CHALMERS
UNIVERSITY OF TECHNOLOGY

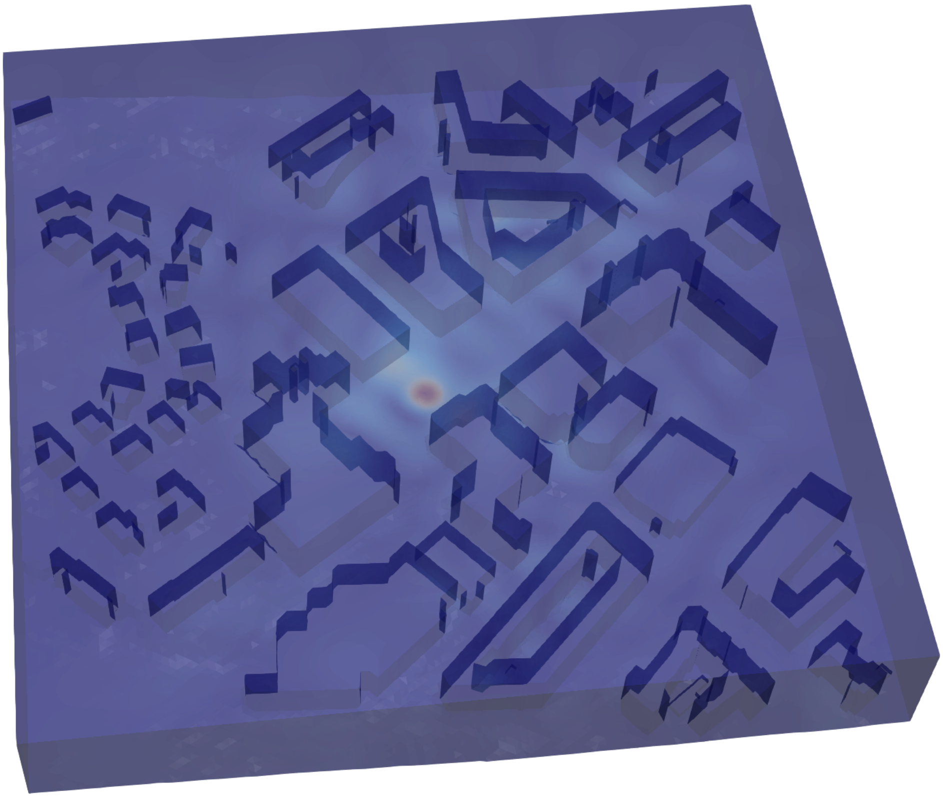

Large-Scale Noise Simulation in Urban Environments Using a Preconditioned Helmholtz Solver

Master's thesis in Complex Adaptive Systems

MARINA SAHAKYAN

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's thesis 2026

Large-Scale Noise Simulation in Urban Environments Using a Preconditioned Helmholtz Solver

MARINA SAHAKYAN

Department of Mathematical Sciences
Applied Mathematics and Statistics
Chalmers University of Technology
Gothenburg, Sweden 2026

Large-Scale Noise Simulation in Urban Environments Using a Preconditioned Helmholtz Solver
MARINA SAHAKYAN

Supervisor: Anders Logg, Department of Mathematical Sciences
Examiner: Anders Logg, Department of Mathematical Sciences

Master's Thesis 2026
Department Mathematical Sciences
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Simulated acoustic pressure amplitude on a realistic city-domain mesh. The source frequency is 5 Hz, and the city mesh has a nominal element size of 5 m.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Large-Scale Noise Simulation in Urban Environments
Using a Preconditioned Helmholtz Solver
MARINA SAHAKYAN
Department of Mathematical Sciences
Chalmers University of Technology

Abstract

This thesis studies the finite element solution of the time-harmonic Helmholtz equation for acoustic wave propagation in urban environments. The discretized Helmholtz equation gives rise to large, complex-valued, indefinite linear systems, for which standard iterative methods may converge very slowly. The aim of this work is to implement and evaluate a preconditioned GMRES solver for such systems, with focus on solver performance, parameter choices, and practical memory limitations.

The solver uses a two-level restricted additive Schwarz preconditioner with a spectral coarse space, following the RAS/MS-GFEM approach of Ma, Alber, Scheichl, and Zhang. The implementation is first verified on an exact plane-wave problem, then tested on a controlled rectangular box domain, and finally applied to realistic city-domain meshes representing part of Gothenburg. The experiments investigate how mesh resolution, frequency, and the choice of coarse space affect convergence and computational cost.

The results show that the two-level coarse correction considerably improves GMRES convergence compared with a one-level method and with plain GMRES without preconditioning. On the box domain, the method gives stable algebraic convergence up to 30 Hz, although the acoustic field is not fully mesh-converged at the higher frequencies. On the city geometry, the 5 Hz case gives the clearest mesh-refinement behaviour, while the 10 Hz, 20 Hz, and 30 Hz cases become increasingly demanding. The largest city runs show that explicit storage of the coarse basis can become memory-limiting. The study therefore shows that the two-level spectral Schwarz preconditioner is promising for large Helmholtz problems in urban acoustics, but that higher-frequency simulations require careful choices of mesh resolution, coarse-space size, and memory-efficient implementation.

Keywords: Helmholtz equation, urban acoustics, finite elements, GMRES, preconditioning.

Acknowledgements

I would like to express my sincere gratitude to my supervisor and examiner, Professor Anders Logg, for his guidance, encouragement, and support throughout this thesis work. I am especially grateful for his patience, thoughtful feedback, and willingness to answer my questions even outside regular working hours. His support has been invaluable during the development of this thesis.

I would also like to thank Professor Jens Forssén for generously taking the time to meet with me and sharing his knowledge. His explanations and advice helped me better understand the broader context of the work and gave me useful direction at an early stage of the project.

I am grateful to Professor Robert Scheichl, Dr. Chupeng Ma, and Dr. Lukas Holbach for their helpful email correspondence and valuable recommendations. Their detailed and thoughtful responses were very helpful and contributed to my understanding of the methods used in this thesis.

Finally, I would like to thank my mother for her unconditional support, encouragement, and care throughout my studies.

Marina Sahakyan, Gothenburg, June 2026

Statement on the Use of AI Tools

AI-based tools, including ChatGPT/Codex, were used during the preparation of this thesis as support for language editing, restructuring of text, LaTeX formatting, debugging assistance, and consistency checks between the written thesis, scripts, and numerical outputs.

The numerical implementation was written by the author. At later stages, AI tools were used to help improve code readability, align parts of the code with common Python style conventions, and assist with debugging and documentation. The algorithms, implementation choices, numerical experiments, and interpretation of the results remain the responsibility of the author.

Parts of the thesis text were developed from handwritten notes, draft paragraphs, terminal outputs, and author-written explanations, which were used as prompts for converting the material into clearer LaTeX text. Terminal outputs from numerical runs were saved as text files and provided to AI tools to help extract iteration counts, residuals, timings, and other reported quantities. All extracted values were checked against the original terminal outputs by the author before being included in the thesis.

AI tools were also used to help formulate presentation material and improve the wording of explanations. No numerical results were generated by AI tools. All reported results are based on simulations run with the implementation described in this thesis. The author takes full responsibility for the final content of the thesis.

Code availability and license

The implementation code developed for this thesis is available below:
`https://github.com/splinter1640/helmholtz-two-level-preconditioner`

The repository contains the Python implementation used for the numerical experiments, including scripts for finite element assembly, domain decomposition, construction of the spectral coarse space, the two-level restricted additive Schwarz preconditioner, and the online GMRES solve. The code is released under the MIT License. Large meshes, generated solution files, preconditioner data files, and full result outputs are not included in the repository because of file-size and data-management limitations.

List of Acronyms

Below is the list of acronyms used throughout this thesis, listed in alphabetical order:

AMG	Algebraic Multigrid
ARPACK	Arnoldi Package
DTCC	Digital Twin Cities Centre
FEM	Finite Element Method
GMRES	Generalized Minimal Residual Method
LU	Lower–Upper factorization
METIS	Mesh partitioning library
MS-GFEM	Multiscale Spectral Generalized Finite Element Method
PoU	Partition of Unity
RAS	Restricted Additive Schwarz
RMS	Root Mean Square

Nomenclature

Below is the nomenclature of indices, sets, parameters, functions, vectors, and matrices used throughout this thesis.

