large sparse linear systems that are often too costly to be solved using a direct procedure. Besides, due to the indefiniteness of the equation, classical iterative methods do not converge or converge only slowly without any theoretical promises of convergence [50].

When solving a numerical problem, a division of the initial computational domain into smaller sub-domains may be natural. It is the case for example when different physics are encountered in different regions of the domain or in the presence of a strongly




FIGURE 4.1: Original domain used by Schwarz; the complex domain is composed of two simple overlapping geometries, a disk Ω_1 and a rectangle Ω_2 .

heterogeneous medium. This partition is also a way to reduce the size of the full problem to a set of smaller local problems that could be solved by a direct procedure. This principle of divide-and-conquer lies at the heart of the domain decomposition methods.

4.2.1 Original Schwarz Methods

In 1870, to prove the existence and uniqueness of a solution of a Poisson problem on irregular domains, Schwarz devised the first domain decomposition method [75]. The problem was the following one:

$$\begin{cases} \nabla^2 u = 0 & \text{in } \Omega, \\ u = g & \text{on } \partial\Omega, \end{cases}$$

where Ω is a union of simple overlapping geometries as illustrated in Fig. 4.1. Schwarz proposed an innovative approach to solve this particular problem. Starting from an initial guess u_1^0 on the interface $\Gamma_{12} = \partial\Omega_1 \cap \Omega_2$, the solutions u_i^{l+1} in domains $i = 1, 2$ are alternatively computed for $l = 0, 1, 2, 3, \dots$, following:

$$\begin{cases} \nabla^2 u_1^{l+1} = 0 & \text{in } \Omega_1, \\ u_1^{l+1} = g & \text{on } \partial\Omega_1 \cap \partial\Omega, \\ u_1^{l+1} = u_2^l & \text{on } \Gamma_{12}, \end{cases} \quad \begin{cases} \nabla^2 u_2^{l+1} = 0 & \text{in } \Omega_2, \\ u_2^{l+1} = g & \text{on } \partial\Omega_2 \cap \partial\Omega, \\ u_2^{l+1} = u_1^{l+1} & \text{on } \Gamma_{21}. \end{cases} \quad (4.4)$$

The portion of the boundary of a sub-domain i that intersects another sub-domain j defines an artificial interface $\Gamma_{i,j}$. It is used to exchange information at each step by the means of Dirichlet conditions; the restriction of the solution in the sub-domain i on the interface $\Gamma_{j,i}$ is used as boundary condition for the sub-problem in the sub-domain j . The sub-problems are sequentially solved until the solutions in the overlap match within some prescribed tolerance. The resulting algorithm is nowadays called the alternating Schwarz method. A slight modification by Lions [108] rendered the algorithm suitable for parallel computing: starting from initial guesses u_1^0 and u_2^0 , the solutions in the two sub-domains can be computed simultaneously by changing the iteration index from

$l + 1$ to l in the transmission condition of Ω_2 in algorithm (4.4): $u_2^{l+1} = u_1^l$. This new algorithm is referred to as the additive Schwarz method.

The convergence of these algorithms was proved for overlapping sub-domains [108, 121]. They lead to reduced memory requirements due to the smaller size of the systems being solved, but their convergence is often slow. Besides, when using non-overlapping sub-domains, the techniques do not converge as the transmission conditions enforce solely the continuity of the solution and not the continuity of the fluxes. In fact, the convergence rate is shown to be a continuous function of the size of the overlap: the larger the overlap, the better it converges [129, 46]. A Fourier analysis [75] also reveals that the overlap allows a quick damping of the high-frequency components of the error. The low frequency part needs special treatment in order to increase the convergence rate. Moreover, the Schwarz methods do not scale with the number of sub-domains: the iteration number and execution time grow linearly with the number of sub-domains along one direction [46]. Another drawback of these methods is that for some partial differential equations such as the Helmholtz equation, convergence is not verified.

4.2.2 Optimised Schwarz Conditions and Schur Methods

Three main reasons motivated the development of new classes of Schwarz methods: extend the Schwarz methods to non-overlapping sub-domains; extend the domain decomposition methods to the Helmholtz equation; and accelerate the convergence of the Schwarz methods. In order to remedy the drawbacks of the original Schwarz methods, a large variety of methods have been devised. The improvements mainly come from the introduction of Krylov methods, more general transmission conditions [109] or coarse spaces allowing a global transfer of information.

Different domain decomposition methods have different ways of treating the variables at the interfaces between sub-domains. They can essentially be distinguished by:

- the presence of an overlap,
- the processing of the interfaces' unknowns (choice of the interface variables and conditions, choice of the algorithm for the solution of the interface problem),
- the choice of the solver for the sub-domain solutions.

Overlapping methods [31, 36] have a simpler algorithmic structure, because there is no need to solve special interface problems between neighbouring sub-domains. However, non-overlapping methods are still popular because the notion of interface is physically more natural to interpret than overlaps. The interface between domains is a concept that is linked to conservation principles stating the equilibrium of fluxes and mass. Moreover, in the case where different models are used in different regions of the domain or in the case of problems with discontinuous properties, it is natural to use a non-overlapping decomposition with interfaces along the discontinuities. Most importantly, the number

of unknowns shared between the sub-domains is limited when using non-overlapping methods. Therefore, in this work, we will focus solely on the class of non-overlapping methods.

For non-overlapping sub-domains, one has to find a way to match the solution and its normal derivative at the interfaces between sub-domains such that:

$$\begin{cases} \nabla^2 u_1 = 0 & \text{in } \Omega_1, \\ u_1 = g & \text{on } \partial\Omega_1 \cap \partial\Omega, \\ u_1 = u_2 & \text{on } \Gamma_{12}, \end{cases} \quad \begin{cases} \nabla^2 u_2 = 0 & \text{in } \Omega_2, \\ u_2 = g & \text{on } \partial\Omega_2 \cap \partial\Omega, \\ \nabla u_2 \cdot \mathbf{n}_2 = \nabla u_1 \cdot \mathbf{n}_1 & \text{on } \Gamma_{21}. \end{cases} \quad (4.5)$$

where Γ_{12} and Γ_{21} refer to the same physical space. However, if one follows the example of the classical Schwarz method and derive an algorithm based on formulation (4.5), the method would not converge as shown by the Fourier analysis in [74].

The first non-overlapping method was introduced by Lions [109]. Lions extended the application range of the alternating Schwarz method to non-overlapping sub-domains by using Robin interface conditions. For example, when the numerical domain is split into two non-overlapping sub-domains, (4.4) becomes for $i = 1$:

$$\begin{cases} \nabla^2 u_1^{l+1} = 0 & \text{in } \Omega_1, \\ u_1^{l+1} = g & \text{on } \partial\Omega_1 \cap \partial\Omega, \\ \nabla u_1^{l+1} \cdot \mathbf{n}_1 + \alpha_1 u_1^{l+1} = \nabla u_2^l \cdot \mathbf{n}_1 + \alpha_1 u_2^l & \text{on } \Gamma_{12}, \end{cases} \quad (4.6)$$

where α_1 is a constant. Lions proved the convergence of his algorithm for an arbitrary number of non-overlapping sub-domains using energy estimates [109, 74]. This new Schwarz method is faster than the original method. It inspired the development of a new class of domain decomposition method called the optimised Schwarz methods and the first decomposition method for the Helmholtz equation.

Another remedy used to extend the Schwarz methods to non-overlapping sub-domains is to keep the initial coupled problem (4.5) and introduce a new set of variables at the interface:

$$\begin{cases} \nabla^2 u_1 = 0 & \text{in } \Omega_1, \\ u_1 = g & \text{on } \partial\Omega_1 \cap \partial\Omega, \\ u_1 = u_2 = \lambda_p & \text{on } \Gamma_{12}, \end{cases} \quad \begin{cases} \nabla^2 u_2 = 0 & \text{in } \Omega_2, \\ u_2 = g & \text{on } \partial\Omega_2 \cap \partial\Omega, \\ \nabla u_2 \cdot \mathbf{n}_2 = \nabla u_1 \cdot \mathbf{n}_1 = \lambda_d & \text{on } \Gamma_{21}. \end{cases} \quad (4.7)$$

where λ_p and λ_d are the newly introduced variables.

This new class of methods is usually referred to as Schur methods. On the one hand, primal Schur methods such as the Neumann-Neumann method presented in [28] and further developed into the Balanced Domain Decomposition (BDD) method [115, 105] adopt u_i (λ_p) as the main unknown at the interface. Using the relation between the normal derivatives in (4.7), a linear equation can be retrieved: $\mathbb{F}_p \lambda_p = \mathbf{d}$. When λ_p is

known, solutions in the sub-domains can be computed. On the other hand, dual Schur methods such as the Finite Element Tearing and Interconnecting (FETI) [60] method employ the gradient (λ_d) as interface unknown. They use the continuity of the primal field at the interface to then build an interface problem. For both primal and dual approaches, the matrix \mathbb{F} has an improved conditioning compared to the original problem [74]. The interface problems are solved iteratively whereas the solutions in each sub-domain are recovered using a direct method. These methods therefore combine the robustness of a direct solver and the advantages in terms of memory requirements of an iterative solver. The properties of the methods, their robustness compared to traditional iterative solvers and their low costs (in terms of memory and CPU requirements) compared to direct solvers are shown in [119, 59]. Further analysis of the connection between BDD and FETI methods can be found in [107, 117]. An important observation is that the iteration count increases only weakly with the size of the sub-domains and global problems. Since their first emergence, a large number of variants of the BDD and FETI methods were published. For instance, the FETI-DP (dual primal) and BDDC (by constraints) introduce special treatments at the corners (points shared by more than two sub-domains). For example, at the sub-domain interfaces, the FETI-DP enforces equality of the solution by Lagrange multipliers except at sub-domain corners where the constraints remain in terms of primal variables. The FETI-DP was found to be faster than, and as numerically scalable as, the original FETI [56, 55]. More complete reviews can be found in [46, 129]. The convergence theory of the BDDC method was developed in [116]. The FETI was extended to heterogeneous problems in [142].

The traditional Schwarz methods are successful in the case of Poisson or elasticity problems because they yield symmetric positive-definite matrices. However, issues arise when the matrices are indefinite. For example, when applied to the Helmholtz equation, the Schwarz algorithm may be ill-posed:

$$\begin{cases} -k^2 u_1^{l+1} - \nabla^2 u_1^{l+1} = 0 & \text{in } \Omega_1, \\ u_1^{l+1} = u_2^l & \text{on } \Gamma_{12}, \end{cases} \quad \begin{cases} -k^2 u_2^{l+1} - \nabla^2 u_2^{l+1} = 0 & \text{in } \Omega_2, \\ u_2^{l+1} = u_1^l & \text{on } \Gamma_{21}, \end{cases} \quad (4.8)$$

where we omit from now on that u_i has to satisfy the Dirichlet condition on $\partial\Omega$ for simplicity. In this case, if k^2 is an eigenvalue of the Laplace operator in a sub-domain i , a non zero function v exists such that:

$$\begin{cases} -\nabla^2 v = k^2 v & \text{in } \Omega_i, \\ v = 0 & \text{on } \partial\Omega_i. \end{cases}$$

The problem (4.8) is therefore ill-posed as either there is no solution to it or if a solution u_i^e exists, then any function of the form $u_i^e + \beta v$ where β is a constant is also a solution. A convergence analysis in the case where the problem is well-posed is derived in [76, 46]. The study shows that without overlap the Schwarz method does not converge and with overlap, only the vanishing modes in the error are damped, the propagative modes being

unaffected by the algorithm. An alternative was to add a coarse mesh for the propagative modes as done in [32], or to change the transmission conditions. Després was the first to propose a non-overlapping algorithm for the Helmholtz equation (subsequently for the Maxwell equation) and to prove its convergence using energy estimates [44]. He in fact extended Lions' algorithm to the Helmholtz equation:

$$\begin{cases} -k^2 u_i - \nabla^2 u_i &= 0 & \text{in } \Omega_i, \\ \nabla u_i \cdot \mathbf{n}_i + i k u_i &= \nabla u_j \cdot \mathbf{n}_i + i k u_j & \text{on } \Gamma_{i,j}, \quad \forall j \in \mathcal{N}_i, \end{cases} \quad (4.9)$$

where \mathcal{N}_i is the set of indices of the sub-domains that share an interface with Ω_i (such that $\partial\Omega_j \cap \partial\Omega_i \neq \emptyset$). We now observe that even in the case where k^2 is an eigenvalue of the Laplace operator with a Dirichlet boundary condition in a sub-domain i , the solution of the problem is unique [46]. Després' algorithm is therefore always well-posed. This idea inspired the methods described in [17, 40, 41, 113].

The convergence of the method can be enhanced with a generalisation of the Robin boundary condition. The general transmission condition can be written as:

$$\mu_i \nabla u_i^{l+1} \cdot \mathbf{n}_i + \mathcal{B}_{ij}(u_i^{l+1}) = \mu_i \nabla u_j^l \cdot \mathbf{n}_i + \mathcal{B}_{ij}(u_j^l) \quad \text{on } \Gamma_{ij}, \forall j \in \mathcal{N}_i, \quad (4.10)$$

where the μ_i are real-valued functions and \mathcal{B}_{ij} are operators acting on the interfaces Γ_{ij} . In the case where for all i and j , $\mu_i = 0$ and $\mathcal{B}_{ij} = \mathbb{I}$, (4.10) corresponds to the original Schwarz method. If for all i and j , $\mu_i = 1$ and $\mathcal{B}_{ij} = \alpha$ (respectively $\mathcal{B}_{ij} = i k$), the algorithm corresponds to Lions' (respectively Després') algorithm. Optimised Schwarz methods look for μ_i and \mathcal{B}_{ij} such that the fastest convergence is obtained in terms of iteration counts. The optimal convergence is obtained by replacing the transmission operator \mathcal{B}_{ij} by the non-local Dirichlet to Neumann (DtN) operator [129] with $\mu_i = 1$. In the case of a slice-wise decomposition, the resulting methods converge in N_s steps with N_s the number of sub-domains. Specific studies on the Helmholtz equation can be found in [39, 38, 76]. However, the use of non-local operators in the transmission conditions renders the method costly to run and difficult to implement. Besides, in the general case of variable coefficient operators or complex geometries, the exact form of the DtN operator is hardly known. A solution is to construct local approximations of this operator. In the field of artificial boundary conditions, these approximations constitute a common problem [49]. An approximation using a Taylor expansion for the convection diffusion problem is given in [130]. Local approximations for the optimal transmission condition can be found in [76, 38, 44, 26, 27] for Helmholtz problems.

The Schur methods have also been adapted to Helmholtz problems by incorporating ideas coming from optimized Schwarz methods. In the FETI-H [58], instead of using Neumann transmission conditions at the interface (4.7), Robin conditions are adopted. The resulting interface problem is symmetric non-Hermitian and solved using the generalized conjugate residual method. This method is efficient not only in terms of computational





FIGURE 4.2: Adjustment of the sign on both sides of the interfaces for a chequerboard partition.

performance but also from a numerical scalability point of view: the technique is scalable with respect to the mesh size and sub-domain size when preconditioned. A FETI-DP formulation also exists for Helmholtz problems [52]. The primal sub-structured methods have also been extended to the indefinite problem in [106] and named BDDC-H.

The FETI methods are amongst the most efficient domain decomposition method for large scale systems. In [76], different non-overlapping methods are studied: the FETI-DPH and the BDDC-H (both using preconditioners) are compared against the FETI-H and an optimised Schwarz method on a 3D model problem. The author concludes that in terms of iterations numbers, the FETI-DPH performs the best. However, the FETI methods impose one Robin condition on an interface. Therefore, one set of interface unknowns called Lagrange multipliers is shared by two sub-domains. This means that on either side of the interface, the sign of the interface unknowns has to be changed in order to have a consistent Robin condition (see Fig. 4.2). Otherwise the sub-problems may be ill-posed. Colouring techniques are used for this purpose [58].

An alternative would be to introduce independent Lagrange multipliers on either sides of the interfaces, as initially done in [41]. In [41, 145, 57], a FETI-H approach using two Lagrange multipliers (FETI-2LM) at the interface between the sub-domains is presented. The method is in fact an optimised Schwarz method in sub-structured form. It can be seen as a reformulation of the Després' algorithm [44] with two Lagrange multipliers. This approach leaves out the need for the adjustment of the sign on either sides of an interface and is therefore easier to implement. The resulting interface problem is not symmetric nor Hermitian. A validation of the FETI-2LM method for solving electromagnetic frequency-domain problems can be found in [15]. The FETI-H and FETI-2LM are applied to an acoustic scattering problem with a submarine shaped obstacle in [57]. It is shown that a characteristic of both methods is that, even in the absence of any preconditioner, they are scalable with respect to the mesh size. However, the scalability with respect to the number of sub-domains is obtained only when preconditioned. The

FIGURE 4.3: Example of chequerboard splitting with $N_s = 4$.

preconditioning technique presented in [57] consists in introducing a second level FETI-H problem by forcing the residual of the Krylov method to be orthogonal to a coarse interface matrix subspace. The authors found that this preconditioner is more efficient for the FETI-H than for the FETI-2LM.

In the next section, in order to illustrate how FETI methods work, we present the FETI-2LM method for the Helmholtz equation.

4.3 The FETI-2LM Method for the Helmholtz Equation

In this section, we present the formulation of the Helmholtz sub-problems and the FETI-2LM method.

4.3.1 Formulation of the Sub-problems

We consider the propagation of harmonic, linear sound waves in a uniform, quiescent medium within a domain Ω described by the Helmholtz equation (see section 2.1). Convection propagation problems including mean flows are investigated in the last chapter. Instead of solving one problem in the whole domain Ω , the domain decomposition methods define several local problems. The initial numerical domain is decomposed into N_s non-overlapping smaller sub-domains $\{\Omega_i\}_{i=1, \dots, N_s}$:

$$\Omega = \bigcup_{i=1}^{N_s} \Omega_i, \quad \forall j \in \mathcal{N}_i,$$

where \mathcal{N}_i is the list of indices j , such that Ω_j and Ω_i are neighbours, meaning that after discretisation they share at least a boundary element. We denote the interface between the sub-domains i and j , $\Gamma_{i,j} = \Gamma_{j,i} = \partial\Omega_i \cap \partial\Omega_j$.

The choice of the transmission condition at the interfaces is important to ensure that the problem is well-posed. The characteristics conditions (Robin conditions) are used. As detailed in 2.1.5.1, at the interface between two sub-domains, the solution is separated into incoming and outgoing parts. The idea is then to impose the incoming wave using the solution in the neighbouring sub-domains. In the sub-domain i , the sub-problem reads:

$$\begin{cases} -k^2 u_i - \nabla^2 u_i &= 0 & \text{in } \Omega_i, \\ \nabla u_i \cdot \mathbf{n}_i + iku_i &= g_i & \text{on } \partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i}, \\ \nabla u_i \cdot \mathbf{n}_i + iku_i &= \nabla u_j \cdot \mathbf{n}_i + iku_j & \text{on } \Gamma_{i,j}, \quad \forall j \in \mathcal{N}_i, \end{cases} \quad (4.11)$$

where u_i is the unknown in Ω_i , \mathbf{n}_i refers to the unitary outward normal to $\partial\Omega_i$ and g_i is a given source term. The variational formulation in domain i can be written and the problem becomes: for $g_i \in L^2(\partial\Omega_i)$, find $u_i \in H^1(\Omega_i)$ such as

$$\begin{aligned} \forall w_i \in H^1(\Omega_i), \int_{\Omega_i} -k^2 \bar{w}_i u_i + \nabla \bar{w}_i \cdot \nabla u_i + \int_{\partial\Omega_i} ik \bar{w}_i u_i &= \int_{\partial\Omega_i \setminus \Gamma} \bar{w}_i g_i \\ &+ \int_{\Gamma} \bar{w}_i (-\nabla u_j \cdot \mathbf{n}_j + iku_j), \quad \forall j \in \mathcal{N}_i. \end{aligned} \quad (4.12)$$

This formulation requires the approximation of the normal gradient of the solution at the interface, which may lead to a deterioration of the solution. This can be avoided by introducing additional variables at the interfaces between the sub-domains.

4.3.2 Introduction of the Lagrange Multipliers

In the transmission conditions, Lagrange multipliers $\lambda_{i,j}$ can be introduced to avoid the computation of the RHS normal derivatives along the interface $\Gamma_{i,j}$:

$$\lambda_{i,j} = \nabla u_j \cdot \mathbf{n}_i + iku_j. \quad (4.13)$$

The sub-problems (4.11) become:

$$\forall i = 1, \dots, N_s, \begin{cases} -k^2 u_i - \Delta u_i &= 0 & \text{in } \Omega_i, \\ \nabla u_i \cdot \mathbf{n}_i + iku_i &= g_i & \text{on } \partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i}, \\ \nabla u_i \cdot \mathbf{n}_i + iku_i &= \lambda_{i,j} & \text{on } \Gamma_{i,j}, \quad \forall j \in \mathcal{N}_i, \end{cases} \quad (4.14)$$

with:

$$\begin{cases} \lambda_{i,j} &= -\lambda_{j,i} + 2iku_j, \\ \lambda_{j,i} &= -\lambda_{i,j} + 2iku_i, \end{cases} \quad \text{on } \Gamma_{i,j}, \quad \forall j \in \mathcal{N}_i. \quad (4.15)$$

These additional equations are based on (4.13):

$$\lambda_{i,j} = -\nabla u_j \cdot \mathbf{n}_j + iku_j = -(\nabla u_j \cdot \mathbf{n}_j + iku_j) + 2iku_j = -\lambda_{j,i} + 2iku_j.$$

They have been determined such that the continuity of the local solutions across the interface between the sub-domains is satisfied. These continuity equations on $\Gamma_{i,j}$ are:

$$\begin{aligned} u_i - u_j &= 0, \\ \nabla u_i \cdot \mathbf{n} - \nabla u_j \cdot \mathbf{n} &= 0. \end{aligned}$$

They ensure that each solution u_i is equal to the restriction on Ω_i of the solution of the global problem [41].

Finally, the problem consists in finding $u_i \in H^1(\Omega_i)$ and $\lambda_{i,j} \in H^{-1/2}(\Gamma_{i,j})$ such that:

$$\begin{aligned} \int_{\Omega_i} \left\{ -k^2 \bar{w}_i u_i + \nabla \bar{w}_i \cdot \nabla u_i \right\} d\Omega + \int_{\partial\Omega_i} ik \bar{w}_i u_i ds = \int_{\partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i}} \bar{w}_i g_i ds \\ + \sum_{j \in \mathcal{N}_i} \int_{\Gamma_{i,j}} \bar{w}_i \lambda_{i,j} ds, \quad \forall w_i \in H^1(\Omega_i), \end{aligned} \quad (4.16)$$

with:

$$\begin{cases} \int_{\Gamma_{i,j}} \bar{\xi}_{i,j} (\lambda_{i,j} + \lambda_{j,i} - 2iku_j) ds = 0, & \forall \xi_{i,j} \in H^{-1/2}(\Gamma_{i,j}), \\ \int_{\Gamma_{i,j}} \bar{\xi}_{j,i} (\lambda_{i,j} + \lambda_{j,i} - 2iku_i) ds = 0, & \forall \xi_{j,i} \in H^{-1/2}(\Gamma_{i,j}), \end{cases} \quad (4.17)$$

where $\xi_{i,j}, \xi_{j,i}$ are the test functions associated respectively to $\lambda_{i,j}, \lambda_{j,i}$. It is worth noting that the interface $\Gamma_{i,j}$ is described by two Lagrange multipliers and the derivatives of the solution are not explicitly apparent in the formulation.

4.3.3 Interface Problem and Iterative Solver

The FETI-2LM method [41, 57], which relies on the formulation (4.16) (4.17) is now introduced.

4.3.3.1 Algebraic Form

After the partitioning of each of the sub-domains into non-overlapping finite elements, the discretisation of the variational formulation (4.16) detailed in 2.3.2 leads to a matrix system of the form:

$$\mathbb{K}^{(i)} \mathbf{u}_i = \mathbf{b}_i + \sum_{j \in \mathcal{N}_i} \mathbb{B}_{i,j}^T \mathbb{M}_{i,j} \tilde{\boldsymbol{\lambda}}_{i,j} \quad (4.18)$$

where \mathbf{u}_i and $\tilde{\boldsymbol{\lambda}}_{i,j}$ are the vector of DOFs respectively corresponding to the pressure field in domain i and to the Lagrange multipliers in domain i on the interface $\Gamma_{i,j}$, $\mathbb{B}_{i,j}$ is the restriction operator from Ω_i to $\Gamma_{i,j}$, \mathbf{b}_i is due to the external source on $\partial\Omega_i$ and

$$(\mathbb{K}^{(i)})_{lm} = \int_{\Omega_i} \left\{ -k_0^2 \bar{\varphi}_l \varphi_m + \nabla \bar{\varphi}_l \cdot \nabla \varphi_m \right\} d\Omega + \int_{\partial\Omega_i} ik_0 \bar{\varphi}_l \varphi_m d\Gamma,$$

The additional systems read:

$$\begin{bmatrix} \mathbb{M}_{i,j} & \mathbb{M}_{i,j} \\ \mathbb{M}_{j,i} & \mathbb{M}_{j,i} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\lambda}}_{i,j} \\ \tilde{\boldsymbol{\lambda}}_{j,i} \end{bmatrix} = - \begin{bmatrix} 2ik_0 \mathbb{M}_{i,j} \mathbb{B}_{j,i} \mathbf{u}_j \\ 2ik_0 \mathbb{M}_{j,i} \mathbb{B}_{i,j} \mathbf{u}_i \end{bmatrix}, \quad \forall j \in \mathcal{N}_i, \quad (4.19)$$

where $(\mathbb{M}_{i,j})_{lm} = \int_{\Gamma_{i,j}} \bar{\varphi}_l \varphi_m d\Gamma$. Note that the shape functions are chosen such that:

$$\mathbb{M}_{i,j} = \mathbb{M}_{j,i}.$$

In the remainder of this manuscript, the interface problem is simplified by defining: $\boldsymbol{\lambda}_{i,j} = \mathbb{M}_{i,j} \tilde{\boldsymbol{\lambda}}_{i,j}$. Eq. (4.19) becomes:

$$\begin{bmatrix} \mathbb{I} & \mathbb{I} \\ \mathbb{I} & \mathbb{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}_{i,j} \\ \boldsymbol{\lambda}_{j,i} \end{bmatrix} = - \begin{bmatrix} 2ik_0 \mathbb{M}_{i,j} \mathbb{B}_{j,i} \mathbf{u}_j \\ 2ik_0 \mathbb{M}_{j,i} \mathbb{B}_{i,j} \mathbf{u}_i \end{bmatrix}, \quad \forall j \in \mathcal{N}_i, \quad (4.20)$$

where \mathbb{I} refers to the identity matrix.

4.3.3.2 Interface Problem

Using (4.18), the unknown \mathbf{u}_i can be eliminated from the system (4.20), yielding the following interface problem:

$$\mathbb{F}_I \boldsymbol{\lambda} = \mathbf{d}. \quad (4.21)$$

For conciseness, the expressions of these matrices are given in the case of two sub-domains only:

$$\begin{bmatrix} \mathbb{I} & \mathbb{I} - 2ik_0 \mathbb{M}_{1,2} \mathbb{B}_{2,1} (\mathbb{K}^{(2)})^{-1} \mathbb{B}_{2,1}^T \\ \mathbb{I} - 2ik_0 \mathbb{M}_{2,1} \mathbb{B}_{1,2} (\mathbb{K}^{(1)})^{-1} \mathbb{B}_{1,2}^T & \mathbb{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}_{1,2} \\ \boldsymbol{\lambda}_{2,1} \end{bmatrix} = -2ik_0 \begin{bmatrix} \mathbb{M}_{1,2} \mathbb{B}_{2,1} (\mathbb{K}^{(2)})^{-1} \mathbf{b}_2 \\ \mathbb{M}_{2,1} \mathbb{B}_{1,2} (\mathbb{K}^{(1)})^{-1} \mathbf{b}_1 \end{bmatrix}. \quad (4.22)$$

The extension to a larger number of sub-domains is straightforward. The matrix \mathbb{F}_I is square, complex and contains full blocks. It is of size N_I where N_I designates the number of unknowns on the interface. Solving this sub-structured problem by a direct method would be costly, an iterative method is therefore used. Each iteration of the interface problem requires a matrix-vector product and the direct solution of a local problem with a new right-hand side in every sub-domain.

For such a non-symmetric problem, Krylov methods with full orthogonalization like GMRES [147] or ORTHODIR [163] should be preferred [57]. The GMRES and ORTHODIR are mathematically equivalent [146] for arbitrary non-singular matrices (they lead to the same approximate solutions) but ORTHODIR requires the storage of an additional vector at each iteration compared to GMRES. However, in the framework of the FETI-2LM, compared to the size of the factors, the storage of an additional interface vector at each

iteration is not expensive. The advantage of using ORTHODIR is pointed out in [144] where the solution of multi-source electromagnetic problems is improved in terms of iterations by using the FETI-2LM algorithm and ORTHODIR with full recycling of the search directions.

4.3.3.3 Iterative Algorithm

To iteratively solve the interface problem (4.22), ORTHODIR [163] is used.

1. Initialisation:

Given an initial solution $\boldsymbol{\lambda}^0$, compute:

- the residual at the initial step $\boldsymbol{r}^0 = \boldsymbol{d} - \mathbb{F}_1 \boldsymbol{\lambda}^0$,
- the first basis vector $\boldsymbol{w}^0 = \boldsymbol{r}^0$,
- its product by \mathbb{F}_1 : $\boldsymbol{\eta}^0 = \mathbb{F}_1 \boldsymbol{w}^0$,
- and normalise $\boldsymbol{\eta}^0$ with $\alpha_1 = 1/\sqrt{(\boldsymbol{\eta}^0, \boldsymbol{\eta}^0)}$ and

$$\begin{cases} \boldsymbol{w}^0 &= \alpha_1 \boldsymbol{w}^0 \\ \boldsymbol{\eta}^0 &= \alpha_1 \boldsymbol{\eta}^0 \end{cases}$$

2. Iterations (for $k = 0, 1, 2, \dots$ until convergence):

(a) Compute the $k + 1$ solution:

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \alpha_2 \boldsymbol{w}^k,$$

where $\alpha_2 = (\boldsymbol{r}^k, \boldsymbol{\eta}^k)$.

(b) Compute the corresponding residual:

$$\boldsymbol{r}^{k+1} = \boldsymbol{r}^k - \alpha_2 \boldsymbol{\eta}^k.$$

(c) Check if $\|\boldsymbol{r}^{k+1}\| \leq \varepsilon$, stop; otherwise, continue.

(d) Compute matrix-vector product $\boldsymbol{\eta}^{k+1} = \mathbb{F}_1 \boldsymbol{\eta}^k$.

(e) Modified Gram–Schmidt (GS) [146]:

Initialisation for summation:

$$\boldsymbol{w}^{k+1} = \boldsymbol{\eta}^k,$$

For $l = 0, \dots, k - 1$, compute:

$$\beta_{kl} = -\frac{(\boldsymbol{\eta}^{k+1}, \boldsymbol{\eta}^l)}{(\boldsymbol{\eta}^l, \boldsymbol{\eta}^l)},$$

$$\boldsymbol{w}^{k+1} = \boldsymbol{w}^{k+1} + \beta_{kl} \boldsymbol{w}^l,$$

$$\boldsymbol{\eta}^{k+1} = \boldsymbol{\eta}^{k+1} + \beta_{kl} \boldsymbol{\eta}^l.$$

(f) Normalise $\boldsymbol{\eta}^{k+1}$ with $\alpha_1 = 1/\sqrt{(\boldsymbol{\eta}^{k+1}, \boldsymbol{\eta}^{k+1})}$ and

$$\begin{cases} \boldsymbol{w}^{k+1} &= \alpha_1 \boldsymbol{w}^{k+1} \\ \boldsymbol{\eta}^{k+1} &= \alpha_1 \boldsymbol{\eta}^{k+1} \end{cases}$$

(g) Store the product of the basis vector by \mathbb{F}_1 :

$$\boldsymbol{\eta}^{k+1} = \boldsymbol{\eta}^{k+1} + \sum_{l=0}^k \beta_{kl} \boldsymbol{\eta}^l.$$

Step 2 (f) makes sure that the norm of $\boldsymbol{\eta}^{k+1}$ is equal to 1. Without normalisation, for large number of iterations, the norm of $\boldsymbol{\eta}^{k+1}$ can exceed machine precision.

Most of the computations in this iterative process consist of scalar products and additions

except for the matrix-vector product for which a Helmholtz problem has to be solved in each sub-domain.

4.3.3.4 Computation of the Matrix-Vector Product

In practice, the matrix \mathbb{F}_I is not explicitly known as its assembly would be computationally inefficient. The product of \mathbb{F}_I by a vector only needs local data. In each sub-domain, the matrix-vector product 2 (d) is obtained by solving local Helmholtz problems as followed. At each iteration k , in sub-domain i ,

1. the solution of a local problem given by:

$$\mathbb{K}^i \mathbf{u}_i^k = \sum_{j \in \mathcal{N}_i} \mathbb{B}_{i,j}^T \boldsymbol{\lambda}_{i,j}^k$$

is easily obtained using the LU factors computed prior to the iterative process;

2. the solution is restricted to each interface:

$$\boldsymbol{\nu}_{i,j}^k = \mathbb{B}_{i,j} \mathbf{u}_i^k ;$$

3. $\boldsymbol{\lambda}_{i,j}^k$ is exchanged for $\boldsymbol{\lambda}_{j,i}^k$ with all neighbours;
4. the matrix-vector product can be computed:

$$(\mathbb{F}_I \boldsymbol{\lambda}^k)_{i,j} = \boldsymbol{\lambda}_{i,j}^k + \boldsymbol{\lambda}_{j,i}^k - 2ik_0 \mathbb{M}_{j,i} \boldsymbol{\nu}_{i,j}^k .$$

In other words, the matrix-vector product consists in local forward and backward substitutions followed by data exchanges between neighbouring processors.

4.3.3.5 Partitioning Algorithm

In order to evenly distribute matrix factorisation cost, the decomposition algorithm has to split the initial domain such that all sub-domains have similar numbers of DOFs. Moreover, the communication cost for the iterative solver increases with the number of interface unknown λ . Therefore, the ideal decomposition would also minimise the size of the interfaces between sub-domains. The goal is to balance the computation time among the processors (if each sub-domain is assigned to a processor). For simple problems and geometries, this decomposition can be done by hand. Realistic cases usually feature complex geometries and manual load balancing can become difficult. In the general case, this task can be achieved automatically by using graph partitioning algorithms available in packages such as METIS [100] and SCOTCH [37]. The idea is to model the mesh by a connectivity graph. This graph is then partitioned into several sub-graphs so that the number of unknowns assigned to each sub-domain is similar, and that the number of interface unknown assigned to each sub-domain is minimised. Optimal partitioning is key

to efficient simulations on high performance parallel machine and mesh decompositions are intensively studied using graph theory [100, 127, 45]. However, in this work, we do not go into the details of graph theory. Unless otherwise stated, the automatic graph partitioner METIS [100] is used to split the global domain in this work.

4.3.4 Summary

In practice, the solving procedure used is as follow:

1. Decomposition of the global mesh into N_s partitions.
2. Discretisation of the problem and construction of the matrices \mathbb{K}^i , $\mathbb{M}_{i,j}$, \mathbf{b}_i .
3. Assignment of a processor to each sub-domain.
4. Each sub-domain factorises \mathbb{K}^i to obtain \mathbb{L} and \mathbb{U} and compute its part of \mathbf{d} .
5. Initialisation of the solutions.
6. Solution of the interface problem using ORTHODIR involving at each iteration:
 - the solution of a local problem using the existing LU factorisation of \mathbb{K}^i ;
 - exchange of information (λ) between neighbouring sub-domains.
7. Solution of the sub-domain problems using the existing LU factorisation of \mathbb{K}^i and the solution of the interface problem.

4.4 Objectives

The objectives pursued in this thesis are twofold.

The first objective is to examine the performance of the FETI-2LM methods for solving large scale Helmholtz problems discretised using high order FEM methods. In the literature, domain decomposition methods as preconditioners have been examined in conjunction with the p -version of the FEM for solving positive-definite systems in [23, 136, 81, 24]. For wave propagation problems, in [45], a classical and an optimised Schwarz methods associated with a high-order DGM are investigated. For a 2D problem, the authors conclude that the efficiency in terms of iteration numbers is increased by using an optimised Schwarz method compared to the classical method. At the same time, the performance also improves with the approximation order in the DG method. In chapter 3, the p -FEM has been shown to yield superior approximation properties compared to a wave-based method for a range of problems. However, the applicability range of the method is limited in practice by the inadequacy of direct methods (due to large memory requirements) and iterative methods (due to erratic convergence) at high-frequency. Domain decomposition methods appear to be the method of choice to address this issue and allow solving large problems at high frequencies.

The second objective is to apply the domain decomposition method to large scale connected wave problems. To the best of the author's knowledge, this type of problem has

not been tackled before using domain decomposition methods. FETI-2LM appeared to be the most appropriate choice for flow acoustics, due to the non-symmetric nature of the transmission conditions.

Chapter 5

FETI-2LM and High-Order FEM

In this chapter, we consider the propagation of harmonic, linear sound waves in a uniform, quiescent medium described by the Helmholtz equation (see section 2.1). Convected propagation problems are investigated in the next chapter. The focus of this chapter is on the behaviour of the FETI-2LM solver combined with a high-order discretisation.

The structure of this chapter is as follows. In the first section, a 2D scalability analysis is given. In addition to the dependency on the mesh size, frequency, or number of sub-domains, the behaviour of the method with respect to the order of approximation is also assessed. In the second section, the costs associated to the FETI method is compared to a direct solver for a 3D wave problem. We look in particular at the evolution of the costs when varying the polynomial order and the number of sub-domains for a constant accuracy. The third section is dedicated to the extension of the method to the p -FEM with adaptive order (FEMAO). Finally, in the last section, the FETI-2LM method and the adaptive p -FEM are used to solve a 3D realistic problem.

5.1 Two-dimensional Scalability Analysis

In this section, we look at the influence of some key parameters such as the number of sub-domains, the frequency, the mesh size and the polynomial order of the discretisation on the convergence of the FETI-2LM solver described in the previous chapter, section 4.3.

5.1.1 Description of the Test Case

A 2D guided wave problem is considered. This test case has been used in previous studies [41, 57] to assess the behaviour of FETI methods with classical FEM. The computational

domain is taken to be a square of size $a \times a$ and is represented by a uniform and unstructured mesh of triangular elements of typical size h . A single plane wave:

$$u_{\text{ex}} = e^{-ikx} ,$$

where x is the direction of propagation, is generated using Robin boundary conditions (described in section 2.1.5.1) on the left and right boundaries. Hard wall boundary conditions are applied on the top and bottom boundaries such that the normal component of the acoustic velocity vanishes $\mathbf{v} \cdot \mathbf{n} = 0$ and $\nabla u \cdot \mathbf{n} = 0$.

The problems are discretised using p -FEM with a constant p throughout the whole domain. Adaptive orders and non-uniform meshes are considered in section 5.3. The numerical models are solved using the FETI-2LM method following the procedure described in section 4.3.4. A zero vector is chosen as initial solution. The computation stops when the normalised (using the right-hand-side) residual on the interface problem is below a fixed value $\varepsilon = 10^{-8}$. At the end of the iterative process, we also ensure that the normalised residual on the global problem has converged and is below a fixed value $\varepsilon_{\text{glob}} = 10^{-6}$. As the analytical solution is known for this simple test case, in addition to the interface residual and the global residual, the L^2 and H^1 relative errors on the solution in the whole numerical domain are also computed at the end of each iteration.

In the remainder of this document, we denote:


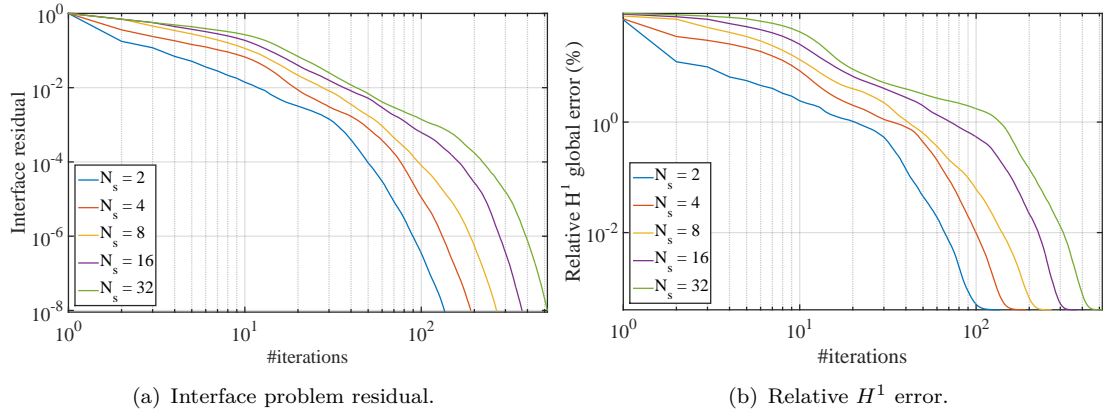
- $N_{\text{DOF}}^{\text{I}}$, the size of the interface problem;
- N_{DOF} , the number of DOFs in the whole numerical domain (after condensation unless stated otherwise);
- D_{λ} , the number of DOFs per wavelength;
- n_c , the number of iterations used to converge with a tolerance ε on the interface residual;
- r_{I} and r_{g} the normalised interface and global residual respectively;
- $L = \sqrt{\text{area}(\Omega)}$, the length of reference.