Indices

i	Subdomain or patch index
j	Local mode or basis-function index
m	GMRES iteration index

Sets and Domains

Ω	Computational acoustic domain
$\partial\Omega$	Boundary of the computational domain
Γ_N	Sound-hard, or Neumann, part of the boundary
Γ_R	Absorbing outer boundary where an impedance term is applied
$\mathcal{T}_h$	Finite element mesh
K	Mesh cell, usually a tetrahedral element
Ω_i^{core}	Non-overlapping core subdomain from the METIS partition
Ω_i	Overlapping patch associated with subdomain i
$\Omega_i^\star$	Oversampled patch associated with subdomain i
V	Continuous variational space
V_h	Finite element space

Physical and Discretization Parameters

c	Speed of sound
f	Frequency

ω	Angular frequency, $\omega = 2\pi f$
k	Wavenumber, $k = \omega/c$
λ	Acoustic wavelength, $\lambda = c/f$
h	Mesh size
ρ_0	Ambient air density
α	Diffusion coefficient in the finite element form
ν	Mass coefficient in the finite element form
β	Boundary coefficient used on absorbing facets
σ	Width parameter of the Gaussian source
$\boldsymbol{x}_0$	Source location
$\boldsymbol{d}$	Unit propagation direction in the plane-wave verification problem

Preconditioner Parameters

M	Number of subdomains
N	Number of global finite element degrees of freedom
n_i	Number of retained local modes on patch i
n_c	Dimension of the global coarse space, $n_c = \sum_{i=1}^{M} n_i$
ℓ_{ov}	Number of overlap layers in the cell-adjacency graph
$d_i(K)$	Cell-graph distance from cell K to the core cells of patch i
ρ	Spectral threshold parameter used in the local mode selection
$\lambda_{i,j}$	Eigenvalue associated with local mode j on patch i

Functions and Fields

$p(\boldsymbol{x}, t)$	Acoustic pressure fluctuation
$p_{\mathrm{tot}}(\boldsymbol{x}, t)$	Total acoustic pressure
$s(\boldsymbol{x}, t)$	Time-dependent acoustic source term
$q(\boldsymbol{x})$	Complex source amplitude in the Helmholtz equation
$u(\boldsymbol{x})$	Complex-valued pressure amplitude
$u_h(\boldsymbol{x})$	Finite element approximation of $u(\boldsymbol{x})$
v	Test function in the weak formulation
φ_j	Finite element basis function
$\phi_{i,j}$	Local eigenmode on patch i

$\widetilde{\chi}_i$	Raw, unnormalized partition-of-unity weight for patch i
χ_i	Normalized partition-of-unity weight associated with patch i
$\psi_{i,j}$	Global coarse basis function obtained from $\chi_i \phi_{i,j}$

Global Vectors and Matrices

$\boldsymbol{A}$	Global finite element Helmholtz matrix
$\boldsymbol{u}$	Vector of unknown complex finite element coefficients
$\boldsymbol{b}$	Right-hand side vector from the source term
$\boldsymbol{r}$	Residual vector
$\boldsymbol{r}_0$	Residual remaining after the local correction
$\boldsymbol{z}$	Output correction from the preconditioner
$\boldsymbol{z}_{\text{local}}$	Local restricted additive Schwarz correction
$\boldsymbol{z}_{\text{coarse}}$	Global coarse correction
$\boldsymbol{\mathcal{M}}^{-1}$	Two-level preconditioner applied as an operation
$\boldsymbol{E}_0$	Global coarse basis matrix
$\boldsymbol{A}_0$	Coarse matrix, $\boldsymbol{A}_0 = \boldsymbol{E}_0^* \boldsymbol{A} \boldsymbol{E}_0$

Local Matrices and Operators

$\boldsymbol{S}_i$	Local stiffness matrix on patch i
$\boldsymbol{H}_i$	Local positive Helmholtz energy matrix on patch i
$\widehat{\boldsymbol{H}}_i$	Partition-of-unity-weighted local energy matrix
$\boldsymbol{B}_i$	Local Helmholtz matrix used in local solves
$\boldsymbol{C}_i$	Matrix selecting interior degrees of freedom of the oversampled patch
$\boldsymbol{K}_i$	Block matrix in the augmented local eigenproblem
$\boldsymbol{G}_i$	Block matrix in the augmented local eigenproblem
$\boldsymbol{D}_i$	Diagonal matrix containing partition-of-unity weights
$\boldsymbol{R}_i$	Restriction operator associated with patch i
$\boldsymbol{R}_i^\star$	Restriction operator to the oversampled patch $\Omega_i^\star$
$\boldsymbol{\eta}_{i,j}$	Auxiliary part of the augmented local eigenvector

Error and Diagnostic Quantities

e_{L^2}	Relative L^2 error in the plane-wave verification problem
e_{H^1}	Relative H^1-seminorm error in the plane-wave verification problem
$\|u_h\|_{\mathrm{rms}}$	Root-mean-square value of the computed acoustic field
$\max \|u_h\|$	Maximum amplitude of the computed acoustic field
Δ_{rms}	Relative change in RMS value between two solutions
$\|\boldsymbol{b} - \boldsymbol{A}\boldsymbol{u}\|_2 / \|\boldsymbol{b}\|_2$	True relative residual of the linear system

Other Symbols

$a(u, v)$	Sesquilinear form of the Helmholtz weak formulation
$\ell(v)$	Linear functional in the weak formulation
I_h	Finite element interpolation or insertion operator
i	Imaginary unit
$\overline{v}$	Complex conjugate of v
$\boldsymbol{n}$	Outward unit normal vector
$\|\cdot\|_2$	Euclidean vector norm
$\|\cdot\|_{L^2(\Omega)}$	L^2 norm over Ω
$\|\cdot\|_{H^1(\Omega)}$	H^1 seminorm over Ω
$(\cdot)^*$	Conjugate transpose for matrices or complex conjugation for scalar functions

Contents

List of Acronyms **xi**

Nomenclature **xiii**

List of Figures **xix**

List of Tables **xxi**

1 Introduction **1**
1.1 Problem formulation 1
1.2 Aim and research questions 2
1.3 Scope and limitations 2
1.4 Thesis outline 3

2 Theory **5**
2.1 Acoustic wave propagation 5
2.2 Time-harmonic Helmholtz equation 6
2.3 Boundary conditions for acoustic domains 7
2.4 Finite element discretization 8
2.5 Why Helmholtz problems are difficult 10
2.6 Iterative solvers and GMRES 11
2.7 Preconditioning 12

3 Methods **13**
3.1 Overview of the numerical method 13
3.2 Finite element model used in the experiments 13
3.3 Domain decomposition and partition of unity 15
3.4 Spectral coarse space 17
3.5 Two-level preconditioner 18
3.6 Offline and online solution procedure 19
3.7 Post-processing and comparison measures 20

4 Numerical experiments and results **23**
4.1 Experiment design 23
4.1.1 Frequencies and wavelength resolution 23
4.1.2 Domains used in the experiments 24
4.1.3 Solver comparisons 25

4.1.4 Main output quantities . . . 25
4.1.5 Computational setup . . . 25
4.2 Verification on an exact plane wave . . . 25
4.3 Box-domain experiments . . . 29
4.3.1 Box-domain setup . . . 29
4.3.2 Mesh-refinement studies at 5, 10, and 20 Hz . . . 29
4.3.3 Receiver convergence on the box domain . . . 31
4.3.4 Effect of the spectral threshold . . . 32
4.3.5 Effect of oversampling . . . 32
4.3.6 One-level and two-level comparison . . . 33
4.3.7 Plain GMRES baseline . . . 33
4.3.8 Fixed local mode count and number of subdomains . . . 34
4.3.9 Effect of oversampling with fixed local mode count . . . 34
4.3.10 Fixed-mode study at 30 Hz . . . 35
4.4 City-domain experiments . . . 35
4.4.1 City-domain setup . . . 36
4.4.2 Mesh-refinement studies at 5 and 10 Hz . . . 36
4.4.3 Receiver convergence on the city domain . . . 39
4.4.4 Effect of retained local modes at 10 Hz . . . 39
4.4.5 Memory limitations . . . 42
4.4.6 City-domain study at 20 Hz . . . 43
4.4.7 City-domain feasibility study at 30 Hz . . . 44
4.4.8 Plain GMRES baseline on the city mesh . . . 45

5 Discussion **47**
5.1 Verification and role of the plane-wave test . . . 47
5.2 Solver convergence and acoustic convergence . . . 47
5.3 Effect of the two-level preconditioner . . . 48
5.4 Influence of frequency and mesh resolution . . . 49
5.5 Influence of preconditioner parameters . . . 49
5.6 Memory limitations of the explicit coarse basis . . . 50
5.7 Practical guidelines for parameter choice . . . 51

6 Conclusion and Outlook **53**

A Additional diagnostics **57**
A.1 Damped Richardson diagnostic . . . 57
A.2 Local eigenvalue diagnostics . . . 59
A.3 Source-amplitude calibration diagnostics . . . 60
Free-space point-source reference . . . 61
Hard-ground image-source diagnostic . . . 62

Bibliography **65**

List of Figures

3.1 Overview of the implemented solution pipeline. The mesh, physical parameters, and solver parameters define the finite element problem. The domain is partitioned into overlapping subdomains, local spectral problems are solved to construct the coarse basis, and the resulting two-level preconditioner is used within GMRES. 14
3.2 Schematic relation between a core subdomain, its overlapping patch, and its oversampled patch. 16

4.1 Geometries used in the numerical experiments: (a) the unit cube used for exact plane-wave verification, (b) the rectangular box domain used for controlled parameter studies, and (c) the realistic city-domain mesh representing part of Gothenburg. 24
4.2 Convergence of the exact plane-wave verification problem at 5 Hz. The L^2 error decreases approximately quadratically, while the H^1-seminorm error decreases approximately linearly, as expected for first-order Lagrange finite elements. 27
4.3 Example box-domain solution at 10 Hz on the $h = 4$ m mesh, showing the magnitude of the computed complex pressure amplitude. The largest amplitude occurs near the Gaussian source and the field decays away from it across the rectangular domain. The view is from below the domain, as indicated by the coordinate axes, so the visible face is the lower side of the box. 30
4.4 Example city-domain solution at 5 Hz on the $h = 5$ m mesh, showing the magnitude of the computed complex pressure amplitude from two viewing directions. The transparent volume view makes the building geometry visible inside the computational domain. The largest amplitude is located near the Gaussian source, while the surrounding buildings and ground shape the spatial distribution of the field. . . . 37
4.5 Effect of the retained number of local modes on the city-domain 10 Hz, $h = 3$ m problem. Increasing n_i strongly reduces the GMRES iteration count, while the online wall time is not monotone because the cost of the coarse correction also increases. 41
4.6 Relative solution difference for the city-domain 10 Hz, $h = 3$ m retained-mode study, measured against the $n_i = 50$ solution. The differences remain of order 10^{-6}, showing that the retained mode count mainly affects solver convergence rather than the final converged solution. . . 42