The code used has been validated against an existing one using a classical sequential direct solver in order to ensure that the solutions using both solvers were consistent.

5.1.2 Influence of the Number of Sub-domains

In this section, the behaviour of the method with respect to the number of sub-domains is studied. For a given Helmholtz number kL and polynomial order, a given mesh is partitioned into $N_s = 2, 4, 8, 16, 32$ using the Metis package [100] (see Fig. 5.1).

The results of the computations are given in Fig. 5.2 and Tab. 5.1. As the problems resolution remain the same, we observe that all four problems converge to the same solution with the same H^1 errors. The observation prevails for the L^2 error which

FIGURE 5.1: Example of different number of partitions for a given mesh $h/L = 1/60$.FIGURE 5.2: Influence of the number of sub-domains N_s . The interface residual and relative H^1 error against the iteration numbers. Fixed: $h/L = 1/60, kL = 100, p = 6, \varepsilon = 10^{-8}$.

N_s	N_{DOF}^I	N_{DOF}	D_λ	n_c	r_g	E_{L^2} it.	E_{H^1} it.
2	878	77144	26	135	3.696e-07	5.190e-05	4.006e-04
4	1630	77560	26	190	3.357e-07	5.190e-05	4.006e-04
8	3206	78320	26	267	2.269e-07	5.190e-05	4.006e-04
16	5320	79216	26	372	2.047e-07	5.190e-05	4.006e-04
32	8094	80738	26	518	1.333e-07	5.190e-05	4.006e-04

TABLE 5.1: Influence of the number of sub-domains N_s . Fixed: $h/L = 1/60, kL = 100, p = 6, \varepsilon = 10^{-8}$.

was not shown here for conciseness. Increasing the number of sub-domains induces a growing interface size. As the size of the interface grows, the amount of information being exchanged between neighbouring processors increases leading to larger memory requirements during the iterative process. This aspect is studied in more detail in the next section. The slight changes in N_{DOF} seen in Tab. 5.1 come from the duplication of the nodes at the interfaces. As the number of interfaces grows, the number of DOFs to be duplicated increases (the sub-domains are all independent and have their own version of a DOF at an interface). Therefore the global number of DOFs also increases.

On the convergence plot in Fig. 5.2 (b), for each curve, we can distinguish three different regimes. They will be denoted A, B and C from left to right.

- Regime A corresponds to the pre-convergence regime, which is characterised by the absence of, or by a very slow, convergence rate for both the interface residual and the relative error. The size of this pre-convergence regime is directly correlated with the number of sub-domains N_s . After one iteration, each sub-domain has exchanged information with its neighbours. After another iteration, each sub-domain indirectly communicate with the neighbours of their neighbours... and so on until the information has been successfully spread over the global domain. More precisely, the pre-convergence regime is directly linked to the number of layers of sub-domains between the inlet and outlet of the duct. The higher number of sub-domains, the longer the time for the information to travel between the sub-domains that are the furthest apart. A larger number of iterations is indeed observed for $N_s = 32$ compared to $N_s = 2$.
- Regime B corresponds to the convergence regime. In this regime, the convergence rates are much faster. The interface residual as well as the global error decrease rapidly with the number of iterations. In fact, within this regime, for each curve, we can distinguish two distinct convergence slopes. It is worth noting that the convergence rates of the interface unknown and error varies only slightly with N_s .
- Regime C corresponds to the post-convergence regime. In this regime, the global errors plateau. The error that dictates the accuracy of the solution is due to the discretisation of the problem and increasing the number of iterations (although it allows to improve the interface residual) does not yield any improvement on the global solution.

The consequence of increasing the number of sub-domains is the increase of the number of iterations to converge.

5.1.3 Influence of the Mesh Size

For a fixed order, frequency and number of partitions, we now investigate the effect of the mesh size on the convergence of the iterative method. The normalised residual of the interface is given in Fig. 5.3 along with the corresponding relative H^1 error on the global problem. The other variables are reported in Tab. 5.2.

h/L	N_{DOF}^I	N_{DOF}	D_λ	n_c	r_g	E_{L^2} it.	E_{H^1} it.
1/10	374	2241	4	131	1.476e-07	77.213	79.327
1/15	566	5143	7	150	1.788e-07	9.412e-01	1.408
1/30	1154	19685	13	193	1.360e-07	6.614e-03	2.532e-02
1/60	2222	77795	26	223	3.456e-07	5.193e-05	4.006e-04
1/80	3110	136659	34	238	3.959e-07	7.846e-06	7.562e-05

TABLE 5.2: Influence of the mesh size h/L . Fixed: $N_s = 5, kL = 100, p = 6, \varepsilon = 10^{-8}$.

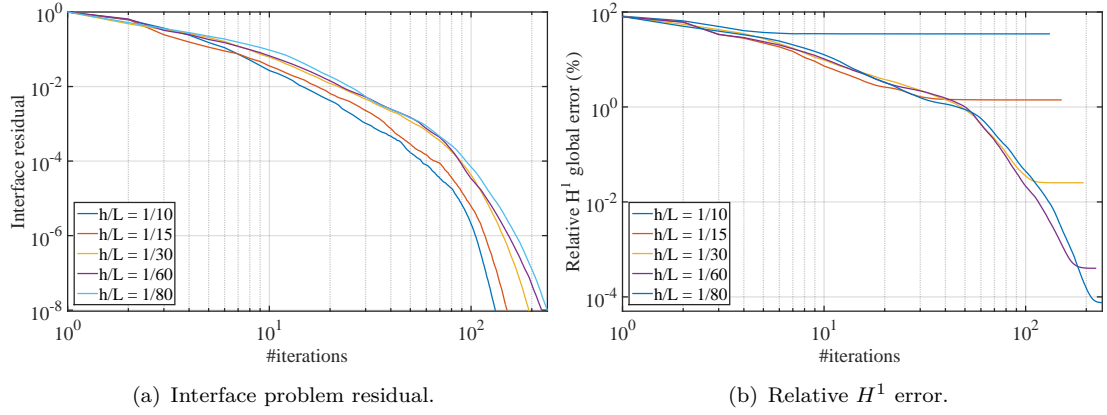


FIGURE 5.3: Influence of the mesh size h/L . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5, kL = 100, p = 6, \varepsilon = 10^{-8}$.

For each curve in Fig. 5.3, we observe the three different regimes mentioned above. Naturally, as the mesh size is decreased, the problem resolution becomes higher and the converged solution is more accurate. As the problem is better and better resolved, the number of iterations needed to reach the regime C becomes higher. The L^2 -error reached in this plateau is similar to the one achieved using a direct solver; we found an average relative difference of 10^{-10} . In other words, the solution obtained from the iterative solver is satisfactory and in this regime, the numerical error is dominated by the discretisation error.

To reach a given interface residual, the number of iterations varies slightly with the mesh size, even though no preconditioner is used. A similar conclusion is drawn for classical h -FEM [57, 41].

5.1.4 Influence of the Wavenumber

In this section, the dependence of the convergence behaviour with respect to the wave number is assessed. For $p = 6$, the results are given in Fig. 5.4 and Tab. 5.3.

As the Helmholtz number is increased, the problem resolution is lowered. Therefore, the L^2 and H^1 relative errors reached in the convergence regime are increased. We note that as the frequency is increased, the iteration counts to achieve a given residual increases slightly. These conclusions on the dependence with the frequency are consistent with the observations reported in the literature. However, contrary to the results found in [41, 57], the number of iterations does not drastically increase with the frequency.

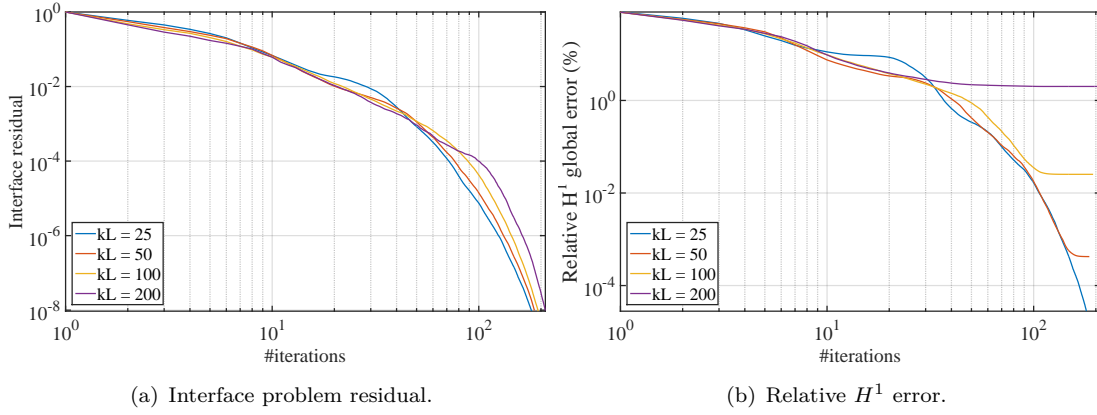


FIGURE 5.4: Influence of the wavenumber kL . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5, h/L = 1/30, p = 6, \varepsilon = 10^{-8}$.

kL	N_{DOF}^I	N_{DOF}	D_λ	n_c	r_g	E_{L^2} it.	E_{H^1} it.
25	1154	19685	53	179	5.981e-07	4.065e-06	2.793e-05
50	1154	19685	26	185	3.691e-07	5.506e-05	4.222e-04
100	1154	19685	13	193	1.486e-07	6.614e-03	2.532e-02
200	1154	19685	7	209	1.194e-07	1.680	2.008

TABLE 5.3: Influence of the wavenumber kL . Fixed: $N_s = 5, h/L = 1/30, p = 6, \varepsilon = 10^{-8}$.

5.1.5 Influence of the Polynomial Order

In order to understand the behaviour of the domain decomposition method with respect to the polynomial order, two different investigations are conducted. In the first study, except for the polynomial order, all the other parameters ($N_s, h/L, kL$) are kept constant. In the second study, all the parameters are kept constant except for the polynomial order and the mesh size which are chosen such that the resolution of the problem is fixed.

5.1.5.1 Varying the Polynomial Order Only

In this section, the triangulation and the frequency are fixed, whereas the polynomial order is varied. The results are given in Fig. 5.5 and Tab. 5.4.

To reach a given residual, the number of iterations increases with the polynomial order. A remark is that as long as the computation is in regime A and B, all curves corresponding to different p collapse into one. The curves differentiate only when the post-convergence plateaus of the global errors Fig. 5.5 (c) are reached i.e. when the regime C is reached. The main observations in terms of iteration counts are:

- to reach a given interface residual, the number of iterations varies with the polynomial order. This variation is larger for small polynomial order. In other words, as

we increase the polynomial order used, the difference of iteration counts between two consecutive order becomes smaller.

- to reach a given H^1 (or L^2) error – as long as one stays in the regime B – the number of iterations varies only slightly with the polynomial order.

On the one hand, the problem resolutions and solution accuracies are increased as we increase the polynomial order. On the other hand, the size of the interface problem and the local problems are larger for higher orders. Therefore, even though the number of iterations is only slightly increased, the final costs of the computation is higher for higher polynomial orders.

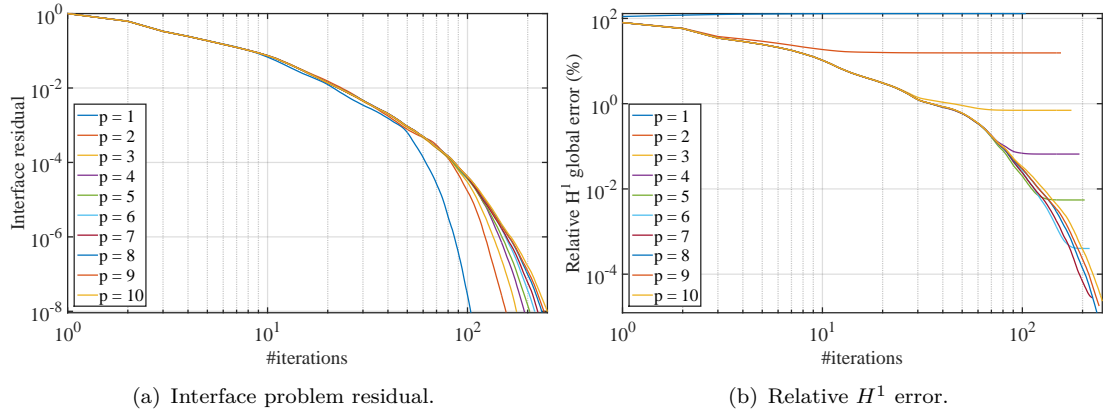


FIGURE 5.5: Influence of the polynomial order p . The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5, h/L = 1/60, kL = 100, \varepsilon = 10^{-8}$.

p	N_{DOF}^I	N_{DOF}	D_λ	n_c	r_g	E_{L^2} it	E_{H^1} it.
1	372	5055	5	103	9.100e-10	129.880	130.1
2	730	19609	9	155	2.310e-09	14.696	15.53
3	1088	34151	12	175	5.536e-09	2.790e-01	6.998e-01
4	1446	48699	18	192	1.420e-08	1.187e-02	6.579e-02
5	1845	63267	22	204	2.249e-08	8.221e-04	5.509e-03
6	2222	77795	26	223	3.456e-07	5.193e-05	4.006e-04
7	2581	92343	31	225	5.487e-08	5.186e-05	4.005e-04
8	2878	106891	35	235	6.993e-08	2.157e-06	1.245e-05
9	3236	121439	39	241	8.339e-08	2.362e-06	1.787e-05
10	3594	135987	44	250	1.468e-07	3.447e-06	2.324e-05

TABLE 5.4: Influence of the polynomial order p . Fixed: $N_s = 5, h/L = 1/60, kL = 100, \varepsilon = 10^{-8}$.

5.1.5.2 Varying the Polynomial Order for a Fixed Resolution

In this second study, the couple $(p, h/L)$ is chosen such that the resolution D_λ (3.3) of the problem is fixed and equals to 12. As the polynomial order is increased, the typical mesh size increases. High orders have a superior ability to control the pollution effect

when solving wave propagation problems. Therefore, for a fixed D_λ , the higher the polynomial order used, the more accurate the solution.

The results are reported in Fig. 5.6 and Tab. 5.5. It is observed that increasing the polynomial order at constant D_λ does not significantly increase the iteration counts to reach a given interface residual. For a given D_λ and N_s , the size of the interface problem remains similar for all the polynomial orders investigated. In terms of iteration counts and number of Lagrange multipliers, there is no advantage in using high-orders. However, thanks to condensation, the sizes of the local and global problems are reduced with larger p which indicates that using higher-order FEM leads to improvements in terms of memory requirements.

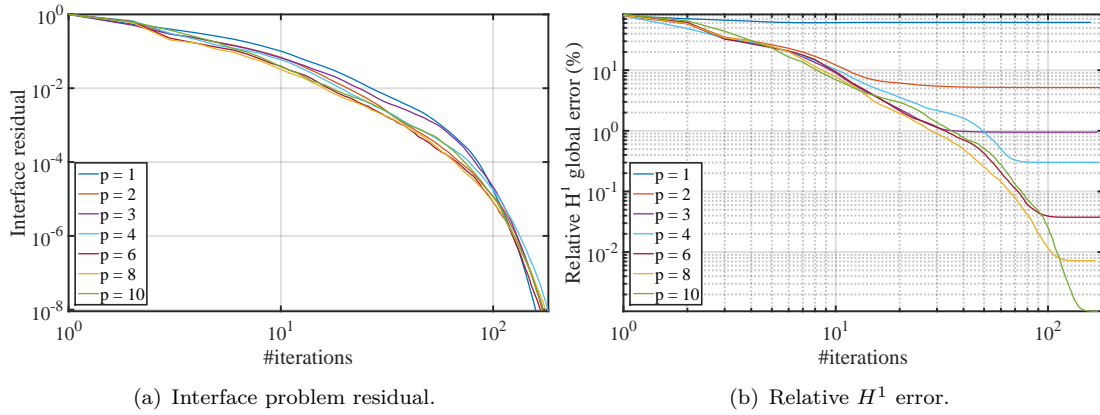


FIGURE 5.6: Influence of the polynomial order p at $D_\lambda \approx 12$. The residual and relative H^1 error against the iteration numbers. Fixed: $N_s = 5, kL = 100, \varepsilon = 10^{-8}$.

p	h/L	N_{DOF}^I	N_{DOF}	n_c	r_g	E_{L^2} it	E_{H^1} it.
1	1/165	1054	36935	158	1.779e-08	60.780	61.505
2	1/83	1038	37229	180	3.238e-08	4.401	5.187
3	1/55	998	28953	170	6.656e-08	4.482e-01	9.505e-01
4	1/41	968	23079	172	7.423e-08	8.617e-02	3.018e-01
6	1/28	1022	17421	181	1.229e-07	9.574e-03	3.753e-02
8	1/21	974	13207	166	4.182e-07	1.976e-03	7.207e-03
10	1/17	968	11813	171	1.450e-06	2.918e-04	1.055e-03

TABLE 5.5: Influence of the polynomial order p . Fixed: $N_s = 5, D_\lambda \approx 12, kL = 100, \varepsilon = 10^{-8}$.

5.2 Comparison of Computational Costs

The costs associated to the FETI-2LM method for solving a guided wave problem is assessed in this section. One MPI process of this solver is used for each sub-domain, and each MPI process is allocated a separate node (64GB, single thread) on an HPC cluster with InfiniBand interconnect.

Firstly, the distribution of the computational expense over the sub-domains/processors is examined. Then, for a given frequency and accuracy, we compare the costs of the FETI-2LM with a direct solver for increasing polynomial orders and number of sub-domains.

The test case used for this comparison is similar to the one described in section 5.1. However, the 3D problem is considered: the numerical domain is a cube of side a . On the boundaries situated at $x = 0$ and $x = a$, characteristics boundary conditions are used (see section 2.1.5.1) to generate a plane wave travelling along the x -direction. On the remaining boundaries, hard wall conditions are imposed.




FIGURE 5.7: Guided wave problem in 3D. Example of solution at $kL = 200$ (left) and mesh with $h/L = 1/26$ and $N_s = 8$ (right).

5.2.1 Cost Distribution

The aim of this study is to assess the quality of the partition. As mentioned in the previous chapter, the ideal partition ensures a uniform distribution of the computational expense across the processes and a minimal amount of information exchanged between neighbouring sub-domains.

Let us look at the cost distribution for the guided wave problem at $kL = 200$, $p = 8$, $N_s = 8$ and $E_{L^2} \approx 1\%$ given in Fig. 5.8. The mesh is partitioned into N_s sub-domains such that the number of DOFs per sub-domain is uniform (as seen on Fig. 5.8(a)) and the global number of interface unknowns is minimised. We note that the number of Lagrange multipliers is not quite similar for every sub-domain. The reason is that some sub-domains are only surrounded by other sub-domains. Therefore, these sub-domains only have boundaries which are interfaces whereas other sub-domains share the external boundaries. However, as the number of sub-domain is increased this disparity is attenuated. The size of the vectors stored at each iteration corresponds to the number of Lagrange multipliers: the lesser interface variables, the lesser the memory requirements. The number of Lagrange multiplier is also directly linked to the amount of information

that has to be exchanged at the end of each iteration. In our cases, the number of Lagrange multipliers is small compared to the local number of DOFs. In the framework of this study, we have not looked at the strategies and advantages of having a more uniform distribution of the number of Lagrange multipliers. Fig. 5.8 (c)-(d) represent the N_{NZ} and factorisation memory needed as a function of the process number. Uniform local N_{DOF} and N_{NZ} (directly linked to the memory required to store the matrices) are observed. As a consequence, the requirements in terms of factorisation memory is similar for all processes (see Fig. 5.8 (d)).

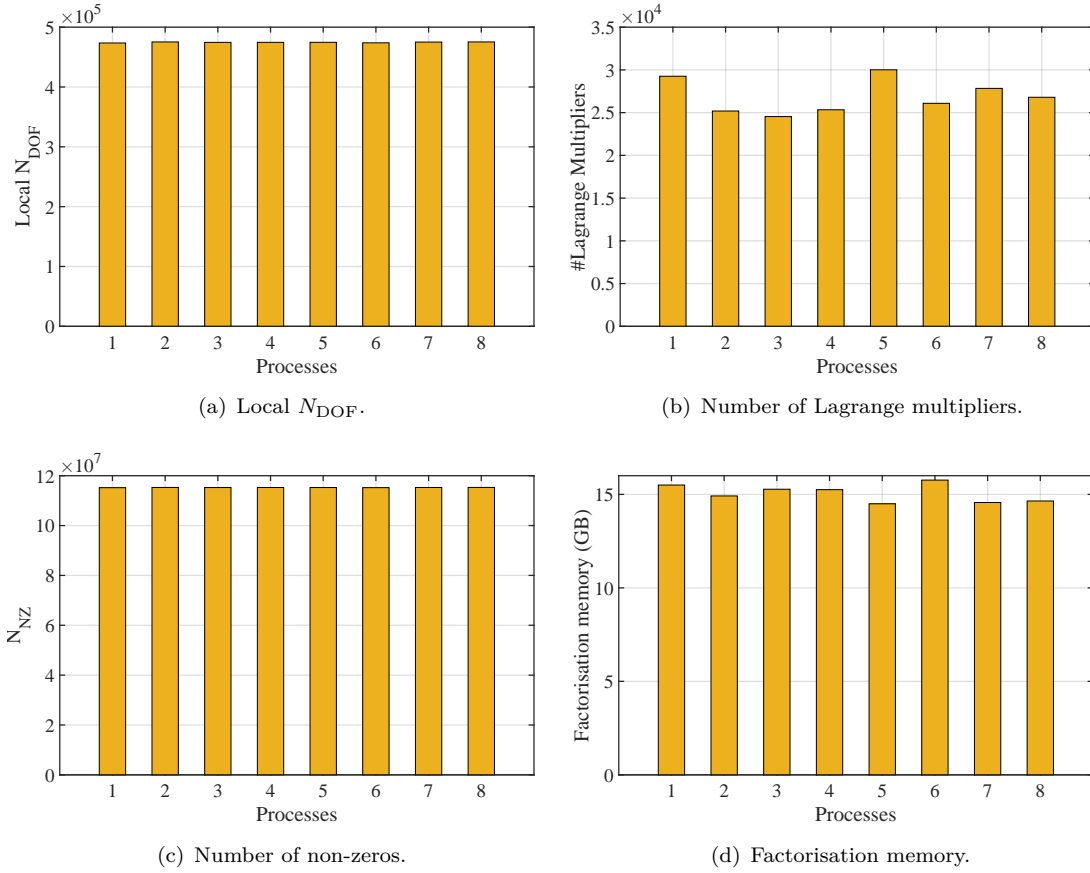
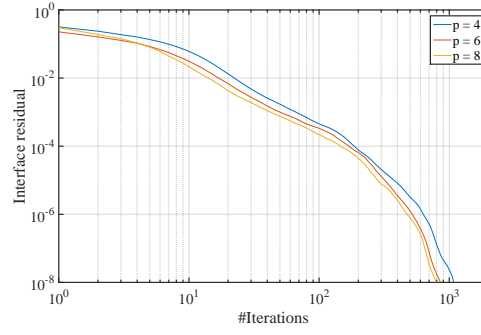




FIGURE 5.8: Guided wave problem at $kL = 200$ and mesh used to obtain $E_{L^2} \approx 1\%$ at $p = 8$ with $N_{\text{DOF}} = 5$ million, $N_s = 8$, $\epsilon = 10^{-8}$.

5.2.2 Influence of the Polynomial Order

The costs for solving the problem at $kL = 100$ (corresponding to 15.9 wavelengths along the edge of the cube) with 8 sub-domains are reported in Fig. 5.9. The problem was solved using $p = 4, 6, 8$. The results are compared to the ones given by MUMPS [1] ran in parallel for the exact same number of processes and configuration (mesh, frequency and order) and given in Fig. 5.10.



(a) Interface residual for the FETI-2LM method.

(b) Local N_{DOF} .

(c) Number of Lagrange multipliers.

FIGURE 5.9: Characteristics of the FETI-2LM computation. Guided wave problem at $kL = 100$, $N_s = 4$, $\varepsilon = 10^{-8}$ with $E_{L^2} \approx 1\%$; $E_{H^1} \approx 2\%$.

As suggested by the previous 2D results, Fig. 5.9 (a) shows that to achieve a solution with a given accuracy, the number of iterations to reach a residual of 10^{-8} only varies slightly with the polynomial order. In terms of pure iteration counts, there is no significant gain in using higher polynomial orders. The convergence rates, in every regime, for all the studied polynomial orders, are similar.

Fig. 5.9 (b) shows the maximum, mean and minimum of the local N_{DOF} over the sub-domains, for each polynomial order. For a given polynomial order, a uniform distribution of the number of DOFs is observed across the processes: all the sub-domains have similar problem sizes. As pointed out in the first two chapters of this report, at constant error, increasing the polynomial order leads to a reduction the total number of DOFs needed. Our results here are consistent with this observation.

In Fig. 5.9 (c), we note that the average number of Lagrange multipliers is also decreased when using higher polynomial order. Considering that the number of iterations does not increase with the polynomial order, using higher polynomial order is advantageous as the memory required to store the vectors for ORTHODIR is reduced and the total amount of information exchanged between processors is reduced.

The costs related to the direct solver and the hybrid solver are reported in Fig. 5.10. Fig. 5.10 (a-b) compares the maximum of factorisation memory needed by the processes and

total time for the FETI-2LM method (in different shades of yellow) and for MUMPS (in red). The total time refers to the factorisation and solving time. The assembly time was not considered in our study. Firstly, for both solver, keeping N_s constant, an increasing order lead to a decrease of both maximum of factorisation memory and computation time. For a given number of processors and polynomial order, FETI-2LM requires less memory but similar computation time compared to MUMPS.





FIGURE 5.10: Comparison of the FETI-2LM with MUMPS. Guided wave problem at $kL = 100$, $N_s = 4$, $\varepsilon = 10^{-8}$ with $E_{L^2} \approx 1\%$; $E_{H^1} \approx 2\%$.

5.2.3 Influence of the Number of Processes/Sub-domains


The polynomial order and triangulation are now fixed: $p = 8$ and $h/L = 1/26$. We vary the number of partitions $N_s = 8, 16, 32, 64$ and give the results in Fig. 5.11 and Fig. 5.12.

Fig. 5.11(a) gives the convergence plots for the FETI-2LM. As observed in the 2D case, the number of iterations increases with the number of sub-domains. Fig. 5.11(b) shows the memory required by the most demanding processor for both solvers for each N_s . The memory includes the factorisation and storage memory required by ORTHODIR. Limited by the computational memory available, the single domain solution with 8 processors was not achievable using MUMPS. However, an estimation of the memory requirement indicated by MUMPS is reported in Fig. 5.11(b). The memory demanded by MUMPS using a single processor was also estimated. The reference in Fig. 5.11(b) correspond to the quotient of this estimated memory and the number of sub-domains. The factorisation memory required by the most demanding processor decreases as we increase the number of sub-domains, for both solvers. We observe that in terms of memory the FETI-2LM performs better than the linear scaling. However, the memory required by the FETI-2LM for storing the vectors for ORTHODIR increases with the number of sub-domains. This effect is explained by the larger number of interfaces for large N_s . For the FETI-2LM we observe that there is a trade-off between the factorisation memory and the storage memory for ORTHODIR. For example, we notice that for


$N_s = 64$, the solver needs three times more storage memory than factorisation memory. Fig. 5.11(c) gives the factorisation and solution time needed by both solvers for the different N_s . Using the FETI-2LM as such does not lead to any significant gain in terms of time compared to MUMPS.



(a) Interface residual.




(b) Max used memory.




(c) Total time.

FIGURE 5.11: Influence of the number of sub-domains N_s with $p = 8$. Guided wave problem at $kL = 200$, $\varepsilon = 10^{-8}$ with $E_{L^2} = 3.05\%$; $E_{H^1} = 4.56\%$.



(a) MUMPS.



(b) FETI-2LM.

FIGURE 5.12: Distribution of the memory requirement over the $N_s = 16$ processes with $p = 8$. Guided wave problem at $kL = 200$, $\varepsilon = 10^{-8}$ with $E_{L^2} = 3.05\%$; $E_{H^1} = 4.56\%$.

5.3 FETI-2LM and an Adaptive Order FEM

5.3.1 About FEMA0

In this section, the FETI-2LM is used to solve problems discretised by an adaptive order p -FEM. The ability of adjusting the order of each element separately is profitable in many situations. For example, in the case of strongly non-uniform flows, in regions of high velocity, the acoustic wavelength can be significantly reduced by the convective effect. With a fixed order FEM, the mesh has to be strongly refined in those regions. If an adaptive order method is used, the mesh can be kept relatively uniform while the polynomial order is locally changed. The main advantage of adapting the order remains for frequency sweeps. When solving a problem for a wide range of frequency, the standard FEM calls for finer meshes at higher frequencies. Indeed, each frequency requires different mesh resolution in order to tackle the dispersion error. However, generating a new mesh for each frequency is time consuming and keeping the finest mesh for the complete range of frequency is computationally inefficient. Instead, for the whole frequency range, an adaptive order method can be used on a unique and carefully chosen mesh (for which the geometrical error is controlled).

As mentioned previously, one advantage of using hierarchical shape functions is that the basis can easily be varied across the mesh without requiring any specific treatment. From the order defined for each element, a unique interpolation order can be assigned to each edge and to each face of the mesh. In recent works [139, 21, 72], an error indicator is presented for the Helmholtz equation. The numerical error on a given element is determined *a priori*, based on the numerical error measured on a single, one-dimensional element with equivalent wave number k and length h , and flow properties if convective phenomena are included. In theory, each element is first assigned an order $p = 1$. On the corresponding 1D element, characteristics boundary conditions are used to generate a propagating plane wave inside this element domain. This problem is then solved and the numerical error is measured. If the 1D solution accuracy is lower than the target accuracy, the order p is increased and another 1D problem is solved. This process is repeated until the numerical error incurred on the element is lower or equal to the expected error. When the target accuracy is reached, the polynomial order is stored and used to solve the actual problem. In practice, the 1D problem is not solved for every element. The order is tabulated as a function of the mesh resolution kh and the target error. And this table is used to directly determine the order needed based on the local values of k , h and flow properties [21]. In between two elements, the maximum rule is applied: the polynomial order assigned to a given face or edge is equal to the maximum order of all adjacent elements sharing this face or edge. This *a priori* error indicator was shown to efficiently control the global numerical error on realistic two-dimensional and three-dimensional Helmholtz problems.

5.3.2 Domain Decomposition and Adaptive Order

Compared to the fixed order FEM, solving a problem using FEMAO and the FETI-2LM require a few additional steps. For the solution of the problems described below, the global computational mesh is first generated using Gmsh [77]. The polynomial orders of the elements are then assigned using the *a priori* error estimator [72]. With METIS [100], the mesh can be split into N_s sub-domains. As mentioned in section 4.3.3.4, we aim at minimising the size of the interface problem and balancing the number of DOFs in each sub-domains. For a three dimensional problem, the number of DOFs in an element scales roughly with $\mathcal{O}(p_e^3)$, where p_e is the local element order. Consequently, in order to improve the load balancing, a weight of p_e^3 is attributed to each element. After the partitioning, the orders of the edges and faces are carefully assigned using the maximum rule. The solutions at the interface between two sub-domains is averaged in order to build the global solution. For simplicity, the solution at the interfaces should be interpolated using the same basis. Therefore, when assigning the orders, two physically identical edges or faces belonging to different sub-domains are enforced with the same order.

After these preliminary steps, the sub-problems can be built and solved following the steps (2-7) of section 4.3.4. This new solving procedure is summarised below.

1. Construction of the global mesh.
2. Assignment of the order of each element using the *a priori* error estimator.
3. Decomposition of the mesh into N_s sub-domains, using the element orders to compute the weight of each element.
4. Discretisation of the problem and construction of the matrices \mathbb{K}^i , $\mathbb{M}_{i,j}$, \mathbf{b}_i .
5. Assignment of a processor to each sub-domain.
6. Each sub-domain factorise \mathbb{K}^i and compute its part of \mathbf{d} .
7. Initialisation of the solutions.
8. Solution of the interface problem using ORTHODIR involving at each iteration:
 - the solution of a local problem using the existing LU factorisation of \mathbb{K}^i ;
 - exchange of information (λ) between neighbouring sub-domains.
9. Solution of the sub-domain problems using the existing LU factorisation of \mathbb{K}^i and the solution of the interface problem.

5.3.3 Plane Wave Scattered by Sphere

5.3.3.1 Description of the Test Case

The test case used to validate the approach and implementation consists of the scattering of a plane wave by a rigid sphere. All quantities are non-dimensionalised, the speed of sound c_0 and the density ρ_0 are set to unity. The reference length L used is

the cubic root of the domain volume. The numerical domain is a cube of dimension $[-2a; 2a] \times [-2a; 2a] \times [-2a; 2a]$ containing a rigid sphere of radius a centered at the origin. An unstructured mesh of second order tetrahedral elements is used to represent the computational domain (see Fig. 5.13). The mesh size h_{\min} at the surface of the sphere is three times smaller than h_{\max} at the external boundary of the domain. h_{\min} is chosen such that at the surface of the sphere, a wavelength is at least described by six nodes. An incident plane wave of amplitude 1 is propagating in the $+z$ direction and is




FIGURE 5.13: Example of an unstructured mesh (left) with a refinement coefficient of $1/3$ and a PML (in red) and solution at $kL = 200$.

scattered by the rigid sphere. The analytical solution is given by [126]:

$$u_{\text{ex}} = e^{-ik_0 z} + \sum_{m=0}^N A_m H_m(k_0 r) \mathcal{L}_m \cos(\theta), \quad (5.1)$$

where H_m is the spherical Hankel function of the first kind of order m , \mathcal{L}_m is the Legendre polynomial of order m and

$$A_m = -(2m+1)i^m \frac{mJ_{m-1}(k_0 a) - (m+1)J_{m+1}(k_0 a)}{mH_{m-1}(k_0 a) - (m+1)H_{m+1}(k_0 a)}$$

with J_m the spherical Bessel function of order m .

To generate the solution, the normal velocity is imposed on the surface of the rigid sphere and a PML is used at the external boundary of the numerical domain to completely absorb the waves that are reflected into the domain (described in section 2.1.5.2). This test case was chosen as it features a simple geometry, an analytic solution, and a mix of wave directions.

5.3.3.2 Varying the Number of Sub-domains

For a given mesh ($h_{\max} = 1/30$) and a Helmholtz number $kL = 200$, we solve the sphere problem with a target accuracy $E_T = 5\%$. The polynomial orders used vary between 3 and 8. The solution is obtained using different numbers of partitions $N_s = 8, 16, 32$.

The mesh decompositions are given in Fig. 5.14. To represent the different partitions, the mesh at the interfaces between the sub-domains and at the surface of the sphere are plotted. Each colour corresponds to a different sub-domain. For clarity, a fictional space has been added in between the sub-domains.





FIGURE 5.14: Influence of the number of sub-domains. Decomposition used. Plane wave scattered by a sphere with $h_{\max}/L = 1/30$.


The costs associated to the computations are represented in Fig. 5.15. Fig. 5.15 (a) shows the convergence behaviours of the method using different numbers of partition. It is observed that varying the number of sub-domains does not significantly change the number of iterations required to reach a given residual ($\varepsilon = 10^{-8}$). We note in Fig. 5.14 that by increasing N_s , the number of sub-domains between the sphere where the velocity boundary condition is applied and the PML is not dramatically changed. This may explain the stagnation of the iteration count. This behaviour is different from the duct problem studied previously where the number of sub-domains between the inlet and outlet of the duct was increased with larger N_s .

Over the whole numerical domain, the number of DOFs is 10.9 million before condensation (approximately 8.7 million after condensation). The size of the interface problem increases with the number of sub-domains (see Fig. 5.15(b)). However, similar to the fixed order case, we observe that the number of Lagrange multipliers, DOFs (Fig. 5.15(c)) and number of non-zeros (Fig. 5.15(d)) is balanced amongst the sub-domains. This results in a balanced factorisation time and memory (Fig. 5.15(e-f)) and storage memory for the matrices and vectors in ORTHODIR.




By increasing the number of sub-domains, the workload (in terms of N_{DOF} or number of Lagrange multipliers for example) assigned to a process is diminished. Indeed, for a given problem, increasing N_s leads to smaller sub-matrices. These sub-matrices therefore require less factorisation time and memory. The same can be said about the solution time and memory. In this case, increasing N_s from 8 to 16 and 32 leads to smaller numbers




(a) Interface residual.



(b) Number of Lagrange multipliers.

(c) Local N_{DOF} .(d) Local N_{NZ} .

(e) Time.



(f) Memory.

FIGURE 5.15: Influence of the number of sub-domains. Costs. Plane wave scattered by a sphere with $kL = 200$, $h_{\max}/L = 1/30$, $E_T = 5\%$.

of Lagrange multipliers per sub-domain. Therefore, at each ORTHODIR iteration, each sub-domain has less information to handle and to exchange.

5.3.3.3 Varying the Frequency at Fixed Accuracy

In this section, we examine the influence of the dimensionless frequency ($kL = 50, 100, 200$) for a given mesh ($h_{\max}/L = 1/30$) and accuracy ($E_T = 5\%$). The number of sub-domains is fixed: $N_s = 8$. The mesh is plotted in Fig. 5.16 using the same principle as in Fig. 5.14. The domain partition and element orders used are also given along with the solution on the line resulting from the intersection of the plane $x = 0$ and the sphere. The angle 0° is taken at the point situated at the top of the sphere. For a given mesh, increasing the Helmholtz number leads to increasing polynomial order used. Indeed each frequency require different resolution in order to tackle the dispersion error. Looking at the solution on the sphere, we notice that the error indicator performs well as the relative L^2 error is consistently below 5%.

The behaviour of the method with increasing angular frequency is given in Fig. 5.17(a). We see that the method scales well for this problem, as the number of iterations needed to converge does not significantly increase with increasing kL .

The number of unknowns (Fig. 5.17(b-c)) is increased for $kL = 200$ compared to lower frequency. This is explained by the necessity of using higher order at higher frequency (as the mesh is fixed). The increasing number of DOFs has repercussions on the time and memory required for solving the problems. Doubling the frequency leads indeed to a factor 8 on the memory and time requirements.

5.3.3.4 Varying the Target Accuracy

We now look at the effect of the target accuracy on the behaviour and costs of the method (see Fig. 5.18). The Helmholtz number $kL = 200$, mesh $h_{\max}/L = 1/30$ and number of partitions $N_s = 16$ are kept constant. The target accuracy is varied with $E_T = 0.5, 5, 15\%$.

In Fig. 5.18 (a), we see that increasing levels of accuracy leads to a slight increase of the number of iterations to reach the target residual. However, this increase is not drastic. In the same way, increasing the frequency leads to increasing numbers of degrees of freedom, the number of λ , N_{DOF} and N_{NZ} is increased when the target error is decreased. We also see that decreasing the error level from 15% to 5% leads to an approximate factor 2 in the memory and time requirements whereas from 5% to 0.5%, this factor is closer to 3.

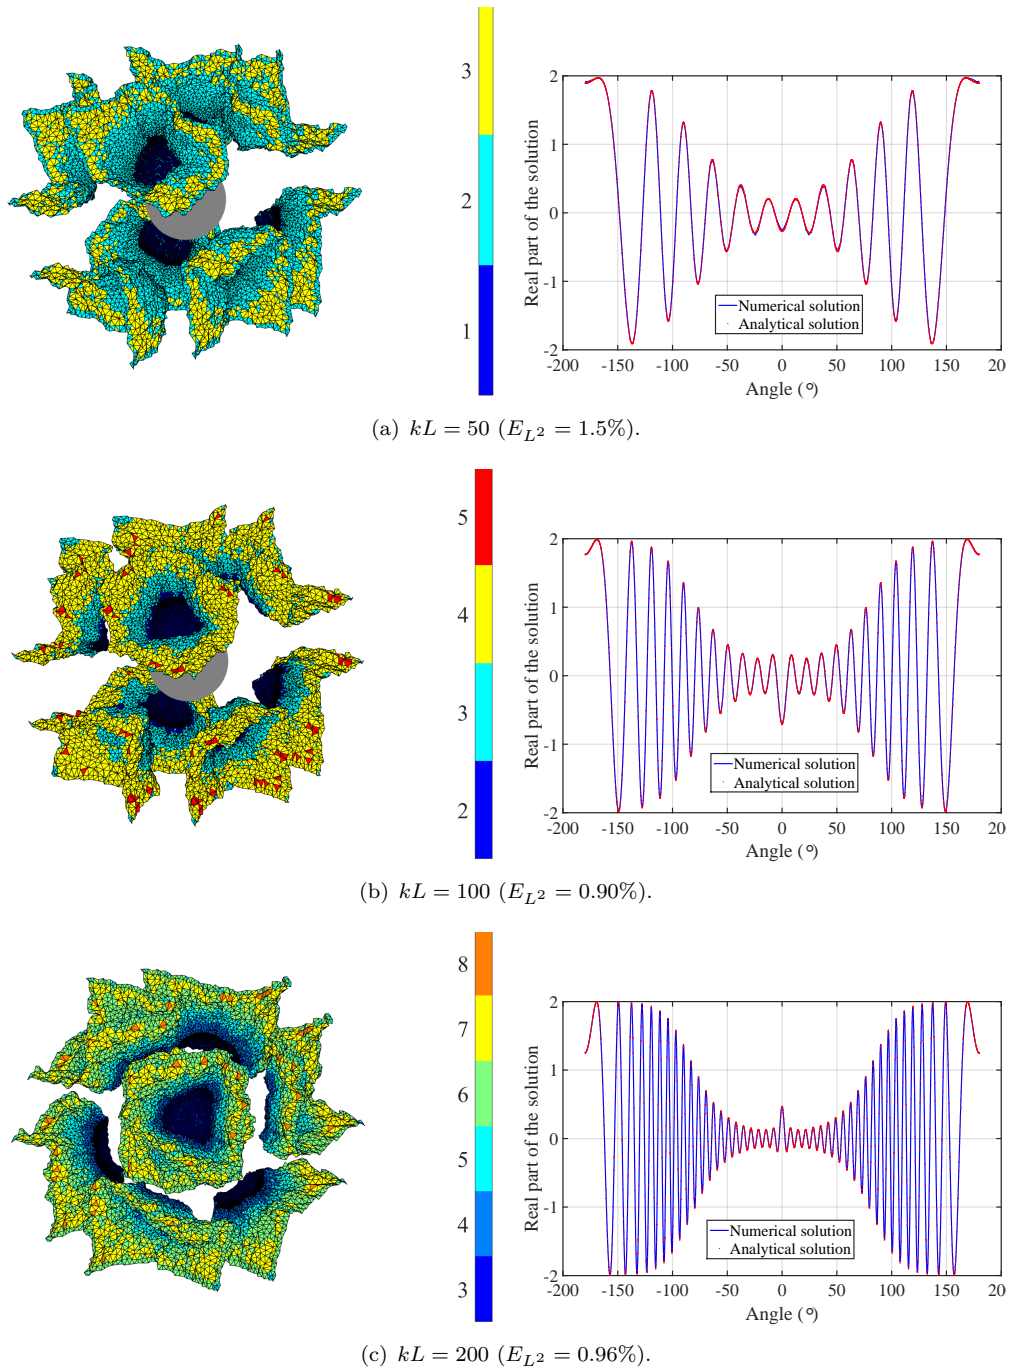







FIGURE 5.16: Influence of the Helmholtz number. Meshes and element orders (left) solution on sphere (right). Plane wave scattered by a sphere with $N_s = 8$, $h_{\max}/L = 1/30$, $E_T = 5\%$.