A.1 Damped Richardson residual histories for the box-domain 20 Hz, $h = 4\,\mathrm{m}$ diagnostic problem. The dashed horizontal line marks the GMRES tolerance used in the main experiments. 58
A.2 Raw local eigenvalue spectrum for patch 13 in the box-domain 20 Hz, $h = 2\,\mathrm{m}$ diagnostic run. The spectrum is plotted on a semilogarithmic scale to inspect the ordering and magnitude of the local eigenvalues entering the spectral coarse-space construction. 59
A.3 Number of retained local eigenmodes for the box-domain 20 Hz, $h = 2\,\mathrm{m}$ diagnostic run with $M = 16$, one overlap layer, three oversampling layers, and $\rho = 1.10$. 60

List of Tables

4.1 Frequencies used in the experiments and corresponding wavelengths for $c = 343\,\mathrm{m\,s^{-1}}$. . . . 24
4.2 Exact plane-wave verification test at 5 Hz on the unit cube. . . . 27
4.3 Box-domain meshes used in the experiments. . . . 29
4.4 Box-domain mesh-refinement study at 5 Hz. . . . 30
4.5 Box-domain mesh-refinement study at 10 Hz. . . . 30
4.6 Box-domain mesh-refinement study at 20 Hz. . . . 30
4.7 Receiver convergence for the box-domain mesh-refinement studies. . . 31
4.8 Effect of the spectral threshold for the box domain at 10 Hz on the $h = 2.5$ mesh. . . . 32
4.9 Effect of oversampling for the box domain at 10 Hz on the $h = 2.5$ mesh. . . . 32
4.10 Comparison between the one-level and two-level methods for the box domain at 5 Hz. . . . 33
4.11 Comparison between the one-level and two-level methods for the box domain at 10 Hz on the $h = 2.5$ mesh. . . . 33
4.12 Plain GMRES baseline without preconditioning on the box domain with $h = 2.5$. . . . 34
4.13 Effect of the number of subdomains for the box domain at 20 Hz and $h = 2.5$, using $n_i = 100$ retained local modes per subdomain. . . . 34
4.14 Effect of oversampling for the box domain at 20 Hz, $h = 2.5$, $M = 64$, and fixed $n_i = 100$. . . . 35
4.15 Box-domain mesh-refinement study at 30 Hz using a fixed-mode coarse space with $M = 64$, two oversampling layers, and $n_i = 100$. . . . 35
4.16 City-domain meshes used in the experiments. . . . 36
4.17 City-domain mesh-refinement study at 5 Hz with $\rho = 1.10$, $M = 16$, one overlap layer, and three oversampling layers. . . . 38
4.18 Completed city-domain mesh-refinement study at 10 Hz. The $h = 5\,\mathrm{m}$ and $h = 4\,\mathrm{m}$ runs use the threshold-based setup with $M = 16$, while the $h = 3\,\mathrm{m}$ run uses the fixed-mode setup with $M = 64$ and $n_i = 50$ retained modes per subdomain. . . . 38
4.19 Receiver convergence for the city-domain 5 Hz study using the $h = 3\,\mathrm{m}$ solution as reference. . . . 39
4.20 Receiver convergence for the city-domain 10 Hz study using the $h = 3\,\mathrm{m}$ solution as reference. . . . 39

4.21 Effect of the number of retained local modes for the city-domain problem at 10 Hz on the $h = 3$ m mesh. The setup uses $M = 64$, one overlap layer, and two oversampling layers. 40
4.22 City-domain runs at 20 Hz using a fixed number of retained local modes. The $h = 3$ m run uses zero oversampling layers to reduce memory cost. 44
4.23 City-domain feasibility runs at 30 Hz using $M = 64$, one overlap layer, one oversampling layer, and $n_i = 60$ retained local modes per subdomain. 45
4.24 Reduced retained-mode online solves for the city-domain 30 Hz, $h = 3$ m case. The reduced preconditioners were obtained from the saved $n_i = 60$ offline data. 45
4.25 Plain GMRES baseline without preconditioning on the city mesh with $h = 5$ m and $N = 330073$ degrees of freedom. 46

A.1 Damped Richardson diagnostic for the box-domain 20 Hz case with $h = 4$ m, $M = 16$, one overlap layer, three oversampling layers, and $\rho = 1.10$. 58
A.2 Free-field source normalization for the Gaussian source with $\sigma = 5$ m and $A_s = 1$. The scale factor α multiplies the numerical solution so that the equivalent free-space source gives 1 Pa at $r = 1$ m. 61
A.3 Hard-ground Green's-function calibration diagnostic for the 20 Hz, $h = 2$ m box-domain case. The fitted source strength should be approximately constant if the numerical field behaves like an ideal point source above a perfectly reflecting infinite plane. 62

1
Introduction

Urban noise is a major environmental problem in modern cities. Road traffic, railways, construction work, and industrial activity all contribute to sound fields that interact with buildings, terrain, and the ground. Long-term exposure to environmental noise is associated with negative effects on health and quality of life, including annoyance and sleep disturbance [19]. Reliable prediction of sound propagation is therefore important for urban planning, environmental assessment, and the design of quieter cities.

In practical environmental acoustics, many prediction tools are based on geometrical or semi-empirical propagation models. Such methods are efficient and are widely used in engineering practice, for example in standardized outdoor noise prediction methods [10]. However, wave effects such as interference, diffraction, and multiple reflections become important in complex urban geometries, especially at low frequencies. Wave-based models provide a more direct representation of these effects. For steady time-harmonic sound fields, the governing equation is the Helmholtz equation, which is a standard frequency-domain model in acoustics [13, 16, 9].

Solving the Helmholtz equation numerically on realistic three-dimensional urban geometries is computationally difficult. Finite element discretization leads to large sparse complex-valued linear systems. As the frequency increases, the wavelength decreases, so the mesh must be refined to resolve the wave field. This increases the number of degrees of freedom. At the same time, the Helmholtz operator is indefinite, and standard iterative solvers can converge slowly or unreliably without suitable preconditioning [6, 7].

The motivation for this thesis is to investigate whether a recently proposed two-level domain decomposition preconditioner can make such Helmholtz solves practical for large-scale urban noise simulation. The implementation follows the restricted additive Schwarz and multiscale spectral generalized finite element method preconditioner proposed by Ma, Alber, Scheichl, and Zhang [12]. The method combines local subdomain corrections with a global coarse correction constructed from local spectral basis functions. In this work, the method is implemented and tested on both controlled box geometries and real city-domain meshes.

1.1 Problem formulation

The central problem studied in this thesis is how to solve finite element discretizations of the time-harmonic acoustic Helmholtz equation efficiently and robustly on large three-dimensional geometries. The application setting is urban noise simulation, where the computational domain may contain buildings, ground surfaces, and open-air regions.

The main computational challenge is not only to assemble the finite element system, but to solve the resulting linear system with reasonable memory use and wall-clock time. Direct solvers are often too expensive for large three-dimensional problems, while unpreconditioned iterative solvers may converge very slowly for indefinite Helmholtz systems. This makes preconditioning a central part of the numerical method.

This motivates the main task of the thesis: evaluating whether a two-level spectral domain decomposition preconditioner can make finite element Helmholtz solves practical for large-scale urban acoustic simulations, and identifying the limitations that appear as the frequency, mesh resolution, and coarse-space dimension increase.

1.2 Aim and research questions

The aim of the thesis is to implement, test, and evaluate a two-level preconditioned finite element solver for the Helmholtz equation, with emphasis on large-scale acoustic problems on realistic city geometries. It is addressed through the following research questions, which guide the design of the numerical experiments in Chapter 4:

1. How does the two-level preconditioner affect GMRES convergence for Helmholtz problems on controlled box geometries and realistic city-domain meshes?

2. How do frequency and mesh resolution influence the solution behaviour, receiver convergence, and solver performance?

3. How do preconditioner parameters such as the spectral threshold, oversampling size, number of subdomains, and number of retained local modes affect convergence and computational cost?

4. What memory limitations appear when the spectral coarse space is stored explicitly for large three-dimensional city problems?

1.3 Scope and limitations

The thesis focuses on the numerical solution of finite element discretizations of the Helmholtz equation. The main object of study is therefore the linear solver and the two-level preconditioner, rather than a complete physical model for outdoor noise prediction.

The acoustic model is kept deliberately simple. The experiments use time-harmonic linear acoustics with constant material parameters, a localized Gaussian source term, sound-hard solid boundaries, and absorbing conditions on selected outer boundary facets. Detailed source calibration, finite-impedance or porous ground models, atmospheric effects, and validation against measured sound levels are outside the scope of the thesis.

The numerical experiments are designed to study solver behaviour: GMRES convergence, mesh and frequency dependence, preconditioner parameters, and memory limitations. The city-domain simulations are therefore interpreted as realistic geometry tests for the numerical method, not as final validated predictions of urban sound levels. Receiver values are used as internal diagnostics for mesh refinement and solution consistency, not as comparison with physical measurements.

1.4 Thesis outline

Chapter 2 introduces the theoretical background: acoustic wave propagation, the time-harmonic Helmholtz equation, boundary conditions, finite element discretization, GMRES, and preconditioning. Chapter 3 describes the implemented numerical method, including the finite element model, domain decomposition, spectral coarse space, two-level restricted additive Schwarz preconditioner, and offline-online solution procedure.