(a) Interface residual.



(b) Number of Lagrange multipliers.


(c) Local N_{DOF} .(d) Local N_{NZ} .

(e) Time.




(f) Memory.




FIGURE 5.17: Influence of the Helmholtz number. Plane wave scattered by a sphere with $N_s = 8$, $h_{\text{max}}/L = 1/30$, $E_T = 5\%$.




(a) Interface residual.



(b) Number of Lagrange multipliers.

(c) Local N_{DOF} .(d) Local N_{NZ} .

(e) Time.



(f) Memory.

FIGURE 5.18: Influence of the target accuracy. Plane wave scattered by a sphere with $kL = 200$, $h_{\max}/L = 1/30$, $N_s = 16$.

5.4 Realistic 3D Application: Plane Wave Scattered by a Submarine

5.4.1 Description of the Test Case

We consider the three-dimensional solution of an acoustic scattering problem. A plane wave is propagating in the direction given by $[\cos(\theta_0)\cos(\alpha_0); \sin(\theta_0)\cos(\alpha_0); \sin(\alpha_0)]$ and is scattered by an obstacle. The shape of the obstacle is inspired by the shark-proof submarine from the Belgian comic *The Adventures of Tintin*. The main body has a length of 5.2 m, a height of 1.25 m, and a maximal diameter of 1 m, see Fig. 5.19. The surrounding fluid is water which is characterised by a density and a speed of sound corresponding to $\rho_0 = 1000 \text{ kg/m}^3$ and $c_0 = 1500 \text{ m/s}$ respectively. The reference length used is the length of the main body: $L = 5.2 \text{ m}$.

The computational domain is a 6 m high cylinder with a radius of 1.50 m containing the submarine. On the surface of the obstacle, the normal velocity is imposed. A PML (described in section 2.1.5.2) is used to completely absorb the waves that are reflected into the numerical domain. An unstructured mesh of second order tetrahedral elements is used to represent the computational domain.




FIGURE 5.19: Shark-proof submarine, courtesy of Dav-dav (<https://grabcad.com/library/submarine-shark-1>).

5.4.2 Results

In this section the problem at $kL = 200$ with a target accuracy of 5% is solved using the FETI-2LM approach. The angles of the plane wave were set to $\theta_0 = 60^\circ, \alpha_0 = 45^\circ$. The mesh and polynomial order used and converged solution are represented on Fig. 5.20. The mesh sizes on the surface of the submarine vary from 10mm to 3mm. The

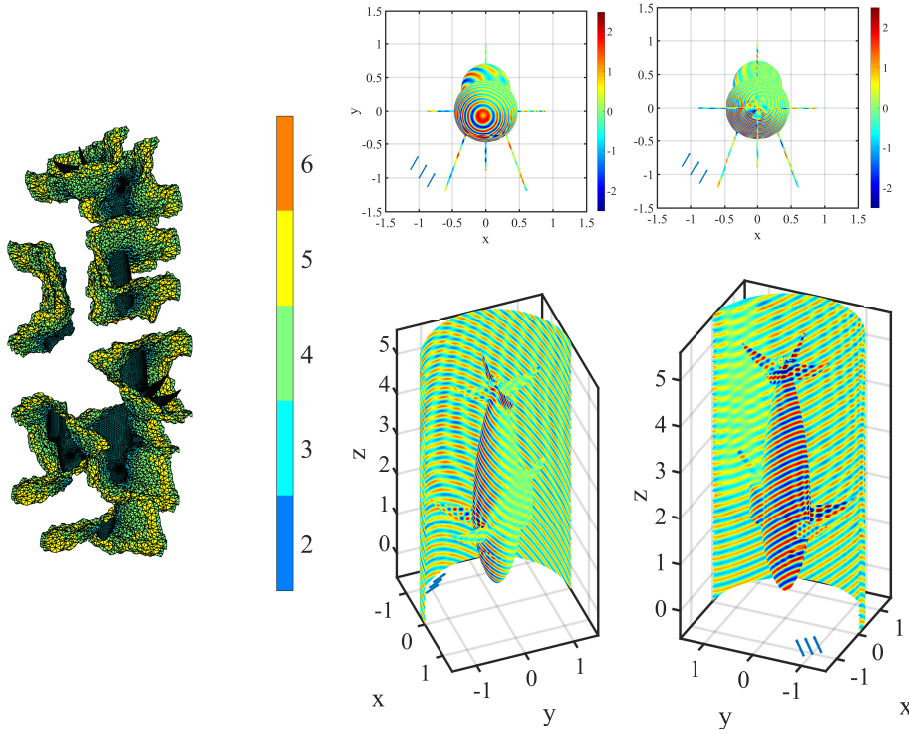



FIGURE 5.20: Mesh and order distribution (left) and solution at $kL = 200$ (center and right). The black arrow shows the direction of the incident plane wave.


elements at the external boundary measure approximatively 10cm. The polynomial orders required to solve the problem with an accuracy of 5% range from 2 to 6. The total size of the global system including the bubble shape functions amounts to about 4.5 million DOFs. After condensation of the internal degrees of freedom, the global matrix size is reduced to 4.3 million. The numerical domain was decomposed into 8 sub-domains resulting in local matrix sizes of about 540000 DOFs and in a balanced partition in terms of N_{NZ} (Fig. 5.21(b-c)) The most demanding processor took 30mins to factorise the local matrix. The solution time was 27 hrs and 22 mins with an iteration count of 14053 (Fig. 5.21(a)). For this particular problem where the geometry is more complex (small elements and details compared to the size of the whole domain) we notice that the method has difficulty converging. The pre-convergence regime is much longer compared to the previous test case, even though the size of the problems tackled are similar. The high number of iterations required leads to demanding computations: for most of the processes, the memory demanded by ORTHODIR is more than 50% of the total memory required for solving the problem.

The slow convergence may also come from the choice of splitting and would surely benefit from more optimal transmission conditions and the use of coarse spaces preconditioners [46, 76, 118, 40]. In [114], for the Helmholtz equation, the authors investigate algebraic optimization techniques of the interface operator for the FETI-2LM method using, for one sub-domain, approximations of the Schur complement of the outer domain. One of the method was then preconditioned and applied to a submarine problem similar to the




one described above. They find that in the case of a one-way splitting — meaning that the sub-domains are created by slicing the initial domain in one direction only — the model converges with $(N_s - 1)$ iterations at most and they retrieve a scalability with respect to the number of sub-domains.



(a) Interface residual.



(b) Number of Lagrange multipliers.

(c) Local N_{DOF} .(d) Local N_{NZ} .

(e) Memory.

FIGURE 5.21: Submarine problem at $kL = 200$ with a target accuracy of 5%.

5.5 Summary

In this chapter, we focused on the convergence behaviour of the FETI-2LM method when combined with a high-order discretisation method. A fixed order approach was

first adopted. The scalability analysis on a simple 2D duct problem showed the same trends that are observed in the literature for low-order FEM. For example, the FETI-2LM scales with the mesh size but not with the number of sub-domains [57, 41]. A surprising result is the effect of the Helmholtz number on the convergence of the method for the plane wave problem. Previous studies for low order FEM indicate that the method requires a coarse space to retrieve the scalability with respect to the frequency. In the results presented here, only a slight increase of the iteration counts was reported for increasing frequency. Different investigations were conducted in order to understand the behaviour of the domain decomposition method with respect to the polynomial order:

- keeping kL and h/L fixed, varying p ;
- keeping kL and D_λ fixed, varying h/L and p ;

For the first study, we find that increasing the order p leads to larger numbers of iterations required to reach a given residual. On the contrary, when the resolution is kept constant, the iteration count is not significantly impacted by the variation of approximation order.

When solving a 3D duct problem, the costs associated with the hybrid solver were compared to the costs induced by MUMPS. We showed the advantages of using the FETI-2LM method in terms of memory requirements. However, the solution time for both solver were similar. Resorting to higher-order discretization at fixed accuracy reduces the overall costs (time and memory) of a simulation. The number of iterations has been shown not to increase with the approximation order.

The FETI approach was then used in combination with an adaptive order FEM. For the sphere test case, the scalability of the method with respect to the number of partitions was retrieved. At fixed accuracy, increasing the frequency does not slow down the convergence. Varying the accuracy leads to slight changes in the iteration counts. In both cases, the variation of the frequency at fixed accuracy or the variation of the accuracy at fixed frequency significantly increase the costs of a single iteration.

Finally, the method was applied to the scattering problem by a 3D complex obstacle. In that case, we observed a slow convergence of the method that are thought to be induced by the non-uniformity of the mesh (presence of strong mesh refinements) and/or the elongated shape of the geometry.

Chapter 6

FETI-2LM for the Linearised Potential Equation

The behaviour of a domain decomposition method combined with a high-order discretisation is assessed for flow acoustics. The chapter is structured as follows. In a first part, the propagation model for waves in a potential base flow is introduced, followed by a description of the FETI framework. In the second part, a scalability analysis is conducted for a 2D problem. Finally, the method is used to predict the noise radiated from a generic turbofan intake.

6.1 Introducing the Flow

The prediction of the propagation and radiation of noise from aero-engines is crucial for the acoustic optimisation of various components of the turbofan engines. One of the most widespread model for this purpose is the linearised potential equation solved in the frequency-domain with the standard finite elements. Even though this model does not include the refraction effect that would be induced by a sheared mean flow, it is well suited for applications such as the radiation of sound from turbofan intakes [7, 69].

6.1.1 Formulation of the Sub-problems

For an irrotational flow field, the convective acoustic waves can be described using the linearised potential equation given in section 2.1.3. In each sub-domain Ω_i defined by its boundary $\partial\Omega_i$, the sub-problem is:

$$\begin{cases} \rho_0 \frac{D_0}{Dt} \left(\frac{1}{c_0^2} \frac{D_0 \phi_i}{Dt} \right) - \nabla \cdot (\rho_0 \nabla \phi_i) & = 0 & \text{in } \Omega_i, \\ \nabla \phi_i \cdot \mathbf{n}_i + ik_i \phi_i & = g_i & \text{on } \partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i} \end{cases} \quad (6.1)$$

where ϕ_i is the acoustic velocity potential in Ω_i . Given the velocity potential, the pressure, velocity and density in the sub-domain are obtained using $u_i = -\rho_0 D_0 \phi_i / Dt$, $\mathbf{v}_i = \nabla \phi_i$, $\rho_i = -(\rho_0 / c_0^2) D_0 \phi_i / Dt$. Assuming an implicit time dependence $e^{i\omega t}$, the material derivative becomes $D_0(\cdot) / Dt = i\omega(\cdot) + \mathbf{v}_0 \cdot \nabla(\cdot)$. The second equation in (6.1) represents a Robin condition applied on the external boundary of the domain (see section 2.1.5.1), where \mathbf{n}_i refers to the unit outward normal to $\partial\Omega_i$. g_i is a given source term and the convected wavenumber:

$$k_i = \frac{\omega}{(\mathbf{v}_0 \cdot \mathbf{n}_i + c_0)}. \quad (6.2)$$

\mathcal{N}_i denotes the set of sub-domain indices that share an interface with Ω_i .

At the interface between two sub-domains i and j , transmission conditions have to be imposed such that $\Gamma_{i,j}$ acts as a non-reflecting boundary. As for the Helmholtz equation, Robin conditions are used:

$$\nabla \phi_i \cdot \mathbf{n}_i + ik_i \phi_i = \nabla \phi_j \cdot \mathbf{n}_i + ik_i \phi_j \quad \text{on } \Gamma_{i,j}, \forall j \in \mathcal{N}_i. \quad (6.3)$$

The use of other transmission conditions is possible and may lead to more efficient methods (see for example [76]). At the interface $\Gamma_{i,j}$, Lagrange multipliers $\lambda_{i,j}$ can be introduced to avoid the computation of the RHS normal derivatives:

$$\lambda_{i,j} = (1 - (\mathbf{M}_a \cdot \mathbf{n}_i)^2)(\nabla \phi_j \cdot \mathbf{n}_i + ik_i \phi_j), \quad (6.4)$$

where $\mathbf{M}_a = \mathbf{v}_0 / c_0$.

The sub-problem (6.1) becomes:

$$\forall i = 1, \dots, N_s, \quad \begin{cases} \rho_0 \frac{D_0}{Dt} \left(\frac{1}{c_0^2} \frac{D_0 \phi_i}{Dt} \right) - \nabla \cdot (\rho_0 \nabla \phi_i) &= 0 & \text{in } \Omega_i, \\ \nabla \phi_i \cdot \mathbf{n}_i + ik_i \phi_i &= g_i & \text{on } \partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i}, \\ \nabla \phi_i \cdot \mathbf{n}_i + ik_i \phi_i &= \frac{1}{1 - (\mathbf{M}_a \cdot \mathbf{n}_i)^2} \lambda_{i,j} & \text{on } \Gamma_{i,j}, \forall j \in \mathcal{N}_i, \end{cases} \quad (6.5)$$

with:

$$\begin{cases} \lambda_{i,j} &= -\lambda_{j,i} + 2ik_0 \phi_j, \\ \lambda_{j,i} &= -\lambda_{i,j} + 2ik_0 \phi_i, \end{cases} \quad \text{on } \Gamma_{i,j}, \forall j \in \mathcal{N}_i, \quad (6.6)$$

where $k_0 = \omega / c_0$. These additional equations are obtained using Eq. (6.4) and the interface condition in Eq. (6.5).

Finally, the variational problem consists in finding $\phi_i \in H^1(\Omega_i)$ and $\lambda_{i,j} \in H^{-1/2}(\Gamma_{i,j})$ such that:

$$\begin{aligned} \int_{\Omega_i} \left\{ -\frac{1}{c_0^2} \frac{\overline{D_0 \psi_i}}{Dt} \frac{D_0 \phi_i}{Dt} + \nabla \overline{\psi_i} \cdot \nabla \phi_i \right\} d\Omega + \int_{\partial\Omega_i} i \overline{\psi_i} \left\{ k_i (1 - (\mathbf{M}_a \cdot \mathbf{n}_i)^2) + k_0 (\mathbf{M}_a \cdot \mathbf{n}_i) \right\} \phi_i ds \\ + \int_{\partial\Omega_i} \overline{\psi_i} (\mathbf{M}_a \cdot \mathbf{n}_i) (\mathbf{M}_a \cdot \boldsymbol{\tau}_i) (\nabla \phi_i \cdot \boldsymbol{\tau}_i) ds = \int_{\partial\Omega_i \setminus \{\Gamma_{i,j}\}_{j \in \mathcal{N}_i}} \overline{\psi_i} (1 - (\mathbf{M}_a \cdot \mathbf{n}_i)^2) g_i ds \\ + \sum_{j \in \mathcal{N}_i} \int_{\Gamma_{i,j}} \overline{\psi_i} \lambda_{i,j} ds, \quad \forall \psi_i \in H^1(\Omega_i), \end{aligned} \quad (6.7)$$

with:

$$\begin{cases} \int_{\Gamma_{i,j}} \overline{\xi_{i,j}} (\lambda_{i,j} + \lambda_{j,i} - 2ik_0 \phi_j) ds = 0, & \forall \xi_{i,j} \in H^{-1/2}(\Gamma_{i,j}), \\ \int_{\Gamma_{i,j}} \overline{\xi_{j,i}} (\lambda_{i,j} + \lambda_{j,i} - 2ik_0 \phi_i) ds = 0, & \forall \xi_{j,i} \in H^{-1/2}(\Gamma_{i,j}), \end{cases} \quad (6.8)$$

where $\xi_{i,j}, \xi_{j,i}$ are the test functions associated respectively to $\lambda_{i,j}, \lambda_{j,i}$.

The additional equations are not explicitly dependent on the base flow, the flow characteristics have in fact been included in the definition of the Lagrange multipliers.

6.1.2 Algebraic Form

Each sub-domain is partitioned into non-overlapping finite elements. The discretisation of Eq. (6.7) leads to the following matrix system:

$$\mathbb{K}^{(i)} \boldsymbol{\phi}_i = \mathbf{b}_i + \sum_{j \in \mathcal{N}_i} \mathbb{B}_{i,j}^T \mathbb{M}_{i,j} \boldsymbol{\lambda}_{i,j}, \quad (6.9)$$

which is similar to the system (4.18) obtained in the no flow case. The differences lie in the construction of the elementary matrices:

$$\begin{aligned} (\mathbb{K}^{(i)})_{lm} = \int_{\Omega_i} \left\{ -\frac{1}{c_0^2} \frac{\overline{D_0 \varphi_l}}{Dt} \frac{D_0 \varphi_m}{Dt} + \nabla \overline{\varphi_l} \cdot \nabla \varphi_m \right\} d\Omega \\ + \int_{\partial\Omega_i} i \overline{\varphi_l} \left\{ k_i (1 - (\mathbf{M}_a \cdot \mathbf{n}_i)^2) + k_0 (\mathbf{M}_a \cdot \mathbf{n}_i) \right\} \varphi_m d\Gamma \\ + \int_{\partial\Omega_i} \overline{\varphi_l} (\mathbf{M}_a \cdot \mathbf{n}_i) (\mathbf{M}_a \cdot \boldsymbol{\tau}_i) (\nabla \varphi_m \cdot \boldsymbol{\tau}_i) d\Gamma, \end{aligned} \quad (6.10)$$

and the external sources on $\partial\Omega_i$. The additional systems read:

$$\begin{bmatrix} \mathbb{M}_{i,j} & \mathbb{M}_{i,j} \\ \mathbb{M}_{j,i} & \mathbb{M}_{j,i} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}_{i,j} \\ \boldsymbol{\lambda}_{j,i} \end{bmatrix} = \begin{bmatrix} 2ik_0 \mathbb{M}_{i,j} \mathbb{B}_{j,i} \boldsymbol{\phi}_j \\ 2ik_0 \mathbb{M}_{j,i} \mathbb{B}_{i,j} \boldsymbol{\phi}_i \end{bmatrix}, \quad \forall j \in \mathcal{N}_i, \quad (6.11)$$

where $(\mathbb{M}_{i,j})_{lm} = \int_{\Gamma_{i,j}} \overline{\varphi_l} \varphi_m d\Gamma$.

Thanks to the choice of suitable Lagrange multipliers (6.4), the form of the problem for the LPE is similar to the one derived for the Helmholtz equation. The systems are solved using the procedure described in section 5.3.2.

6.2 Two-Dimensional Study

The test case used to validate the approach and implementation consists of the scattering of a point source by a rigid cylinder.

6.2.1 Description of the Test Case

A cylinder is placed in a steady, two-dimensional, incompressible, inviscid and irrotational flow such that far enough from the cylinder, the velocity vector is aligned with the x -direction and uniform with a norm $\|V_\infty\|$. The expression of the velocity can be derived using the Laplace equation for its potential. In cylindrical coordinates, it is given by:

$$v_{0,r}(r, \theta) = V_\infty \left(1 - \frac{R^2}{r^2}\right) \cos(\theta), \quad v_{0,\theta}(r, \theta) = -V_\infty \left(1 + \frac{R^2}{r^2}\right) \sin(\theta), \quad (6.12)$$

where R is the radius of the cylinder. The flow induces zero velocities at the stagnation points ($r = R$ and $\theta = 0, \pi$) and a maximum velocity corresponding to $2\|V_\infty\|$ at the top and bottom parts of the cylinder ($r = R$ and $\theta = \pi/2, 3\pi/2$). In the computations performed here, the potential mean flow is prescribed analytically. The speed of sound c_0 and the density ρ_0 are set to unity. We use $R = 1$, $V_\infty = -0.4$ m/s and $M_\infty = V_\infty/c_0$. The resulting mean flow direction and magnitude is plotted in Fig. 6.1.

The location and amplitude of the point source are equal respectively to $(x_s, y_s) = (-R - 0.4, 0)$ and $A = 1$. It is imposed at the RHS of Eq. (6.1):

$$f_s(x, y) = A\delta(x - x_s, y - y_s),$$

where δ is a Dirac delta function.

The test case is solved on a square domain of length $2R$ containing the cylinder which is centered at the origin. A hard wall boundary condition is imposed at the surface of the cylinder ($\nabla\phi \cdot \mathbf{n} = 0$) and a perfectly matched layer (described in section 2.1.5.2 and shown in red in Fig. 6.1) is used to completely absorb the waves that may be reflected back into the numerical domain. An unstructured mesh of quadratic triangular elements is used to represent the computational domain (see Fig. 6.1). The mesh is refined on the surface of the cylinder. Compared to the typical mesh size h at the outer boundary,




FIGURE 6.1: Mean flow direction and magnitude with $M_\infty = -0.4$ (left); example of mesh $h = 0.125$ (right).




the mesh size on the cylinder is close to $h/3$. A refinement is also used around the point source where the solution is expected to be irregular: $h_s = h/3$.

6.2.2 Validation on Non-Homogeneous Flow

In the absence of analytical solution, we are comparing the solution obtained with the FETI method to the solution produced using the same mesh and a direct solver. The polynomial order distribution used for the hybrid and direct solver are identical for a given problem.

For a given number of sub-domains, three different frequencies are considered $\omega = 10, 25, 100$ for a target accuracy of $E_T = 1\%$ and $\varepsilon = 10^{-8}$. Accounting for the convective effect induced by the mean flow, these frequencies correspond respectively to wave number magnitudes of approximately 50, 125 and 500 in the regions where the mean flow magnitude is the largest. For each frequency, the typical mesh size is chosen such that the shortest wavelength at the surface of the cylinder is described by at least 6 nodes.

The solutions in the domain are given in Fig. 6.2 along with the real part of the solutions at each frequency on the perimeter of the cylinder (at $r = R$). While large wavelengths are found in the region downstream of the source, short creeping waves are observed along the cylinder curvature. We see that there is an excellent agreement between the solutions, even at high frequency and in the most oscillatory regions corresponding to $\theta = \pm\pi/2$.

(a) $\omega = 10$.(b) $\omega = 25$.(c) $\omega = 100$.FIGURE 6.2: Solutions in the domain (left) and solutions on the cylinder (right) with $E_T = 1\%$.

6.2.3 Scalability Analysis

A set of numerical experiments is now presented in order to assess the convergence of the proposed method for convective problems with increasing number of subdomains, Mach numbers, frequency or level of accuracy.

In a first experiment, the Mach number, frequency, mesh size and target accuracy are fixed. Only the number of sub-domains is varied and the results are computed in Fig. 6.3 and Tab. 6.1. The converged solution is represented in Fig 6.2 (b) for $N_s = 9$. As we increase the number of sub-domains, the size of the interface problem N_{DOF}^I increases markedly (by a factor 17 between $N_s = 2$ and 64), at the same time the number of iterations n_c also increases, but not as drastically (only a factor 2 between $N_s = 2$ and 64). It is also worth noting that the iterative procedure converges rapidly (i.e. n_c is always smaller than N_{DOF}^I), even in the presence of corners (points shared between three or more sub-domains).

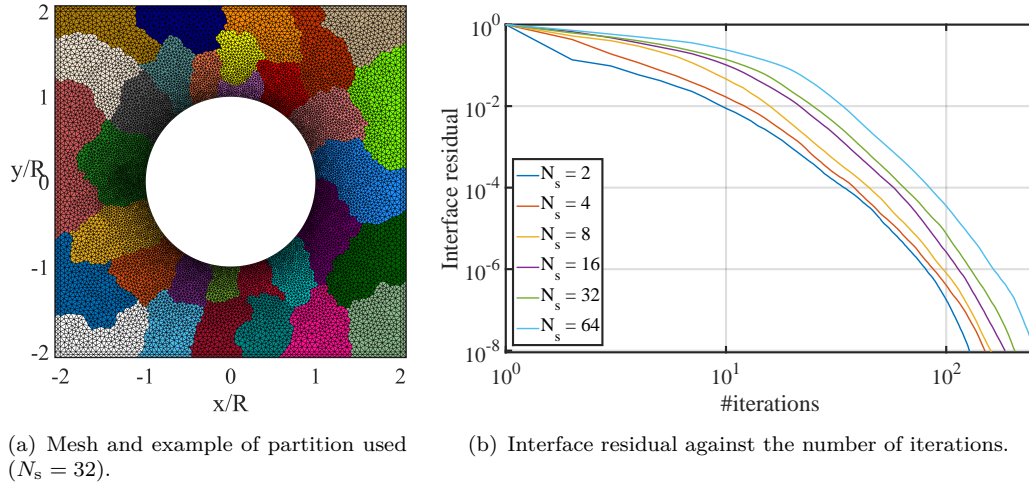


FIGURE 6.3: Influence of the number of sub-domains N_s with $M_\infty = -0.4$, $\omega = 25$, $h/R = 0.06$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$.

N_s	N_{DOF}^I	N_{DOF}	n_c	r_g
2	572	93298	127	1.920e-07
4	1042	93533	149	3.657e-07
8	2258	94141	159	2.656e-07
16	4100	95050	184	2.388e-07
32	6566	96263	209	1.837e-07
64	9990	97922	253	1.074e-07

TABLE 6.1: Influence of the number of sub-domains N_s with $M_\infty = 0.4$, $\omega = 25$, $h/R = 0.06$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$.

The second experiment consists in keeping all the parameters constant except for the Mach number M_∞ . The polynomial basis uses orders ranging from $p = 2$ to $p = 3, 4$ (see Fig. 6.4). The proportion of higher order is increased for larger velocities. In fact,

we note that as the flow velocity is increased, the orders used around the source and the stagnation regions are constant whereas the orders are increased in the regions where the flow accelerates. As the flow velocity is increased, the numbers of unknowns in the global problem and the interface problem also increase, see Tab. 6.5. However, the iteration counts only slightly increase (see Fig. 6.5).

The results for a third experiment where the frequency is increased while the Mach number is kept fixed is given in Fig. 6.6. The number of degrees of freedom increases as ω is increased, both for the global problem and for the interface problem. This is explained by the fact that in order to maintain the same level of error, the polynomial order in the elements is increased when the frequency is increased. An interesting observation is that the number of iterations decreases with the frequency.




FIGURE 6.4: Element orders for different Mach numbers M_∞ with $N_s = 8, \omega = 25, h = 0.06, E_T = 0.5\%$ and $\varepsilon = 10^{-8}$.

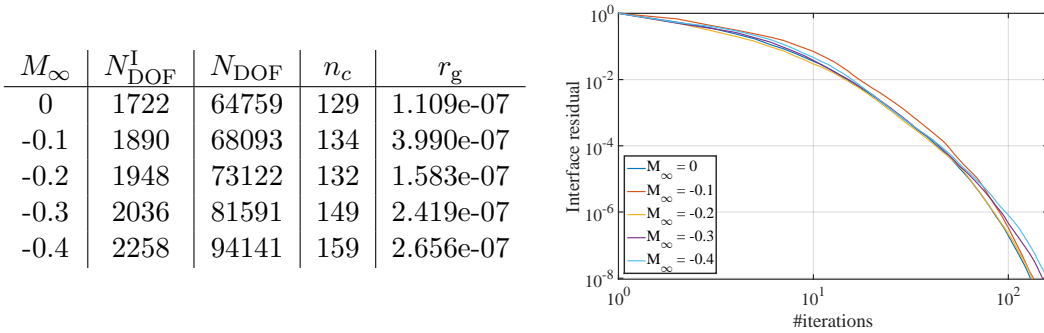


FIGURE 6.5: Influence of the Mach number M_∞ with $N_s = 8, \omega = 25, h = 0.06, E_T = 0.5\%$ and $\varepsilon = 10^{-8}$.

6.3 Realistic 3D Application

In this section, the methodology is applied to solve the problem of sound radiation from a generic turbofan engine intake.




FIGURE 6.6: Influence of the frequency with $M_\infty = -0.4$, $N_s = 8$, $h = 0.025$, $E_T = 0.5\%$ and $\varepsilon = 10^{-8}$.




FIGURE 6.7: Geometry of the engine intake (left) and local Mach numbers for the sideline case (right); the black dotted lines indicate the start of the PML.

6.3.1 Description of the Test Case

The axisymmetric generic engine intake geometry used in this work is given in Fig. 6.7. The computational domain consists of a cylinder of length 3.5 m and radius 2.5 m and a perfectly matched layer. The thicker part of the spinner and the fan have radii of $R_s = 0.36$ m and $R_f = 1.2$ m, respectively. The outer radius of the nacelle is 1.5 m. For the purpose of the mean flow calculation and the acoustic simulations, a straight, annular ‘fan duct’ is added downstream of the fan face. The noise source is introduced through a duct mode boundary condition which is defined on the source plane located at $z = 0$ m. A cylindrical PML formulation (presented in section 2.1.5.2) is used to efficiently absorb the outgoing waves at the outer boundary. More detailed information concerning this test case are available in [139, 128]. This geometry was previously used to benchmark the performance of the *a priori* p -adaptive strategy for flow acoustics in [69, 19].

For noise certification tests, three standard flight conditions are usually considered: sideline (at take-off when the engines are at full-power), cutback (at take off, when the engine power is reduced) and approach (landing). We focus here on the sideline condition which is a computationally demanding condition to solve as the Mach number




FIGURE 6.8: Generic intake at 1BPF, mode (1, 24): polynomial orders used (left); real part of the velocity potential (right).

at the fan face reaches $M_f = 0.55$. The flow field is strongly non-uniform within the domain and the local flow Mach numbers are given in Fig. 6.7. It was computed using a potential, compressible, inviscid flow model. The conditions at infinity are as follows: $c_\infty = 340$ m/s, $\rho_\infty = 1.2$ kg/m³ and $M_\infty = u_\infty/c_\infty = 0.25$.

6.3.2 Results

Tonal noise is present at the Blade Passing Frequency (BPF) and its harmonics. We consider the BPF of the fan $f = 1300$ Hz. Using the following definition of the Helmholtz number [156]:

$$kR_f = \frac{\omega R_f}{c_0(1 - M_f)}, \quad (6.13)$$

we find a Helmholtz number for the present problem of 64. Results are presented here for the first radial mode with azimuthal order 24 (see Fig. 6.8). The mesh consists of quadratic tetrahedral elements, with a characteristic length of $h = 20$ cm in the far field and $h = 0.4$ cm close to the walls of the nacelle. Following the procedure detailed in section 5.3.2, the *a priori* error indicator is used to assign the order in each element across the mesh with a target accuracy of 5 %. The orders range from $p = 2$ close to the nacelle where the mesh is refined to $p = 5 - 6$ in the far field and near the throat of the inlet where the Mach number exceeds 0.8. The mesh is then split into $N_s = 6$ sub-domains. There is one MPI process of this solver for each subdomain, and each MPI process is allocated a separate node (64GB, single thread) on an HPC cluster with InfiniBand interconnect. The total number of primal unknowns (does not include the number of Lagrange multipliers) is 8.2 million DOFs (7.8 million after condensation), and the global linear system contains more than 520 million non-zero entries. The total size of the interface problem is 202188.




FIGURE 6.9: Generic intake at 1BPF, mode (1,24): interface residual as a function of the number of iterations (left); memory requirements for each process (right).

Fig. 6.8 presents the real part of the velocity potential field obtained after recovering the solution in the $N_s = 6$ sub-domains. The convergence curve and the memory requirements per process are given in Fig. 6.9. The iterative procedure converged in 866 iterations to a relative residual of 10^{-8} in two hours and fifty minutes. On the one hand, the most computationally demanding task involves factorising the linear system for a sub-domain with 1.3 million DOFs, which was performed with MUMPS using 53.5GB. The total memory requirements on this process was 54.7GB and the local interface problem contained 31141 Lagrange multipliers. On the other hand, the least computationally demanding process required 48.8GB of RAM, which indicates a relatively good load balancing between the different processes.

6.4 Summary

In this chapter, the FETI-2LM approach has been applied to the propagation of sound with a background mean flow described by the linearised potential equation. It was shown that, through the choice of suitable Lagrange multipliers, it is possible to obtain an algebraic formulation that is similar to the formulation associated with the Helmholtz equation. The method has been validated and its scalability has been assessed with respect to the flow velocity, the frequency and the number of sub-domains. Finally, a 3D generic intake propagation problem was solved using the proposed FETI-2LM formulation.

Chapter 7

Conclusions

The aim of the Ph.D is to investigate efficient methods for the solution of frequency-domain aero-acoustics problems. The work is mainly divided into two parts involving:

1. Discretisation methods: comparison of a high-order polynomial method against a wave-based method in order to find the best performing method over a range of frequency for the solution of the Helmholtz equation;
2. Solvers: application of a domain decomposition method to the solution of aero-acoustic problems discretised with a high-order method.

Summary

Relative performance of a polynomial and a physics-based methods

The conventional finite element method is not suited when the problems imply short waves, due to the presence of cumulative phase errors. To circumvent this issue, various methods have been developed: for example, spectral methods, high-order polynomial methods using Bernstein, Lobatto or Hermite functions or physics-based methods. In this work, we compared a high-order polynomial method (p -FEM) and the wave-based discontinuous Galerkin method in order to determine the advantages of using each method over different ranges of frequency. These two methods use different approaches. The polynomial p -FEM exploits the superior interpolation properties of some families of polynomials (here, Lobatto functions). It is also well established in the numerical acoustic community and has been recently implemented in a commercial code. On the contrary, the wave-based DGM includes local solutions (plane waves) in the numerical method: plane waves are canonical solutions of the underlying equations and are therefore believed to be better suited than polynomials to construct an approximation basis. Previous studies have shown that both methods may lead to an effective control

of the pollution effect. However, to our knowledge, no detailed comparison can be found in the literature. In this work, a detailed comparison of the two methods is done in terms of interpolation properties, convergence behaviour and costs to achieve a solution with a fixed accuracy. The main observations are:

- For the smooth propagating wave problems, the wave-based DGM and p -FEM have been able to achieve the same levels of accuracy with similar levels of performance. To reach the accuracy required for typical engineering problem, contrary to what is usually said, the wave-based systems are not ill-conditioned. The condition numbers are in fact comparable to that of the p -FEM with condensation. In the case where flow is involved, the wave basis can be adjusted such that the short wavelengths are better described, leading to good performances.
- When dealing with evanescent waves, the wave-based DGM becomes expensive compared to the p -FEM which performs as well as for the propagating waves cases.
- The performance of the wave-based method for the singular problem was only satisfactory: with the given uniform mesh, it was hardly possible to get an L^2 -error below 1% by only increasing the number of plane waves in the basis. The p -FEM is more adapted to represent the non-smooth solution even though the p -convergence observed for smooth problems is lost. The remedy for the p -FEM is to refine the mesh around the singularity point. This approach was also applied to the wave-based DGM but the ill-conditioning of the matrices made it hard to achieve high accuracy solutions.

For a range of frequency and error levels, the p -FEM is able to achieve comparable, and in some cases superior, results compared to the wave-based method, thanks to the application of condensation. The studied physics-based method does not provide a step change in computational performance, even at high frequency compared to the polynomial method. More general aspects have also to be considered. For example, when applied to problems with non-uniform coefficient or involving source terms, the p -FEM does not require additional treatment contrary to the physics-based method as in those cases the plane waves are not solutions of the governing equations. In practice, for industrial applications, the use of the physics-based method is therefore, for the moment, hard to implement as the basis has to be adjusted according to the type of problem that has to be tackled. Last but not least, for frequency sweeps i.e. when varying the frequency of interest solely, the p -FEM does not require a complete calculation of the element matrices for each frequency, unlike for a wave-based method which basis is frequency dependent. There is therefore still significant benefits in using p -FEM, in terms of performance and robustness.

It should be noted that the present work used the wave-based DGM as an example of Trefftz method. Other wave-based methods should also be considered. The wave-based DGM is closely related to a number of other methods, in particular the ultra

weak variational formulation [34, 93] and the least-squares method [125]. Both methods can be recovered in the discontinuous Galerkin framework by using different numerical fluxes. Existing results indicate that the change of numerical flux is unlikely to change the performance of the methods sufficiently to alter our conclusions [71]. Nevertheless, other categories of wave-based methods remain to be compared systematically with high-order polynomial methods. The discontinuous enrichment method (DEM), for example, uses the standard finite element basis together with a set of plane waves defined in each element. Lagrange multipliers are adopted at the element interfaces to impose continuity of the solution [53, 120]. Another example is the Partition of Unity Finite Element Method (PUFEM) [122] using a standard finite element polynomial basis multiplied with a set of plane waves at each node. Extending these comparisons to three-dimensional problems would also be useful. Our conclusions hold for the Helmholtz operator but it would also be interesting to extend the comparison to other wave propagation problems in order to determine if the advantages of the p -FEM still stand. In particular, one could look at vector operators induced by the linearised Euler equations or the Maxwell's equations for electromagnetism.

A domain decomposition solver combined with a high-order discretisation for frequency-domain flow acoustics

The performance of a method relies heavily on efficient solving procedures for the resulting large, sparse, complex linear systems. On the one hand, continued efforts are deployed to develop efficient iterative solvers for this purpose. However, due to the properties of the underlying operator, the design of robust iterative solvers for Helmholtz problems remains a challenge. On the other hand, direct solving procedures do not scale well with respect to problem size and frequency. In practice, the amount of computational resources available is often not sufficient to tackle the frequency range of interest. Domain decomposition methods emerged as powerful ways to bypass this issue. In this work, a Finite Element Tearing and Interconnecting method with two Lagrange Multipliers (FETI-2LM) is proposed to solve large-scale finite element models for the propagation of sound with a background mean flow. The idea is to split the computational domain into a collection of non-overlapping sub-domains. An iterative solution procedure is formulated in terms of Lagrange multipliers which are unknowns defined only on the interfaces between sub-domains. In each sub-domain, a high-order finite element method is used to solve the governing equations, further reducing the computational costs. This approach allows to solve large-scale problems with only a fraction of the memory requirements compared to the standard approach which is to use a direct solver for the complete problem. The FETI approaches have been used extensively for Helmholtz problems and their performance is well documented for conventional finite elements. To our knowledge, the behaviour of the FETI-2LM method

- in conjunction with a high-order FEM;
- for the solution of large scale flow acoustics

is still unclear.

First, the performance of the FETI-2LM method combined with the high order FEM was examined for the solution of Helmholtz problems. A scalability analysis has been conducted to verify the behaviour of the solver with respect to the mesh size, the frequency and the number of sub-domains. In particular, the influence of the polynomial order was assessed. The method was found scalable with respect to the polynomial order when the mesh size and Helmholtz number are kept fixed. The same conclusion can be drawn when the Helmholtz number and the number of DOF per wavelength are kept constant. However, when the target L^2 error is fixed for a given frequency, we find that increasing the polynomial order results in a reduction of the required number of iterations for convergence of the FETI-2LM algorithm. This shows the superiority of the higher-order methods. The costs associated with the hybrid solver were also compared to the costs induced by MUMPS, for a simple 3D wave guide problem. The advantages of using the FETI-2LM method in terms of memory requirements were clear. However, solution times were comparable for both solvers. The FETI approach was then used in combination with an adaptive order FEM. For the sphere test case, the scalability of the method with respect to the number of partitions was retrieved. At fixed accuracy, increasing the frequency does not slow down the convergence. Varying the accuracy leads to slight changes in the iteration counts. However, in both experiments, increasing the frequency at fixed accuracy or increasing the accuracy at fixed frequency increases the costs of a single iteration.

The FETI-2LM approach was applied to the propagation of sound with a background mean flow described by the linearised potential equation. It was shown that, through the choice of suitable Lagrange multipliers, it is possible to obtain an algebraic formulation that is similar to the formulation associated with the Helmholtz equation. The method has been validated and its scalability has been assessed with respect to the flow velocity, the frequency and the number of sub-domains. We found that keeping the mesh, frequency and target error constant and varying the Mach number leads to a slight increase of the iterations counts. Finally, a 3D generic intake propagation problem was solved using the proposed FETI-2LM formulation.