Chapter 4 presents the numerical experiments and results. It starts with verification on an exact plane-wave problem, then studies the method on a controlled rectangular box domain, and finally applies it to realistic city-domain meshes. Chapter 5 discusses the results in relation to the research questions, including solver convergence, acoustic convergence, parameter choices, and memory limitations. Chapter 6 summarizes the main findings and outlines possible directions for future work.

2

Theory

2.1 Acoustic wave propagation

Sound propagation in air can be described as a small perturbation of an ambient equilibrium state. Let $p_{\mathrm{tot}}(\boldsymbol{x}, t)$ denote the total pressure at position $\boldsymbol{x} \in \Omega \subset \mathbb{R}^3$ and time t. It is written as

$$p_{\mathrm{tot}}(\boldsymbol{x}, t) = p_0 + p(\boldsymbol{x}, t), \tag{2.1}$$

where p_0 is the constant atmospheric pressure and $p(\boldsymbol{x}, t)$ is the acoustic pressure perturbation.

Similarly, the total density is written as

$$\rho_{\mathrm{tot}}(\boldsymbol{x}, t) = \rho_0 + \rho_1(\boldsymbol{x}, t), \tag{2.2}$$

where ρ_0 is the ambient air density and $\rho_1(\boldsymbol{x}, t)$ is the density perturbation. In the small-amplitude regime, where nonlinear effects can be neglected, acoustic wave propagation is governed by the linearized equations of mass and momentum conservation [13, 16]. If $\boldsymbol{v}(\boldsymbol{x}, t)$ is the acoustic particle velocity, these equations can be written as

$$\frac{\partial \rho_1}{\partial t} + \rho_0 \nabla \cdot \boldsymbol{v} = 0, \tag{2.3}$$

and

$$\rho_0 \frac{\partial \boldsymbol{v}}{\partial t} + \nabla p = \mathbf{0}. \tag{2.4}$$

To close the system, pressure and density perturbations are related by the linear equation of state

$$p = c^2 \rho_1, \tag{2.5}$$

where c is the speed of sound.

Combining (2.3), (2.4), and (2.5) gives the acoustic wave equation

$$\frac{1}{c^2} \frac{\partial^2 p}{\partial t^2} - \Delta p = s \quad \text{in } \Omega, \tag{2.6}$$

where $s(\boldsymbol{x}, t)$ represents an acoustic source term. This equation models the propagation of pressure waves through a homogeneous medium. It is a standard starting point for outdoor sound propagation models when the acoustic amplitudes are small compared with the ambient pressure [13].

In the context of urban noise simulation, the pressure field is affected by reflections from buildings, ground surfaces, and terrain. Geometrical acoustic methods are widely used in environmental noise prediction and are also part of standardized engineering models [10, 16]. However, in geometrically complicated regions such as street canyons, courtyards, and dense urban environments, wave effects and multiple reflections can become important. A wave-based model such as (2.6) is therefore a natural mathematical model when the aim is to resolve the acoustic pressure field in a complex geometry.

The full time-dependent problem is often more information than is needed for steady or tonal sources. For such sources, it is common to study the response at one frequency at a time. This leads to a frequency-domain formulation, which is the starting point for the Helmholtz equation introduced in the next section.

2.2 Time-harmonic Helmholtz equation

For steady or tonal sources, the acoustic field can be studied at a single frequency. This frequency-domain formulation is standard in acoustic scattering and time-harmonic wave propagation [13, 9]. In this case, both the acoustic pressure and the source are assumed to oscillate sinusoidally in time with angular frequency ω. Instead of solving for the full time-dependent pressure field, one solves for complex amplitudes that contain the spatial variation, phase, and magnitude of the oscillation.

The pressure perturbation is written as

$$p(\boldsymbol{x}, t) = \operatorname{Re}\left(u(\boldsymbol{x})e^{-i\omega t}\right), \tag{2.7}$$

where $u(\boldsymbol{x})$ is the complex-valued pressure amplitude. Similarly, the source term is written as

$$s(\boldsymbol{x}, t) = \operatorname{Re}\left(q(\boldsymbol{x})e^{-i\omega t}\right), \tag{2.8}$$

where $q(\boldsymbol{x})$ is the complex source amplitude. The physical quantities are obtained by taking the real parts of these expressions. The complex notation is used only as a convenient way to represent both amplitude and phase.

The angular frequency ω, frequency f, wavenumber k, and wavelength λ are related by

$$\omega = 2\pi f, \qquad k = \frac{\omega}{c}, \qquad \lambda = \frac{2\pi}{k} = \frac{c}{f}. \tag{2.9}$$

Thus, increasing the frequency decreases the wavelength. This is important for the numerical problem, since a shorter wavelength requires a finer mesh in order to resolve the oscillations of the solution.

Substituting the expressions for $p(\boldsymbol{x}, t)$ and $s(\boldsymbol{x}, t)$ into the acoustic wave equation gives

$$\operatorname{Re}\left[\frac{1}{c^2}\frac{\partial^2}{\partial t^2}\left(u(\boldsymbol{x})e^{-i\omega t}\right) - \Delta\left(u(\boldsymbol{x})e^{-i\omega t}\right)\right] = \operatorname{Re}\left(q(\boldsymbol{x})e^{-i\omega t}\right). \tag{2.10}$$

Since $u(\boldsymbol{x})$ is independent of time,

$$\frac{\partial^2}{\partial t^2}\left(u(\boldsymbol{x})e^{-i\omega t}\right) = -\omega^2 u(\boldsymbol{x})e^{-i\omega t}. \tag{2.11}$$

After cancelling the common time-dependent factor, one obtains

$$-\Delta u(\boldsymbol{x}) - \frac{\omega^2}{c^2} u(\boldsymbol{x}) = q(\boldsymbol{x}). \tag{2.12}$$

Using $k = \omega/c$, this becomes

$$-\Delta u(\boldsymbol{x}) - k^2 u(\boldsymbol{x}) = q(\boldsymbol{x}) \quad \text{in } \Omega. \tag{2.13}$$

Since the right-hand side $q(\boldsymbol{x})$ represents an acoustic source, (2.13) is an inhomogeneous Helmholtz equation. If no source were present, then $q(\boldsymbol{x}) = 0$, and the equation would be homogeneous. The Helmholtz equation is the standard frequency-domain model for linear time-harmonic acoustics [13, 9, 16].

It is used in this thesis because it describes the acoustic field at one fixed frequency. This is different from the wave equation, which describes how the pressure changes in time. For a time-harmonic source, the time dependence is already known: the field oscillates like $e^{-i\omega t}$. Therefore, the unknown part is only the spatial complex amplitude $u(\boldsymbol{x})$.

Solving for $u(\boldsymbol{x})$ gives the magnitude and phase of the acoustic pressure at the chosen frequency. This is useful for urban noise simulations when selected frequencies are studied separately, for example low frequencies or tonal source components [16, 13]. The disadvantage is that the Helmholtz equation is numerically difficult to solve. The term $-k^2 u$ makes the operator indefinite, and the difficulty increases as the frequency f, and therefore the wavenumber k, increases [9, 6]. This issue is discussed later in this chapter.

2.3 Boundary conditions for acoustic domains

To obtain a well-posed acoustic boundary value problem, the Helmholtz equation must be supplemented with boundary conditions. In outdoor sound propagation, different parts of the boundary have different physical meanings. Solid surfaces, such as rigid building facades or hard ground, may be modeled as reflecting boundaries. Artificial outer boundaries, introduced only because the computational domain must be finite, should instead allow outgoing waves to leave the domain with as little reflection as possible [13, 16, 9].

A common idealization for a rigid surface is the sound-hard boundary condition. Physically, this means that the normal component of the acoustic particle velocity is zero. For the frequency-domain pressure amplitude $u(\boldsymbol{x})$, this gives the Neumann condition

$$\nabla u(\boldsymbol{x}) \cdot \boldsymbol{n} = 0 \quad \text{on } \Gamma_N, \tag{2.14}$$

where $\boldsymbol{n}$ is the outward unit normal vector and Γ_N is the sound-hard part of the boundary. This condition represents perfect reflection at the boundary and is a standard first approximation for rigid acoustic surfaces [13, 16].

For an open acoustic domain, the physical domain is unbounded, while the numerical domain is finite. The boundary introduced by truncating the domain is not a physical wall. It should therefore be modeled so that waves can leave the computational domain. A simple first-order absorbing, or impedance, boundary condition is

$$\nabla u(\boldsymbol{x}) \cdot \boldsymbol{n} - iku(\boldsymbol{x}) = 0 \quad \text{on } \Gamma_R. \tag{2.15}$$

This condition is exact for a plane wave leaving the boundary in the normal direction, but only approximate for more general outgoing waves [9]. More accurate treatments of artificial boundaries include absorbing layers, perfectly matched layers, and infinite elements [9, 16].

The boundary conditions used in this thesis are based on these two ideas: sound-hard conditions for reflecting surfaces and first-order absorbing conditions for artificial open boundaries. The precise assignment of these conditions in the computations is described in Chapter 3. More detailed acoustic boundary models, such as frequency-dependent facade absorption or impedance models for porous ground, are outside the scope of this work. The aim here is to keep the physical model controlled while studying the numerical solution of the resulting Helmholtz systems.

2.4 Finite element discretization

Closed-form solutions of the Helmholtz equation are available only in special cases, such as simple geometries, homogeneous media, and idealized boundary conditions. For realistic three-dimensional urban geometries, no such analytical solution is available, and a numerical discretization is required. The finite element method is well suited for this purpose because it can handle unstructured meshes and complicated geometrical boundaries [2, 9].

Consider the Helmholtz problem

$$-\Delta u - k^2 u = q \quad \text{in } \Omega, \tag{2.16}$$

with sound-hard boundary condition on Γ_N and absorbing boundary condition on Γ_R. To derive the weak form, the equation is multiplied by a complex-valued test function v and integrated over the domain:

$$\int_\Omega \left(-\Delta u - k^2 u\right) \overline{v} \, \mathrm{d}\boldsymbol{x} = \int_\Omega q\overline{v} \, \mathrm{d}\boldsymbol{x}. \tag{2.17}$$

Here $\overline{v}$ denotes the complex conjugate of v. The normal derivative is defined by

$$\frac{\partial u}{\partial n} = \nabla u \cdot \boldsymbol{n}, \tag{2.18}$$

where $\boldsymbol{n}$ is the outward unit normal vector. Integrating the Laplacian term by parts gives

$$\int_\Omega \nabla u \cdot \nabla \overline{v} \, \mathrm{d}\boldsymbol{x} - k^2 \int_\Omega u\overline{v} \, \mathrm{d}\boldsymbol{x} - \int_{\partial\Omega} \frac{\partial u}{\partial n} \overline{v} \, \mathrm{d}s = \int_\Omega q\overline{v} \, \mathrm{d}\boldsymbol{x}. \tag{2.19}$$

On the sound-hard boundary, $\partial u/\partial n = 0$. On the absorbing boundary, using (2.15) gives

$$\frac{\partial u}{\partial n} = iku \quad \text{on } \Gamma_R. \tag{2.20}$$

Therefore the weak formulation is: find $u \in V$ such that

$$a(u, v) = \ell(v) \quad \forall v \in V, \tag{2.21}$$

where

$$a(u, v) = \int_\Omega \nabla u \cdot \nabla \overline{v} \, \mathrm{d}\boldsymbol{x} - k^2 \int_\Omega u\overline{v} \, \mathrm{d}\boldsymbol{x} - ik \int_{\Gamma_R} u\overline{v} \, \mathrm{d}s, \tag{2.22}$$

and

$$\ell(v) = \int_\Omega q\overline{v} \, \mathrm{d}\boldsymbol{x}. \tag{2.23}$$

The finite element method approximates the infinite-dimensional space V by a finite-dimensional space V_h. In this thesis, V_h consists of continuous piecewise linear Lagrange basis functions on a tetrahedral mesh. Let $\{\phi_j\}_{j=1}^N$ be the finite element basis. The discrete solution is written as

$$u_h(\boldsymbol{x}) = \sum_{j=1}^{N} u_j \phi_j(\boldsymbol{x}), \tag{2.24}$$

where $u_j \in \mathbb{C}$ are scalar unknown coefficients. These coefficients are collected in the vector

$$\boldsymbol{u} = \begin{bmatrix} u_1 & u_2 & \cdots & u_N \end{bmatrix}^T \in \mathbb{C}^N. \tag{2.25}$$

Substituting this finite element approximation into the weak formulation and testing with each basis function ϕ_i, for $i = 1, \ldots, N$, gives a linear system

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b}. \tag{2.26}$$

The entries of the matrix $\boldsymbol{A}$ are

$$A_{ij} = a(\phi_j, \phi_i), \qquad i, j = 1, \ldots, N, \tag{2.27}$$

and the entries of the right-hand side vector $\boldsymbol{b}$ are

$$b_i = \ell(\phi_i), \qquad i = 1, \ldots, N. \tag{2.28}$$

It is useful to decompose the matrix into stiffness, mass, and boundary contributions:

$$\boldsymbol{A} = \boldsymbol{K} - k^2\boldsymbol{M} - ik\boldsymbol{B}. \tag{2.29}$$

Here the entries of the stiffness matrix $\boldsymbol{K}$ are

$$K_{ij} = \int_\Omega \nabla\phi_j \cdot \nabla\overline{\phi_i} \, \mathrm{d}\boldsymbol{x}, \tag{2.30}$$

the entries of the mass matrix $\boldsymbol{M}$ are

$$M_{ij} = \int_\Omega \phi_j\overline{\phi_i} \, \mathrm{d}\boldsymbol{x}, \tag{2.31}$$

and the entries of the absorbing boundary matrix $\boldsymbol{B}$ are

$$B_{ij} = \int_{\Gamma_R} \phi_j\overline{\phi_i} \, \mathrm{d}s. \tag{2.32}$$

This matrix structure is important for the numerical solution. The matrix $\boldsymbol{K}$ comes from the Laplacian term, while $\boldsymbol{M}$ comes from the k^2u term. The negative contribution $-k^2\boldsymbol{M}$ is one reason why the Helmholtz matrix $\boldsymbol{A}$ is indefinite. This makes the resulting linear systems more difficult to solve than the positive definite systems that arise from elliptic problems such as the Poisson equation [9, 6].

2.5 Why Helmholtz problems are difficult

The matrix structure in (2.29) explains why Helmholtz problems are harder to solve than many elliptic finite element problems. The main difficulties are the oscillatory nature of the solution, the indefiniteness of the algebraic system, and the accumulation of numerical phase errors [9, 6].

To explain indefiniteness, it is useful to first recall the meaning of positive definiteness. A Hermitian matrix $\boldsymbol{K}$ is positive definite if

$$\boldsymbol{w}^*\boldsymbol{K}\boldsymbol{w} > 0 \quad \text{for every nonzero } \boldsymbol{w}. \tag{2.33}$$

The quantity $\boldsymbol{w}^*\boldsymbol{K}\boldsymbol{w}$ can be interpreted as an energy associated with the vector $\boldsymbol{w}$. For positive definite elliptic problems, such as the Poisson equation, this energy is always positive. This gives the problem a stable minimization structure and is one reason why many standard iterative methods and preconditioners work well for elliptic problems [15, 18].

A matrix is indefinite if the corresponding quadratic form can be positive in some directions and negative in others. That is, for an indefinite matrix there may exist vectors $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ such that

$$\boldsymbol{w}_1^*\boldsymbol{A}\boldsymbol{w}_1 > 0, \tag{2.34}$$

while

$$\boldsymbol{w}_2^*\boldsymbol{A}\boldsymbol{w}_2 < 0. \tag{2.35}$$

In geometric terms, a positive definite problem behaves like a bowl with a single energy minimum, while an indefinite problem behaves more like a saddle: some directions increase the quadratic form and other directions decrease it.

For the Helmholtz equation, this indefiniteness comes from the negative k^2u term. Ignoring the absorbing boundary contribution for the moment, the finite element matrix has the form

$$\boldsymbol{A} = \boldsymbol{K} - k^2\boldsymbol{M}, \tag{2.36}$$

where $\boldsymbol{K}$ is the stiffness matrix and $\boldsymbol{M}$ is the mass matrix. Both $\boldsymbol{K}$ and $\boldsymbol{M}$ are positive matrices in the usual finite element setting, but the Helmholtz operator subtracts the mass contribution. Therefore,

$$\boldsymbol{w}^*\boldsymbol{A}\boldsymbol{w} = \boldsymbol{w}^*\boldsymbol{K}\boldsymbol{w} - k^2\boldsymbol{w}^*\boldsymbol{M}\boldsymbol{w}. \tag{2.37}$$

For some vectors $\boldsymbol{w}$, the stiffness contribution dominates and the expression is positive. For other vectors, the mass contribution dominates and the expression is negative. This is the algebraic meaning of the indefiniteness of the Helmholtz problem.

This matters for iterative solution. For positive definite problems, many methods can be understood as reducing an energy or damping error components. For indefinite Helmholtz problems, this simple damping picture breaks down. Some error components are not naturally reduced by classical relaxation or elliptic preconditioners, and they may persist or even be amplified. This is one reason why classical iterative methods that are effective for Poisson-type problems can perform poorly for Helmholtz problems [6].

The difficulty increases with frequency. The wavenumber is

$$k = \frac{2\pi f}{c}. \tag{2.38}$$

Thus, increasing the frequency f increases k, and the negative mass term $-k^2\boldsymbol{M}$ becomes stronger. At the same time, the wavelength is

$$\lambda = \frac{c}{f}. \tag{2.39}$$

A higher frequency therefore gives a shorter wavelength, so the acoustic field oscillates more rapidly in space. The mesh must be fine enough to resolve these oscillations. A useful measure is the number of mesh intervals per wavelength,

$$N_\lambda = \frac{\lambda}{h} = \frac{c}{fh}, \tag{2.40}$$

where h is the mesh size. If N_λ is too small, the discrete space cannot accurately represent the wave field.

Even when the wave is locally resolved, Helmholtz discretizations can suffer from pollution error. This means that the numerical wave may propagate with a slightly incorrect phase velocity. Over many wavelengths, small phase errors can accumulate into a significant global error. As a result, Helmholtz problems often require finer meshes than a simple local points-per-wavelength argument would suggest [9].

Thus, increasing the frequency makes the problem harder in two ways. First, the mesh must be refined, which increases the number of degrees of freedom. Second, the algebraic system becomes more indefinite and more difficult for iterative solvers. The computational cost therefore grows not only because the matrix is larger, but also because the matrix is harder to solve. This is why robust preconditioning is essential for large-scale Helmholtz simulations.

2.6 Iterative solvers and GMRES

For the large sparse systems considered in this thesis, direct solvers may become too expensive in both memory and computational time. Iterative methods are therefore commonly used for large sparse systems [15].