The next investigations should involve preconditioning techniques and new transmission conditions at the interface between sub-domains. In [76], optimal transmission conditions leading to the convergence of the Schwarz method in a finite number of iterations are found by Fourier analysis. They show that as these optimal conditions are non-local, local approximations are necessary in practice. Second order approximations are derived and used for the solution of a three-dimensional realistic Helmholtz problem. The results are compared to the zero-th order approximation of the optimal condition

which is equivalent to a Robin condition. Domain decomposition algorithms are strongly influenced by the quality of the approximation of the optimal transmission condition. Using a higher order approximation could therefore accelerate the convergence of the method. As mentioned previously, the present method does not scale with the number of sub-domains for most investigated problems: as the number of sub-domains is increased, the number of iterations needed to converge to a solution is increased. In order to accelerate the convergence and recover a scalability with respect to the number of sub-domains, coarse spaces can be considered [118]. The coarse space facilitates a global exchange of information. They can provide a convergence rate that is independent of the number of sub-domains [162]. In two-level methods, a coarse space correction is introduced: a small problem of size of the order of the number of sub-domains couples all sub-domains at each iteration. For example, in [57], the method is equipped with a coarse problem based on plane waves. The authors prove that the resulting method called two-level FETI-2LM method is scalable with respect to the wave number and the number of sub-domains which make it suitable for solving high frequency acoustic problems.

List of publications

Conference paper

- A comparison of high-order finite element method and wave-based discontinuous Galerkin method for Helmholtz problems, A. Lieu, G. Gabard and H. Bériot. *The 12th International Conference on Mathematical and Numerical Aspects of Wave Propagation*, Karlsruhe, Germany, 2015.
- A performance study of high-order finite element and wave-based discontinuous Galerkin methods, A. Lieu, G. Gabard and H. Bériot. *The 22th International Congress on Sound and Vibration*, Florence, Italy, 2015.
- Performance of a domain decomposition method for high-order FEM, A. Lieu, H. Bériot, G. Gabard and F.-X. Roux, *European Congress on Computational Methods in Applied Sciences and Engineering*, Crete Island, Greece, 2016.
- A domain decomposition method with high-order finite elements for flow acoustics, A. Lieu, G. Gabard, H. Bériot, *25th AIAA/CEAS Aeroacoustics Conference*, Delft, The Netherlands, 2019.

Journal paper

- A comparison of high-order polynomial and wave-based methods for Helmholtz problems, A. Lieu, G. Gabard and H. Bériot. *Journal of Computational Physics*, 321:105-125, 2016.

Annexe A

Scattering of a Plane Wave by a Cylinder

An additional test case has been considered to assess the convergence behaviour of the p -FEM and the wave-based DGM. The problem of the scattering of a plane wave by a cylinder has an analytic solution and features a mix of solution modes and plane wave directions. Examples of results for the solution and the numerical error are shown in Figure A.1 and a convergence graph can be found in Figure A.2.

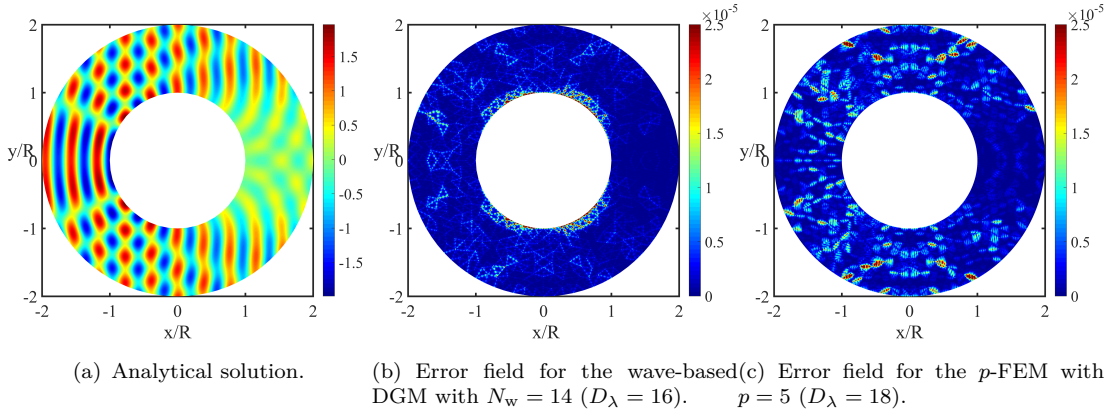


FIGURE A.1: Scattering of a plane wave by a cylinder at $kL = 50$ and $h/L = 1/20$.

We see on the example of solution in Figure A.1(a) that the incident plane wave remains significant. This implies that the results from the wave-based method will depend strongly on the alignment of the plane-wave basis with this incident plane wave.

Figures A.1(b) and A.1(c) show the magnitude of the difference between exact and numerical solutions. For the wave-based DGM, this error is clearly dominated by the solution close to the cylinder. This is explained by the results from the individual spinning waves shown on Figures 3.1(b), 3.1(c): the evanescent part of the solution (localised in the vicinity of the cylinder) is more difficult to represent accurately with the wave-based




FIGURE A.2: Scattering of a plane wave by a cylinder. Relative L^2 -error (%) against D_λ at $kL = 50$; red dashed lines: wave-based DGM; black solid lines: p -FEM with condensation. For each curve, D_λ is increased by refining the mesh while p or N_w is fixed.

method. In contrast, the error for the p -FEM is more uniformly distributed over the computational domain which is also consistent with the results for the spinning waves since with this method there is little difference in accuracy between the propagating and evanescent waves.

In figure A.2 the convergence results for the p -FEM are very similar with those for the spinning waves, which is again consistent with the fact that the polynomial method describes both the evanescent and propagating waves with a similar level of accuracy. The relative error for the wave-based method is slightly above the error for the propagating spinning wave, and below the error for the evanescent wave. This is also expected since the scattered field from the cylinder is made of spinning waves.

The behaviour of the two numerical methods can be anticipated based on the results already included for the evanescent and propagating modes considered individually. These results illustrate that it is preferable to consider the propagating and evanescent spinning waves separately as by considering the propagating and evanescent spinning waves separately, one is able to gain more detailed insight into the behaviour of the numerical methods.

Annexe B

Odd and Even Number of Plane Waves

In the literature, there are both numerical evidences and theoretical results showing that $N_w = 2n$ and $N_w = 2n - 1$ provide the same order of convergence. For our study, in addition to the results for odd numbers of plane waves, even numbers of plane waves have also been considered. The convergence plots for the problem of the propagating spinning waves are shown in Fig. B.1. It was indeed observed that $N_w = 2n$ and $N_w = 2n - 1$ provide the same order of convergence and the performance of the wave-based DGM with even N_w was found to be slightly below those obtained with odd values of N_w (see Fig. B.2.). The results are slightly improved in terms of conditioning and efficiency. However, this improvement is limited and does not alter the discussion of the results nor the conclusions of the comparison between p -FEM and the wave-based DGM. number of plane waves offers a gain in performance [95, 87].




FIGURE B.1: Spinning wave problem with $kL = 50$, $m = 10$. The relative L^2 error against D_λ (left) and against the condition number (right). Black lines: p -FEM with condensation; blue lines: p -FEM without condensation; red dashed lines: wave-based DGM with $N_w = 2n$, magenta dashed lines wave-based DGM with $N_w = 2n - 1$.

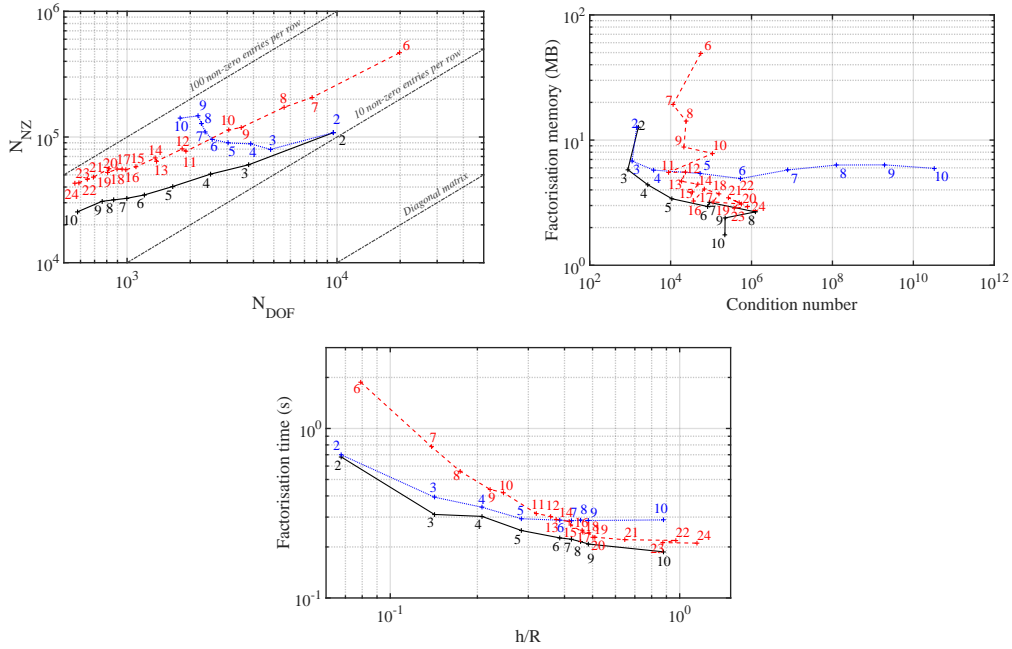


FIGURE B.2: Spinning wave problem with $kL = 50$, $m = 10$. The relative L^2 error against D_λ (left) and against the condition number (right). Black lines: p -FEM with condensation; blue lines: p -FEM without condensation; red dashed lines: wave-based DGM with $N_w = 2n$, magenta dashed lines wave-based DGM with $N_w = 2n - 1$.

Bibliography

- [1] MUMPS 4.10.0. *Users' guide*, 2011.
- [2] A. Agarwal, P. J. Morris, and R. Mani. Calculation of sound propagation in nonuniform flows: suppression of instability waves. *AIAA Journal*, 42(1):80–88, January 2004.
- [3] M. Ainsworth. Discrete dispersion relation for *hp*-version finite element approximation at high wave number. *SIAM Journal on Numerical Analysis*, 42(2):553–575, 2004.
- [4] C. J. S. Alves and C. S. Chen. A new method of fundamental solutions applied to nonhomogeneous elliptic problems. *Advances in Computational Mathematics*, 23(1-2):125–142, 2005.
- [5] P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41, 2001.
- [6] D. Appelo, M. Inkman, T. Hagstrom, and T. Colonius. Hermite methods for aeroacoustics: recent progress. *17th AIAA/CEAS Aeroacoustics Conference, Portland, Oregon*, June 2011.
- [7] R. J. Astley. Numerical methods for noise propagation in moving flows, with application to turbofan engines. *Acoustical Science and Technology*, 30:227–39, July 2009.
- [8] R. J. Astley and P. Gamallo. Special short wave elements for flow acoustics. *Computer Methods in Applied Mechanics and Engineering*, 194(2-5, SI):341–353, 2005.
- [9] W. Babisch, D. Houthuijs, G. Pershagen, E. Cadum, K. Katsouyanni, M. Velonakis, M.-L. Dudley, H.-D. Marohn, W. Swart, O. Breugelmans, G. Bluhm, J. Selander, F. Vigna-Taglianti, S. Pisani, A. Haralabidis, K. Dimakopoulou, I. Zachos, and L. Järup. Annoyance due to aircraft noise has increased over the years—Results of the HYENA study . *Environment International*, 35(8):1169 – 1176, 2009.

- [10] I. Babuška and M. Suri. The p and hp versions of the finite element method, basic principles and properties. *SIAM review*, 36(4):578–632, 1994.
- [11] I. Babuška, B. Andersson, B. Guo, J. M. Melenk, and H. S. Oh. Finite element method for solving problems with singular solutions. *Journal of Computational and Applied Mathematics*, 74(1-2):51–70, November 1996.
- [12] I. Babuška and B. Q. Guo. The h , p and h - p version of the finite element method - Basic theory and applications. *Advances in Engineering Software*, 15(3-4):159–174, 1992.
- [13] I. Babuška, F. Ihlenburg, E. T. Paik, and S. A. Sauter. A generalized finite element method for solving the Helmholtz equation in two dimensions with minimal pollution. *Computer Methods in Applied Mechanics and Engineering*, 128(3-4):325–359, December 1995.
- [14] I. Babuška and S. A. Sauter. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? *SIAM Review*, 42(3):451–484, September 2000.
- [15] A. Barka and F. X. Roux. Scalability of FETI-2LM methods on HPC clusters for antenna RCS applications. In *2014 International Symposium on Antennas and Propagation Conference Proceedings*, pages 19–20. HAL CCSD, November 2014.
- [16] E. Bécache, A.-S. Bonnet-Ben Dhia, and G. Legendre. Perfectly matched layers for the convected Helmholtz equation. *SIAM Journal on Numerical Analysis*, 42(1):409–433, 2004.
- [17] J. D. Benamou and B. Després. A domain decomposition method for the Helmholtz equation and related optimal control problems. *Journal of Computational Physics*, 136(1):68–82, September 1997.
- [18] J.-P. Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114:185–200, 1994.
- [19] H. Bériot and G. Gabard. Anisotropic adaptivity of the p -fem for time-harmonic acoustic wave propagation. *Journal of Computational Physics*, 378:234 – 256, 2019.
- [20] H. Bériot, G. Gabard, and E. Perrey-Debain. Analysis of high-order finite elements for convected wave propagation. *International Journal for Numerical Methods in Engineering*, 96(11):665–688, December 2013.
- [21] H. Bériot, A. G. Prinn, and G. Gabard. Efficient implementation of high-order finite elements for Helmholtz problems. *International Journal for Numerical Methods in Engineering*, 106:213–240, April 2016.

- [22] H. Bériot and M. Tournour. *Numerical methods for acoustics problems*, chapter Perfectly matched layers for time harmonic acoustics. Saxe-Coburg Publications, January 2016.
- [23] S. Beuchler and M. Purrucker. Schwarz type solvers for *hp*-FEM discretizations of mixed problems. *Computational Methods in Applied Mathematics*, 112:369–390, 2012.
- [24] I. Bica. Nonoverlapping domain decomposition algorithms for the p-version finite element method for elliptic problems. *Contemporary Mathematics*, 218:231–237, 1998.
- [25] P. Bientinesi, V. Eijkhout, K. Kim, J. Kurtz, and R. van de Geijn. Sparse direct factorizations through unassembled hyper-matrices. *Computer Methods in Applied Mechanics and Engineering*, 199(9–12):430 – 438, 2010.
- [26] Y. Boubendir. An analysis of the BEM-FEM non-overlapping domain decomposition method for a scattering problem. *Journal of Computational and Applied Mathematics*, 204(2):282 – 291, 2007. Special Issue: The Seventh International Conference on Mathematical and Numerical Aspects of Waves (WAVES’05) .
- [27] Y. Boubendir, X. Antoine, and C. Geuzaine. A quasi-optimal non-overlapping domain decomposition algorithm for the Helmholtz equation. *Journal of Computational Physics*, 231(2):262 – 280, 2012.
- [28] J.-F. Bourgat, R. Glowinski, P. Le Tallec, and M. Vidrascu. Variational Formulation and Algorithm for Trace Operator in Domain Decomposition Calculations. In Chan, T. and Glowinski, R. and= Périaux, J. and Widlund, O., editor, *Domain Decomposition Methods*, pages 3–16, Philadelphia, PA, 1989. SIAM.
- [29] J. P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover, 2000.
- [30] A. Buffa and P. Monk. Error estimates for the ultra weak variational formulation of the Helmholtz equation. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 42(6):925–940, 2008.
- [31] X. Cai. Overlapping domain decomposition methods. In H. P. Langtangen and A. Tveito, editors, *Advanced Topics in Computational Partial Differential Equations*, volume 33 of *Lecture Notes in Computational Science and Engineering*, pages 57–95. Springer Berlin Heidelberg, 2003.
- [32] X. Cai and O. B. Widlund. Domain decomposition algorithms for indefinite elliptic problems. *SIAM Journal on Scientific and Statistical Computing*, 13(1):243–258, January 1992.
- [33] J. E. Caruthers, J. C. French, and G. K. Raviprakash. Green function discretization for numerical solution of the Helmholtz equation. *Journal of Sound And Vibration*, 187(4):553–568, November 1995.

- [34] O. Cessenat and B. Després. Application of an ultra weak variational formulation of elliptic PDEs to the two-dimensional Helmholtz problem. *SIAM Journal on Numerical Analysis*, 35(1):255–299, February 1998.
- [35] O. Cessenat and B. Després. Using plane waves as base functions for solving time harmonic equations with the ultra weak variational formulation. *Journal of Computational Acoustics*, 11(2):227–238, June 2003.
- [36] T. F. Chan and T. P. Mathew. Domain decomposition algorithms. *Acta Numerica*, 3:61–143, 1994.
- [37] C. Chevalier and F. Pellegrini. PT-Scotch: A tool for efficient parallel graph ordering. *CoRR*, abs/0907.1375, 2009.
- [38] P. Chevalier and F. Nataf. Symmetrized method with optimized second-order conditions for the Helmholtz equation. *Contemporary Mathematics*, 218:400–407, 1998.
- [39] F. Collino. Conditions absorbantes d’ordre élevé pour des modèles de propagation d’onde dans des domaines rectangulaires. Research Report RR-1790, INRIA, 1992. Projet IDENT.
- [40] F. Collino, S. Ghanemi, and P. Joly. Domain decomposition method for harmonic wave propagation: a general presentation. *Computer Methods in Applied Mechanics and Engineering*, 184(2-4):171–211, 2000.
- [41] A. de La Bourdonnaye, C. Farhat, A. Macedo, F. Magoulès, and F. X. Roux. A non-overlapping domain decomposition method for the exterior Helmholtz problem. *Contemporary Mathematics*, 218:42 – 66, 1998.
- [42] A. Deraemaeker, I. Babuška, and P. Bouillard. Dispersion and pollution of the FEM solution for the Helmholtz equation in one, two and three dimensions. *International Journal for Numerical Methods in Engineering*, 46(4):471–499, October 1999.
- [43] W. Desmet. *A wave based prediction technique for coupled vibro-acoustic analysis*. PhD thesis, KU Leuven, Belgium, 1998.
- [44] B. Després. Domain decomposition method and the Helmholtz problem. *SIAM Mathematical and numerical aspects of wave propagation phenomena*, pages 44–52, 1991.
- [45] V. Dolean, M. E. Bouajaji, M. J. Gander, S. Lanteri, and R. Perrussel. Domain decomposition methods for electromagnetic wave propagation problems in heterogeneous media and complex domains. In Y. Huang, R. Kornhuber, O. B. Widlund, and J. Xu, editors, *Domain Decomposition Methods in Science and Engineering*

- XIX*, volume 78 of *Lecture Notes in Computational Science and Engineering*, pages 15–26. Springer Berlin Heidelberg, 2011.
- [46] V. Dolean, P. Jolivet, and F. Nataf. An Introduction to Domain Decomposition Methods: algorithms, theory and parallel implementation. Lecture, January 2015.
- [47] I. S. Duff. Direct methods. Technical report, Rutherford Appleton Laboratory, 1998. Technical Report RAL-98-054.
- [48] I. S. Duff and J. K. Reid. The multifrontal solution of indefinite sparse symmetric linear equations. *ACM Transactions on mathematical software*, 9(3):302–325, 1983.
- [49] B. Engquist and A. Majda. Absorbing boundary conditions for numerical simulation of waves. *Proceedings of the National Academy of Sciences*, 74(5):1765–1766, 1977.
- [50] O. G. Ernst and M. J. Gander. *Why it is difficult to solve Helmholtz problems with classical iterative methods*, pages 325–363. Springer Berlin Heidelberg, 2012.
- [51] Q. H. Fang. A perfectly matched layer absorbing boundary condition for linearized Euler equations with a non-uniform mean flow. *Journal of Computational Physics*, 208(2):469 – 492, 2005.
- [52] C. Farhat, P. Avery, R. Tezaur, and J. Li. FETI-DPH: a dual-primal domain decomposition method for acoustic scattering. *Journal of Computational Acoustics*, 13:499 – 524, 2005.
- [53] C. Farhat, I. Harari, and L. P. Franca. The discontinuous enrichment method. *Computer Methods in Applied Mechanics and Engineering*, 190(48):6455–6479, 2001.
- [54] C. Farhat, I. Harari, and U. Hetmaniuk. A discontinuous Galerkin method with Lagrange multipliers for the solution of Helmholtz problems in the mid-frequency regime. *Computer Methods in Applied Mechanics and Engineering*, 192(11-12):1389–1419, 2003.
- [55] C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen. FETI-DP: a dual-primal unified FETI method—part I: A faster alternative to the two-level FETI method. *International journal for numerical methods in engineering*, 50(7):1523–1544, 2001.
- [56] C. Farhat, M. Lesoinne, and K. Pierson. A scalable dual-primal domain decomposition method. *Numerical linear algebra with applications*, 7(7-8):687–714, 2000.
- [57] C. Farhat, A. Macedo, M. Lesoinne, F.-X. Roux, F. Magoulès, and A. de La Bourdonnaie. Two-level domain decomposition methods with Lagrange multipliers for the fast iterative solution of acoustic scattering problems. *Computer Methods in Applied Mechanics and Engineering*, 184(2–4):213 – 239, 2000.