GMRES is a Krylov subspace method for general nonsymmetric linear systems [14, 15]. Given an initial guess $\boldsymbol{u}_0$, the initial residual is

$$\boldsymbol{r}_0 = \boldsymbol{b} - \boldsymbol{A}\boldsymbol{u}_0. \tag{2.41}$$

The m-th Krylov subspace is

$$\mathcal{K}_m(\boldsymbol{A}, \boldsymbol{r}_0) = \operatorname{span}\left\{\boldsymbol{r}_0, \boldsymbol{A}\boldsymbol{r}_0, \ldots, \boldsymbol{A}^{m-1}\boldsymbol{r}_0\right\}. \tag{2.42}$$

GMRES chooses the approximate solution from $\boldsymbol{u}_0 + \mathcal{K}_m(\boldsymbol{A}, \boldsymbol{r}_0)$ that minimizes the residual norm:

$$\boldsymbol{u}_m = \arg \min_{\boldsymbol{w} \in \boldsymbol{u}_0 + \mathcal{K}_m(\boldsymbol{A}, \boldsymbol{r}_0)} \left\| \boldsymbol{b} - \boldsymbol{A}\boldsymbol{w} \right\|_2. \tag{2.43}$$

GMRES is appropriate for the Helmholtz systems considered in this thesis because the matrices are complex-valued and indefinite, and they are not generally symmetric positive definite. However, for Helmholtz problems, the number of iterations can grow rapidly as the frequency and problem size increase [6]. Therefore, GMRES is usually used together with a preconditioner.

2.7 Preconditioning

For Helmholtz problems, the convergence of GMRES depends strongly on the properties of the matrix $\boldsymbol{A}$. Since this matrix is large, sparse, complex-valued, and indefinite, applying GMRES directly may require many iterations or may converge unreliably. A preconditioner is introduced to transform the system into an equivalent system that is easier for the iterative method to solve [15, 18].

A left preconditioner is a matrix or linear operator $\boldsymbol{\mathcal{M}}^{-1}$ that approximates the action of $\boldsymbol{A}^{-1}$. Instead of applying GMRES directly to the original system, one applies GMRES to the left-preconditioned system

$$\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{A}\boldsymbol{u} = \boldsymbol{\mathcal{M}}^{-1}\boldsymbol{b}. \tag{2.44}$$

The exact solution is unchanged, but the convergence behaviour of the iterative method can be significantly improved if $\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{A}$ has more favorable properties than $\boldsymbol{A}$.

In an ideal case, one would choose $\boldsymbol{\mathcal{M}}^{-1} = \boldsymbol{A}^{-1}$. Then

$$\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{A} = \boldsymbol{I}, \tag{2.45}$$

and the preconditioned system would be trivial to solve. In practice, constructing and applying the exact inverse is too expensive. A useful preconditioner must therefore be much cheaper to apply than solving the original system, while still capturing enough of the structure of $\boldsymbol{A}$ to improve convergence.

For positive definite elliptic problems, many standard preconditioners are very effective. Helmholtz problems are more difficult because the operator is indefinite and wave-like. In particular, local error components and global wave components can both affect convergence. A preconditioner that works well for the Poisson equation may therefore perform poorly for the Helmholtz equation [6].

This motivates preconditioners that combine local and global information. Local components are useful because they are relatively cheap and can exploit the sparsity of the finite element matrix. However, local corrections alone may not remove error components that extend across the domain. For large Helmholtz problems, especially at higher frequencies, an effective preconditioner must also include a mechanism for representing global wave-like error components.

The specific preconditioner used in this work is a two-level domain decomposition preconditioner with a spectral coarse space. The next chapter describes this method, following the construction of Ma, Alber, Scheichl, and Zhang [12], and explains how it is implemented for the numerical experiments.

3

Methods

3.1 Overview of the numerical method

This chapter describes how the finite element Helmholtz formulation from Chapter 2 is used in the implemented solver. The purpose is not to repeat the derivation of the Helmholtz equation or GMRES, but to describe how the discrete problem is assembled, decomposed into local patch problems, enriched by a spectral coarse space, and solved with a two-level preconditioned Krylov method. The overall workflow is shown in Figure 3.1.

The mesh, frequency, source term, and boundary markers define the global finite element system

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b}. \tag{3.1}$$

The mesh is then decomposed into non-overlapping core subdomains. These are enlarged by cell layers to form overlapping patches and oversampled patches,

$$\Omega_i^{\mathrm{core}} \subset \Omega_i \subset \Omega_i^{\star}. \tag{3.2}$$

Here Ω_i^{core} is the core subdomain, Ω_i is the overlapping patch used for local Schwarz corrections, and $\Omega_i^{\star}$ is the oversampled patch used for local operators and local spectral problems.

Selected local spectral modes are weighted by partition-of-unity functions and assembled into a global coarse basis $\boldsymbol{E}_0$. GMRES is then applied to the global system using a two-level preconditioner that combines local restricted additive Schwarz corrections with a global coarse correction.

3.2 Finite element model used in the experiments

All main experiments use first-order Lagrange finite elements on tetrahedral volume meshes. The discrete space is

$$V_h = \mathrm{span}\{\varphi_1, \dots, \varphi_N\} \subset H^1(\Omega; \mathbb{C}), \tag{3.3}$$

where N is the number of finite element degrees of freedom. The numerical solution is written as

$$u_h(\boldsymbol{x}) = \sum_{j=1}^{N} u_j \varphi_j(\boldsymbol{x}), \qquad \boldsymbol{u} = (u_1, \dots, u_N)^T \in \mathbb{C}^N. \tag{3.4}$$

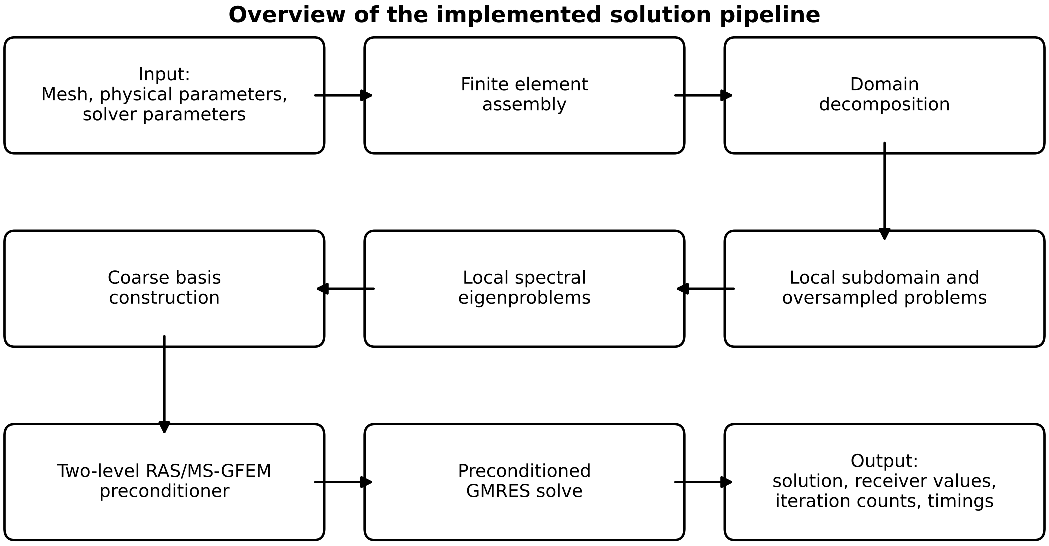

Figure 3.1: Overview of the implemented solution pipeline. The mesh, physical parameters, and solver parameters define the finite element problem. The domain is partitioned into overlapping subdomains, local spectral problems are solved to construct the coarse basis, and the resulting two-level preconditioner is used within GMRES.

The coefficients u_j are complex pressure amplitudes.

The implementation assembles the discrete variational problem

$$a_h(u_h, v_h) = \ell_h(v_h) \qquad \text{for all } v_h \in V_h. \tag{3.5}$$

The bilinear form used in the reported experiments is

$$\begin{aligned} a_h(u, v) = &\int_\Omega \alpha \nabla u \cdot \nabla \overline{v} \, \mathrm{d}\boldsymbol{x} - k^2 \int_\Omega \nu^2 u \overline{v} \, \mathrm{d}\boldsymbol{x} \\ &- \mathrm{i}k \sum_{\tau \in \mathcal{T}_R} \int_{\Gamma_\tau} \beta u \overline{v} \, \mathrm{d}s. \end{aligned} \tag{3.6}$$

Here $k = 2\pi f / c$, with sound speed $c = 343\,\mathrm{m\,s^{-1}}$. In the reported experiments,

$$\alpha = 1, \qquad \nu = 1, \qquad \beta = 1. \tag{3.7}$$

The absorbing boundary term is applied only on the outer boundary facets whose mesh markers belong to

$$\mathcal{T}_R = \{-2, -3, -4, -5, -6\}. \tag{3.8}$$

All other physical boundary facets are left as natural Neumann boundaries, corresponding to sound-hard walls or ground surfaces in the simplified acoustic model. No Dirichlet boundary condition or anchor point is used in the reported experiments.

The source term in the main experiments is a localized Gaussian forcing term. The right-hand side is

$$\ell_h(v) = -\int_\Omega g_h(\boldsymbol{x})\overline{v}\,\mathrm{d}\boldsymbol{x}, \tag{3.9}$$

where

$$g_h(\boldsymbol{x}) = A_s \exp\left(-\frac{\|\boldsymbol{x}-\boldsymbol{x}_0\|^2}{2\sigma^2}\right). \tag{3.10}$$

The numerical source parameters are

$$A_s = 1, \qquad \sigma = 5\,\mathrm{m}. \tag{3.11}$$

With the sign convention used here, the effective source in the Helmholtz equation is $q_h = -g_h$. This sign affects the phase of the complex solution, but not the magnitude-based diagnostics used in the results.

Unless the source location is specified explicitly, it is chosen from the mesh bounding box. If

$$\boldsymbol{x}_{\min} = (x_{\min}, y_{\min}, z_{\min}), \qquad \boldsymbol{x}_{\max} = (x_{\max}, y_{\max}, z_{\max}), \tag{3.12}$$

then the default source position is

$$\boldsymbol{x}_0 = \left(\frac{x_{\min}+x_{\max}}{2}, \frac{y_{\min}+y_{\max}}{2}, 0.9z_{\min}+0.1z_{\max}\right). \tag{3.13}$$

Using the basis functions in (3.3), the discrete problem becomes

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b}, \tag{3.14}$$

with entries

$$A_{ij} = a_h(\varphi_j, \varphi_i), \qquad b_i = \ell_h(\varphi_i). \tag{3.15}$$

3.3 Domain decomposition and partition of unity

The preconditioner is based on a cell-wise domain decomposition of the tetrahedral mesh. The mesh cells are first divided into M non-overlapping core subdomains using METIS.[1] The partition is constructed on the cell adjacency graph, so neighbouring tetrahedra define the graph connectivity.

For each core subdomain, two larger regions are constructed by adding layers in the cell adjacency graph. The first enlargement gives the overlapping patch Ω_i, which is used for the local Schwarz correction. The second enlargement gives the oversampled patch $\Omega_i^\star$, which is used when assembling local operators and solving the local spectral problems:

$$\Omega_i^{\mathrm{core}} \subset \Omega_i \subset \Omega_i^\star. \tag{3.16}$$

The overlap and oversampling have different roles. The overlapping patch determines where the local correction is inserted back into the global vector. The

[1] METIS is a graph partitioning library used here to divide the finite element cell-adjacency graph into non-overlapping subdomains with approximately balanced cell counts and relatively small interfaces between subdomains.

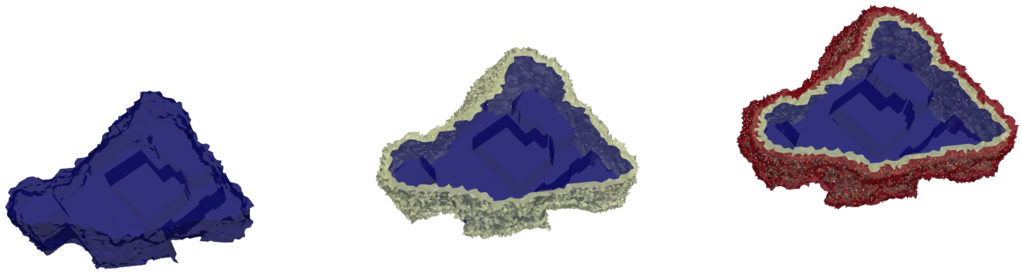

Figure 3.2: Schematic relation between a core subdomain, its overlapping patch, and its oversampled patch.

oversampled patch gives a larger local region for the spectral problem, so that the selected modes are less sensitive to artificial patch boundaries.

Because overlapping patches share degrees of freedom, local corrections and local basis functions are combined using partition-of-unity weights. The implementation uses a graph-layer ramp based on the distance from the core subdomain inside the overlapping patch. Let ℓ_{ov} denote the number of overlap layers. For a cell $K \subset \Omega_i$, let $d_i(K)$ be the cell-graph distance from K to the core cells of Ω_i^{core}, measured inside Ω_i. Thus $d_i(K) = 0$ on the core cells. The raw cell weight is

$$\widetilde{\chi}_i(K) = \max\left(0, 1 - \frac{d_i(K)}{\ell_{\mathrm{ov}} + 1}\right). \tag{3.17}$$

The raw nodal weight is then obtained by taking the maximum raw cell weight over the cells incident to the degree of freedom,

$$\widetilde{\chi}_i(a) = \max_{K \ni a} \widetilde{\chi}_i(K), \tag{3.18}$$

and the raw weight is set to zero outside the overlapping patch. Degrees of freedom belonging to the core are explicitly assigned raw weight one.

The raw weights are normalized over all patches that contain the degree of freedom:

$$\chi_i(a) = \frac{\widetilde{\chi}_i(a)}{\sum_{j=1}^{M} \widetilde{\chi}_j(a)}. \tag{3.19}$$

Consequently, on all covered degrees of freedom,

$$\sum_{i=1}^{M} \chi_i(a) = 1. \tag{3.20}$$

It is important to distinguish the raw and normalized weights. The raw weight is one on the core degrees of freedom of a patch, but after normalization the final weight χ_i may be smaller than one at a degree of freedom also covered by other overlapping patches. The final normalized weights are the ones used in the solver.

The same weights are used in two places. First, they weight the local Schwarz corrections before these corrections are inserted into the global vector. Second, they weight the selected local spectral modes before the modes are assembled into the global coarse basis.

3.4 Spectral coarse space

The local Schwarz correction removes error components that are local to each patch. For Helmholtz problems, however, global wave-like error components can extend across many subdomains. The purpose of the coarse space is to represent such components with a smaller set of global basis functions. The construction used here follows the spectral MS-GFEM idea of Ma et al. [12], but is implemented directly on the finite element degrees of freedom of each oversampled patch.

For each oversampled patch $\Omega_i^\star$, the implementation assembles three local matrices. The local stiffness matrix is

$$(\boldsymbol{S}_i)_{mn} = \int_{\Omega_i^\star} \alpha \nabla\varphi_n \cdot \nabla\overline{\varphi_m}\,\mathrm{d}\boldsymbol{x}. \tag{3.21}$$

The positive Helmholtz energy matrix is

$$(\boldsymbol{H}_i)_{mn} = \int_{\Omega_i^\star} \alpha \nabla\varphi_n \cdot \nabla\overline{\varphi_m}\,\mathrm{d}\boldsymbol{x} + k^2 \int_{\Omega_i^\star} \nu^2 \varphi_n \overline{\varphi_m}\,\mathrm{d}\boldsymbol{x}. \tag{3.22}$$

The local Helmholtz matrix is

$$\begin{aligned}(\boldsymbol{B}_i)_{mn} = {} & \int_{\Omega_i^\star} \alpha \nabla\varphi_n \cdot \nabla\overline{\varphi_m}\,\mathrm{d}\boldsymbol{x} - k^2 \int_{\Omega_i^\star} \nu^2 \varphi_n \overline{\varphi_m}\,\mathrm{d}\boldsymbol{x} \\ & - \mathrm{i}k \int_{\Gamma_i^{\star,\mathrm{imp}}} \beta \varphi_n \overline{\varphi_m}\,\mathrm{d}s. \end{aligned} \tag{3.23}$$

Here $\Gamma_i^{\star,\mathrm{imp}}$ denotes the parts of the patch boundary where an impedance term is applied. This includes artificial patch boundaries and physical absorbing facets. Physical sound-hard facets are left as natural Neumann boundaries.

The partition-of-unity weight is included in the local spectral problem. Let

$$\boldsymbol{X}_i = \mathrm{diag}(\chi_i) \tag{3.24}$$

on the degrees of freedom of $\Omega_i^\star$. The weighted energy matrix is

$$\widehat{\boldsymbol{H}}_i = \boldsymbol{X}_i^* \boldsymbol{H}_i \boldsymbol{X}_i. \tag{3.25}$$

The augmented spectral problem also uses a restriction matrix $\boldsymbol{C}_i$, which selects the interior degrees of freedom of the oversampled patch:

$$\boldsymbol{C}_i \in \mathbb{R}^{m_i \times n_i^\star}, \tag{3.26}$$

where $n_i^\star$ is the number of patch degrees of freedom and m_i is the number of interior patch degrees of freedom.

The augmented matrices are

$$\boldsymbol{\mathcal{K}}_i = \begin{bmatrix} \boldsymbol{S}_i & \boldsymbol{B}_i^* \boldsymbol{C}_i^* \\ \boldsymbol{C}_i \boldsymbol{B}_i & \boldsymbol{0} \end{bmatrix} \tag{3.27}$$

and

$$\boldsymbol{\mathcal{G}}_i = \begin{bmatrix} \widehat{\boldsymbol{H}}_i & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \end{bmatrix}. \tag{3.28}$$

The local spectral problem is solved numerically as

$$\boldsymbol{\mathcal{G}}_i \begin{bmatrix} \boldsymbol{\phi}_{i,j} \\ \boldsymbol{\eta}_{i,j} \end{bmatrix} = \lambda_{i,j} \left(\boldsymbol{\mathcal{K}}_i + \varepsilon \boldsymbol{I} \right) \begin{bmatrix} \boldsymbol{\phi}_{i,j} \\ \boldsymbol{\eta}_{i,j} \end{bmatrix}. \tag{3.29}$$

Here ε is a small regularization shift used in the eigenvalue solve. The vector $\boldsymbol{\phi}_{i,j}$ contains the local mode values on the oversampled patch, while $\boldsymbol{\eta}_{i,j}$ is an auxiliary vector associated with the interior constraint. Only $\boldsymbol{\phi}_{i,j}$ is retained for the coarse space.

The computed eigenvalues are ordered by decreasing real part. In the threshold-based version, the retained modes satisfy

$$\mathrm{Re}(\lambda_{i,j}) > \rho^2, \tag{3.30}$$

where ρ is the spectral threshold parameter. In fixed-mode experiments, the same eigenproblem is solved, but instead of applying the threshold rule a prescribed number n_i of modes is kept on each patch. Since the modes are ordered by decreasing real part of the eigenvalue, this keeps the n_i largest modes on each patch.