- [58] C. Farhat, A. Macedo, and R. Tezaur. FETI-H: a scalable domain decomposition method for high frequency exterior Helmholtz problems. In C. Lai, P. Bjørstad, M. Cross, and O. Widlund, editors, *Eleventh International Conference on Domain Decomposition Method*, page 231–241, 1999.
- [59] C. Farhat, J. Mandel, and F.-X. Roux. Optimal convergence properties of the FETI domain decomposition method. *Computer Methods in Applied Mechanics and Engineering*, 115(3–4):365 – 385, 1994.
- [60] C. Farhat and F.-X. Roux. A method of finite element tearing and interconnecting and its parallel solution algorithm. *International Journal for Numerical Method in Engineering*, 32(6):1205–1227, October 1991.
- [61] C. Farhat, R. Tezaur, and P. Weidemann-Goiran. Higher-order extensions of a discontinuous Galerkin method for mid-frequency Helmholtz problems. *International Journal for Numerical Method in Engineering*, 61(11):1938–1956, November 2004.
- [62] R. T. Farouki. The Bernstein polynomial basis: a centennial retrospective. *Computer Aided Geometric Design*, 29(6):379–419, August 2012.
- [63] R. W. Freund and N. M. Nachtigal. QMR: a quasi-minimal residual method for non-Hermitian linear systems. *Numerische Mathematik*, 60(1):315–339, 1991.
- [64] L. Fritschi and WHO Regional Office for Europe. Burden of disease from environmental noise, Quantification of healthy life years lost in Europe. Technical report, World Health Organization, 2011.
- [65] R. A. Fuller, P. H. Warren, and K. J. Gaston. Daytime noise predicts nocturnal singing in urban robins. *Biology Letters*, 3(4):368–370, 2007.
- [66] G. Gabard. Discontinuous Galerkin methods with plane waves for time-harmonic problems. *Journal of Computational Physics*, 225(2):1961–1984, August 2007.
- [67] G. Gabard. Fundamentals of aero-acoustics. Technical report, University of Southampton, 2014. Lecture notes.
- [68] G. Gabard, R. J. Astley, P. Gamallo, and G. Kennedy. Physics-based computational methods for aero-acoustics. In Astley, R. J. and Gabard, G, editor, *IUTAM Symposium on computational aero-acoustics for aircraft noise prediction*, volume 6 of *Procedia Engineering*, pages 183–192, 2010.
- [69] G. Gabard, H. Bériot, A. G. Prinn, and K. Kucukcoskun. Adaptive, high-order finite-element method for convected acoustics. *AIAA Journal*, 56(8):3179–3191, 2018.
- [70] G. Gabard and O. Dazel. A discontinuous Galerkin method with plane waves for sound absorbing materials. *International Journal for Numerical Methods in Engineering*, 2015.

- [71] G. Gabard, P. Gamallo, and T. Huttunen. A comparison of wave-based discontinuous Galerkin, ultra-weak and least-square methods for wave problems. *International Journal for Numerical Methods in Engineering*, 85(3):380–402, January 2011.
- [72] G. Gabard, A. G. Prinn, H. Bériot, and K. Kucukcoskun. An adaptive, high-order finite element method for aeroengine acoustics. *22nd AIAA/CEAS Aeroacoustics Conference*, 2016.
- [73] P. Gamallo and R. J. Astley. A comparison of two Trefftz-type methods: the ultraweak variational formulation and the least-squares method, for solving short-wave 2-D Helmholtz problems. *International Journal for Numerical Methods in Engineering*, 71(4):406–432, July 2007.
- [74] M. J. Gander. Optimized Schwarz methods. *SIAM Journal on Numerical Analysis*, 44(2):699–731, 2006.
- [75] M. J. Gander. Schwarz methods over the course of time. *Electronic transactions on numerical analysis*, 31:228–255, 2008.
- [76] M. J. Gander, F. Magoules, and F. Nataf. Optimized Schwarz methods without overlap for the Helmholtz equation. *SIAM Journal on Scientific Computing*, 24(1):38–60, August 2002.
- [77] C. Geuzaine and J. F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79:1309–1331, 2009.
- [78] C. Gittelsohn and R. Hiptmair. Dispersion analysis of plane wave discontinuous Galerkin methods. *International Journal for Numerical Methods in Engineering*, 98(5):313–323, 2014.
- [79] C. Gittelsohn, R. Hiptmair, and I. Perugia. Plane wave discontinuous Galerkin methods: analysis of the h -version. *ESAIM: Mathematical Modelling and Numerical Analysis*, 43(02):297–331, 2009.
- [80] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. Society for Industrial and Applied Mathematics, 1997.
- [81] B. Guo and W. Cao. Additive Schwarz methods for the h - p version of the finite element method in two dimensions. *SIAM Journal on Scientific Computing*, 18(5):1267–1288, 1997.
- [82] K. K. Gupta and J. L. Meek. A brief history of the beginning of the finite element method. *International Journal for Numerical Methods in Engineering*, 39(22):3761–3774, November 1996.

- [83] K. Hamiche, G. Gabard, and H. Bériot. A high-order finite element method for the linearised Euler equations. *Acta Acustica united with Acustica*, 102, April 2016.
- [84] W. Han and W. K. Liu. Flexible piecewise approximations based on partition of unity. *Advances in Computational Mathematics*, 23:191–9, July 2005.
- [85] R. Hiptmair, A. Moiola, and I. Perugia. Plane wave discontinuous Galerkin methods for the 2D Helmholtz equation: analysis of the p -version. *SIAM Journal on Numerical Analysis*, 49(1):264–284, 2011.
- [86] R. Hiptmair, A. Moiola, and I. Perugia. Trefftz discontinuous Galerkin methods for acoustic scattering on locally refined meshes. *Applied Numerical Mathematics*, 79:79–91, 2014.
- [87] R. Hiptmair, A. Moiola, and I. Perugia. *A survey of Trefftz methods for the Helmholtz equation*, pages 237–279. Springer International Publishing, Cham, 2016.
- [88] R. Hiptmair, A. Moiola, I. Perugia, and C. Schwab. Approximation by harmonic polynomials in star-shaped domains and exponential convergence of Trefftz hp -dGFEM. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48(03):727–752, 2014.
- [89] P. Hood. Frontal solution program for unsymmetric matrices. *International Journal for Numerical Methods in Engineering*, 10(2):379–399, 1976.
- [90] Q. Hu and L. Yuan. A weighted variational formulation based on plane wave basis for discretization of Helmholtz equations. *International Journal of analysis and modeling*, 11(3):587–607, 2014.
- [91] D. L. Huff. Noise reduction technologies for turbofan engines. *Internoise*, 2007.
- [92] T. Huttunen, P. Gamallo, and R. J. Astley. Comparison of two wave element methods for the Helmholtz problem. *Communications in Numerical Methods in Engineering*, 25(1):35–52, January 2009.
- [93] T. Huttunen, P. Monk, and J. P. Kaipio. Computational aspects of the ultra-weak variational formulation. *Journal of Computational Physics*, 182(1):27–46, October 2002.
- [94] F. Ihlenburg and I. Babuška. Finite element solution of the Helmholtz equation with high wave number part II: the h - p version of the FEM. *SIAM Journal on Numerical Analysis*, 34(1):315–358, February 1997.
- [95] L.-M. Imbert-Gérard and B. Després. A generalized plane-wave numerical method for smooth nonconstant coefficients. *IMA Journal of Numerical Analysis*, 34(3):1072–1103, July 2014.

- [96] B. M. Irons. A frontal solution program for finite element analysis. *International Journal for Numerical Methods in Engineering*, 2:15–32, 1970.
- [97] J.-M. Jin and W. C. Chew. Combining PML and ABC for the finite-element analysis of scattering problems. *Microwave and Optical Technology Letters*, 12(4):192–197, 1996.
- [98] S. G. Jonhson. Notes on perfectly match layers. Technical report, MIT, 2010. Lecture notes.
- [99] G. Karniadakis and S. Sherwin. *Spectral/hp element methods for computational fluid dynamics*. Oxford University Press, 2013.
- [100] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, 20(1):359–392, December 1998.
- [101] C. R. Kight and J. P. Swaddle. How and why environmental noise impacts animals: an integrative, mechanistic review. *Ecology letters*, 14(10):1052–1061, 2011.
- [102] A. Kuzmin, M. Luisier, and O. Schenk. Fast Methods for Computing Selected Elements of the Green’s Function in Massively Parallel Nanoelectronic Device Simulations. In F. Wolf, B. Mohr, and D. an Mey, editors, *Euro-Par 2013 Parallel Processing*, pages 533–544, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [103] P. Ladevèze, L. Arnaud, P. Rouch, and C. Blanzé. The variational theory of complex rays for the calculation of medium frequency vibrations. *Engineering Computations*, 18(1/2):193–214, 2001.
- [104] O. Laghrouche, P. Bettess, and R. J. Astley. Modelling of Short Wave Diffraction Problems Using Approximating Systems of Plane Waves. *International Journal for Numerical Methods in Engineering*, 54(10):1501–1533, 8 2002.
- [105] P. Le Tallec, Y. H. De Roeck, and M. Vidrascu. Domain decomposition methods for large linearly elliptic three-dimensional problems. *Journal of Computational and Applied Mathematics*, 34(1):93 – 117, 1991.
- [106] J. Li and X. Tu. Convergence analysis of a balancing domain decomposition method for solving a class of indefinite linear systems. *Numerical Linear Algebra with Applications*, 16(9):745–773, 2009.
- [107] J. Li and O. B. Widlund. FETI-DP, BDDC, and block Cholesky methods. *International Journal for Numerical Methods in Engineering*, 66(2):250–271, 2006.
- [108] P. L. Lions. On the Schwarz alternating method I. In *First International Symposium on domain decomposition methods for partial differential equations*, pages 1–42, 1988.

- [109] P. L. Lions. On the Schwarz alternating method III: a variant for the non-overlapping subdomains. In *Third International Symposium on domain decomposition methods for partial differential equations*, pages 202–223, 1990.
- [110] T. Luostari, T. Huttunen, and P. Monk. The ultra weak variational formulation using Bessel basis functions. *Communications in Computational Physics*, 11(2):400–414, February 2012.
- [111] T. Luostari, T. Huttunen, and P. Monk. Improvements for the ultra weak variational formulation. *International Journal for Numerical Methods in Engineering*, 94(6):598–624, MAY 11 2013.
- [112] J. Mabe, F. Calkins, and G. Butler. Boeing’s variable geometry chevron, morphing aerostructure for jet noise reduction. *American Institute of Aeronautics and Astronautics*, 2142, April 2006.
- [113] F. Magoulès, K. Meerbergen, and J.-P. Coyette. Application of a domain decomposition method with Lagrange multipliers to acoustic problems arising from the automotive industry. *Journal of Computational Acoustics*, 8(3):503–521, 2000.
- [114] F. Magoulès, F.-X. Roux, and S. Salmon. Optimal discrete transmission conditions for a non-overlapping domain decomposition method for the Helmholtz equation. *Siam Journal on Scientific Computing*, 25, January 2004.
- [115] J. Mandel. Balancing domain decomposition. *Communications in Numerical Methods in Engineering*, 9(3):233–241, 1993.
- [116] J. Mandel and C. R. Dohrmann. Convergence of a balancing domain decomposition by constraints and energy minimization. *Numerical Linear Algebra with Applications*, 10(7):639–659, 2003.
- [117] J. Mandel and B. Sousedík. BDDC and FETI-DP under minimalist assumptions. *Computing*, 81(4):269–280, 2007.
- [118] J. Mandel and B. Sousedík. Coarse spaces over the ages. *Lecture Notes in Computational Science and Engineering*, 78, November 2009.
- [119] J. Mandel and R. Tezaur. Convergence of a substructuring method with Lagrange multipliers. *Numerische Mathematik*, 73(4):473–487, 1996.
- [120] P. Massimi, R. Tezaur, and C. Farhat. A discontinuous enrichment method for three-dimensional multiscale harmonic wave propagation problems in multi-fluid and fluid-solid media. *International Journal for Numerical Methods in Engineering*, 76(3):400–425, October 2008.
- [121] T. P. Mathew. Uniform convergence of the Schwarz alternating method for solving singularly perturbed advection-diffusion equations. *SIAM Journal on Numerical Analysis*, 35(4):1663–1683, August 1998.

- [122] J. M. Melenk and I. Babuška. The partition of unity finite element method: basic theory and applications. *Computer Methods in Applied Mechanics and Engineering*, 139(1-4):289–314, DEC 15 1996.
- [123] V. Mengle, U. Ganz, E. Nesbitt, E. Bultemeier, and R. Thomas. Flight test results for uniquely tailored propulsion airframe aeroacoustic chevrons: Shockcell noise. In *12th AIAA/CEAS Aeroacoustics Conference*, volume 1, May 2006.
- [124] T. Mertens, P. Gamallo, R. J. Astley, and P. Bouillard. A mapped finite and infinite partition of unity method for convected acoustic radiation in axisymmetric domains. *Computer Methods in Applied Mechanics and Engineering*, 197(49-50):4273–4283, September 2008.
- [125] P. Monk and D. Q. Wang. A least-squares method for the Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering*, 175(1-2):121–136, June 1999.
- [126] P. M. Morse and K. U. Ingard. *Theoretical Acoustics*. Princeton University Press, Princeton, New Jersey, 1986.
- [127] I. Moulitsas and G. Karypis. Algorithms for graph partitioning and fill reducing ordering for domain decomposition methods. Technical report, University of Minnesota, 01 2006.
- [128] P. Mustafi. *Improved turbofan intake liner design and optimization*. PhD thesis, University of Southampton, 2013.
- [129] F. Nataf. Interface connections in domain decomposition methods. In A. Bourlioux, M.J. Gander, and G. Sabidussi, editors, *Modern methods in scientific computing and applications*, volume 75 of *NATO Science Series II-Mathematics Physics and Chemistry*, pages 323–364. NATO Adv Study Inst; Math Superieures, July 2002.
- [130] F. Nataf, F. Rogier, and E. de Sturler. Optimal interface conditions for domain decomposition methods. *CMAP (Ecole Polytechnique)*, 301:1–18, 1994.
- [131] J. T. Oden, S. Prudhomme, and L. Demkowicz. A posteriori error estimation for acoustic wave propagation problems. *Archives of Computational Methods in Engineering*, 12(4):343–389, December 2005.
- [132] High Level Group on Aviation Research. Flightpath 2050 Europe’s Vision for Aviation. Technical report, European Commission, 2011.
- [133] D. Osei-Kuffuor and Y. Saad. Preconditioning Helmholtz linear systems. *Applied Numerical Mathematics*, 60(4):420–431, 2010.

- [134] D. Pardo, J. Álvarez Aramberri, M. Paszynski, L. Dalcin, and V.M. Calo. Impact of element-level static condensation on iterative solver performance. *Computers & Mathematics with Applications*, 70(10):2331 – 2341, 2015.
- [135] K. M Parris, M. Velik-Lord, and J. M. A. North. Frogs call at a higher pitch in traffic noise. *Ecology and Society*, 14, 2009.
- [136] L. F. Pavarino. Additive Schwarz methods for the p -version finite element method. *Numerische Mathematik*, 66(1):493–515, 1994.
- [137] S. Petersen, D. Dreyer, and O. von Estorff. Assessment of finite and spectral element shape functions for efficient iterative simulations of interior acoustics. *Computer Methods in Applied Mechanics and Engineering*, 195(44-47):6463–6478, 2006.
- [138] A. D. Pierce. *Acoustics: An introduction to its physical principles and applications*. Acoustical Society of America, 1989.
- [139] A. G. Prinn. *Efficient finite element methods for aircraft engine noise prediction*. PhD thesis, University of Southampton, United Kingdom, 2014.
- [140] P. P. Rao and P. J. Morris. Use of finite element methods in frequency domain aeroacoustics. *AIAA Journal*, 44(7):1643–1652, 2006.
- [141] S. W. Rienstra and A. Hirschberg. An introduction to acoustics. Technical report, Eindhoven University of Technology, 2014.
- [142] D. J. Rixen and C. Farhat. A simple and efficient extension of a class of substructure based preconditioners to heterogeneous structural mechanics problems. *International Journal for Numerical Methods in Engineering*, 44(4):489–516, 1999.
- [143] S. H. M. Roth, M. H. Gross, S. Turello, and F. R. Carls. A Bernstein-Bezier based approach to soft tissue simulation. *Computer Graphics Forum*, 17(3, CI):C285+, 1998.
- [144] F.-X. Roux and A. Barka. Block krylov recycling algorithms for feti-2lm applied to 3-d electromagnetic wave scattering and radiation. *IEEE Transactions on Antennas and Propagation*, 65(4):1886–1895, 04 2017.
- [145] F.-X. Roux, F. Magoulès, S. Salmon, and L. Series. Optimization of interface operator based on algebraic approach. In I. Herrera, D. E. Keyes, O. B. Widlund, and R. Yates, editors, *Domain Decomposition Methods in Science and Engineering*, pages 297–304, 2003.
- [146] Y. Saad. *Iterative methods for sparse linear systems*. PWS Publishing Company, 1995.

- [147] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, July 1986.
- [148] O. Schenk, M. Bollhoefer, and R. A. Roemer. On large scale diagonalization techniques for the Anderson model of localization. *ArXiv Mathematics e-prints*, August 2005.
- [149] H. Slabbekoorn and A. den Boer-Visser. Cities Change the Songs of Birds. *Current Biology*, 16(23):2326 – 2331, 2006.
- [150] R. Tezaur and C. Farhat. Three-dimensional discontinuous Galerkin elements with plane waves and Lagrange multipliers for the solution of mid-frequency Helmholtz problems. *International Journal for Numerical Methods in Engineering*, 66(5):796–815, 2006.
- [151] R. Tezaur, L. Zhang, and C. Farhat. A discontinuous enrichment method for capturing evanescent waves in multiscale fluid and fluid/solid problems. *Computer Methods in Applied Mechanics and Engineering*, 197(19-20):1680–1698, 2008.
- [152] K. W. Thompson. Time dependent boundary conditions for hyperbolic systems. *Journal of Computational Physics*, 1987.
- [153] L. L. Thompson. A review of finite-element methods for time-harmonic acoustics. *Journal of the Acoustical Society of America*, 119(3):1315–1330, 2006.
- [154] L. L. Thompson and P. M. Pinsky. Complex wavenumber Fourier analysis of the p -version finite element method. *Computational Mechanics*, 1993.
- [155] A. M. Turing. Rounding-off errors in matrix processes. *The Quarterly Journal of Mechanics and Applied Mathematics*, pages 1:287–308, 1948.
- [156] B. Van Antwerpen, R. Leneveu, S. Caro, and P. Ferrante. New advances in the use of actran/tm for nacelle simulations. In *AIAA 14th Aeroacoustics Conference*, 05 2008.
- [157] H. A. Van der Vorst. BI-CGSTAB: A fast and smoothly converging variant of BI-CG for the solution of nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 13(2):631–644, March 1992.
- [158] P. E. J. Vos, S. J. Sherwin, and R. M. Kirby. From h to p efficiently: Implementing finite and spectral/ h - p element methods to achieve optimal performance for low- and high-order discretisations. *Journal of Computational Physics*, 229(13):5161–5181, July 2010.
- [159] P. Šolín, K. Segeth, and I. Doležel. *Higher-Order Finite Element Methods*. Chapman and Hall, 2004.

- [160] D. Wang, R. Tezaur, J. Toivanen, and C. Farhat. Overview of the discontinuous enrichment method, the ultra-weak variational formulation, and the partition of unity method for acoustic scattering in the medium frequency regime and performance comparisons. *International Journal for Numerical Methods in Engineering*, 89:403–17, January 2012.
- [161] G. B. Whitham. *Linear and nonlinear waves*. Wiley-Interscience, 1999.
- [162] O. B. Widlund. The development of coarse spaces for domain decomposition algorithms. In *18th International Conference on Domain Decomposition Methods*, 2009.
- [163] D. M. Young and K. C. Jea. Generalized conjugate-gradient acceleration of non-symmetrizable iterative methods. *Linear Algebra and its Applications*, 34:159 – 194, 1980.
- [164] G. W. Zumbusch. Symmetric hierarchical polynomials for the h - p -version of finite elements. *Technical Report SC-93-32, Konrad-Zuse-Zentrum, Berlin, Germany*, 1993.