After selection, each local mode is normalized using the local stiffness matrix:

$$\mathrm{Re} \left(\boldsymbol{\phi}_{i,j}^{*} \boldsymbol{S}_i \boldsymbol{\phi}_{i,j} \right) = 1. \tag{3.31}$$

Each retained local mode is embedded into the global finite element vector on the degrees of freedom of $\Omega_i^{\star}$, multiplied by the corresponding partition-of-unity weight, and then inserted into the global coarse basis:

$$\boldsymbol{\psi}_{i,j} = I_h \left(\chi_i \boldsymbol{\phi}_{i,j} \right). \tag{3.32}$$

For the first-order Lagrange elements used here, this corresponds to nodal multiplication by χ_i.

The global coarse basis matrix is

$$\boldsymbol{E}_0 = \begin{bmatrix} \boldsymbol{\psi}_{1,1} & \cdots & \boldsymbol{\psi}_{1,n_1} & \boldsymbol{\psi}_{2,1} & \cdots & \boldsymbol{\psi}_{M,n_M} \end{bmatrix}. \tag{3.33}$$

Thus,

$$\boldsymbol{E}_0 \in \mathbb{C}^{N \times n_c}, \qquad n_c = \sum_{i=1}^{M} n_i. \tag{3.34}$$

For fixed-mode experiments with the same number of retained modes on each patch,

$$n_c = M n_i. \tag{3.35}$$

3.5 Two-level preconditioner

The coarse basis from Section 3.4 is used together with local patch solves to define the preconditioner applied inside GMRES. The preconditioner is not stored as a full inverse matrix. It is applied as an operation mapping an input residual to an approximate correction.

Let $\boldsymbol{r} \in \mathbb{C}^N$ be the input vector. The first part is a local restricted additive Schwarz correction. For each patch, the residual is restricted to the oversampled patch, a local Helmholtz problem is solved, the result is restricted back to the overlapping patch, weighted by the partition-of-unity weights, and inserted into the global vector:

$$\boldsymbol{z}_{\mathrm{loc}} = \sum_{i=1}^{M} \boldsymbol{R}_{i,\Omega}^{T} \boldsymbol{D}_i \boldsymbol{P}_i \boldsymbol{B}_i^{-1} \boldsymbol{R}_{i,\star} \boldsymbol{r}. \tag{3.36}$$

Here $\boldsymbol{R}_{i,\star}$ restricts a global vector to the oversampled patch $\Omega_i^{\star}$, $\boldsymbol{B}_i$ is the local Helmholtz matrix on that patch, $\boldsymbol{P}_i$ restricts the local patch solution to the overlapping patch Ω_i, $\boldsymbol{D}_i$ is the diagonal matrix of partition-of-unity weights on Ω_i, and $\boldsymbol{R}_{i,\Omega}^{T}$ injects the weighted local correction into the global vector.

After the local correction, the remaining residual is

$$\boldsymbol{r}_0 = \boldsymbol{r} - \boldsymbol{A}\boldsymbol{z}_{\mathrm{loc}}. \tag{3.37}$$

The coarse matrix is formed by Galerkin projection:

$$\boldsymbol{A}_0 = \boldsymbol{E}_0^{*} \boldsymbol{A} \boldsymbol{E}_0. \tag{3.38}$$

The coarse correction is

$$\boldsymbol{z}_{\mathrm{coarse}} = \boldsymbol{E}_0 \boldsymbol{A}_0^{-1} \boldsymbol{E}_0^{*} \boldsymbol{r}_0. \tag{3.39}$$

The complete two-level preconditioner action is therefore

$$\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{r} = \boldsymbol{z}_{\mathrm{loc}} + \boldsymbol{E}_0 \boldsymbol{A}_0^{-1} \boldsymbol{E}_0^{*} \left(\boldsymbol{r} - \boldsymbol{A}\boldsymbol{z}_{\mathrm{loc}}\right). \tag{3.40}$$

Thus, the coarse correction is applied to the residual remaining after the local Schwarz correction.

In the reported experiments, the local patch problems and the coarse problem are solved by sparse LU factorization.[2] For one-level comparison runs, the same local correction is used but the coarse basis is disabled:

$$\boldsymbol{\mathcal{M}}_{\mathrm{loc}}^{-1}\boldsymbol{r} = \boldsymbol{z}_{\mathrm{loc}}. \tag{3.41}$$

3.6 Offline and online solution procedure

Each preconditioned solve is separated into an offline stage and an online stage. The offline stage constructs and stores the data needed by the preconditioner. The online stage loads this data, assembles the global system for the chosen case, and applies GMRES with the preconditioner from Section 3.5.

The offline stage depends on the mesh, frequency, boundary markers, domain decomposition parameters, and spectral coarse-space parameters. It constructs the finite element matrix, builds the domain decomposition and partition-of-unity

[2] LU factorization writes a matrix as a product of a lower triangular matrix and an upper triangular matrix. Once this factorization has been computed, linear systems with the same matrix can be solved efficiently. A sparse LU solver is used here because each GMRES iteration requires repeated local solves and one coarse solve.

weights, assembles the local patch operators, solves the local spectral problems, and forms the coarse basis

$$\boldsymbol{E}_0 \in \mathbb{C}^{N \times n_c}. \tag{3.42}$$

The stored offline data include the decomposition, the partition-of-unity weights, the local patch matrices, the retained local spectral modes and eigenvalues, and the explicit coarse basis matrix.

The online stage starts by loading the stored data and checking that the number of finite element degrees of freedom agrees with the current mesh. The global finite element system is then assembled:

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b}. \tag{3.43}$$

The two-level preconditioner is rebuilt as an operator from the stored local patch data and the coarse basis.

GMRES is applied to the left-preconditioned system

$$\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{A}\boldsymbol{u} = \boldsymbol{\mathcal{M}}^{-1}\boldsymbol{b}. \tag{3.44}$$

The matrix $\boldsymbol{\mathcal{M}}^{-1}\boldsymbol{A}$ is not assembled explicitly. Each GMRES matrix-vector product is applied as

$$\boldsymbol{x} \mapsto \boldsymbol{\mathcal{M}}^{-1}\left(\boldsymbol{A}\boldsymbol{x}\right). \tag{3.45}$$

The right-hand side is also preconditioned:

$$\boldsymbol{b}_{\mathrm{prec}} = \boldsymbol{\mathcal{M}}^{-1}\boldsymbol{b}. \tag{3.46}$$

The relative tolerance used for the reported GMRES solves is 10^{-6}. After GMRES terminates, the solution is checked using the true, unpreconditioned relative residual defined in Section 3.7.

3.7 Post-processing and comparison measures

The numerical results are compared using both solver diagnostics and solution diagnostics. The main solver diagnostics are GMRES iteration count, final true relative residual, number of finite element degrees of freedom, and coarse-space dimension.

The true relative residual is computed after the solve using the original, unpreconditioned system:

$$\eta_{\mathrm{true}} = \frac{\|\boldsymbol{b} - \boldsymbol{A}\boldsymbol{u}\|_2}{\|\boldsymbol{b}\|_2}. \tag{3.47}$$

This is the residual reported in the result tables.

The size of the computed acoustic field is summarized using the coefficient RMS value

$$u_{\mathrm{rms}} = \left(\frac{1}{N}\sum_{j=1}^{N} |u_j|^2\right)^{1/2}, \tag{3.48}$$

and the maximum coefficient amplitude

$$u_{\max} = \max_{1 \le j \le N} |u_j|. \tag{3.49}$$

These are numerical measures of the computed complex pressure amplitude; they are not calibrated physical sound pressure levels.

For comparisons between two solutions on the same finite element space, the relative coefficient difference is

$$e_{\mathrm{rel}} = \frac{\|\boldsymbol{u}^{(1)} - \boldsymbol{u}^{(2)}\|_2}{\|\boldsymbol{u}^{(2)}\|_2}. \tag{3.50}$$

This is used only when the two coefficient vectors are defined on the same mesh.

For mesh-refinement studies, coefficient vectors cannot be compared directly, because different meshes have different finite element spaces. Instead, the solution is evaluated at fixed receiver points in the physical domain. If $\boldsymbol{u}_{\mathrm{rec}}^{(h)}$ is the vector of complex solution values at receiver points for mesh size h, and $\boldsymbol{u}_{\mathrm{rec}}^{(\mathrm{ref})}$ is the corresponding vector on the reference mesh, the receiver error is

$$e_{\mathrm{rec}}^{(h)} = \frac{\left\|\boldsymbol{u}_{\mathrm{rec}}^{(h)} - \boldsymbol{u}_{\mathrm{rec}}^{(\mathrm{ref})}\right\|_2}{\left\|\boldsymbol{u}_{\mathrm{rec}}^{(\mathrm{ref})}\right\|_2}. \tag{3.51}$$

Receiver errors are used as mesh-refinement diagnostics, not as validation against measured sound levels.

For the exact plane-wave verification problem, the numerical solution is compared with the known exact solution. The relative L^2 error is

$$e_{L^2} = \frac{\|u_{\mathrm{ex}} - u_h\|_{L^2(\Omega)}}{\|u_{\mathrm{ex}}\|_{L^2(\Omega)}}, \tag{3.52}$$

and the relative H^1-seminorm error is

$$e_{H^1} = \frac{|u_{\mathrm{ex}} - u_h|_{H^1(\Omega)}}{|u_{\mathrm{ex}}|_{H^1(\Omega)}}. \tag{3.53}$$

These errors are used only for the verification problem with a known exact solution.

4

Numerical experiments and results

This chapter presents the numerical experiments used to evaluate the finite element Helmholtz solver and the two-level preconditioner. The experiments are organized in increasing order of complexity: first an exact plane-wave verification problem, then controlled box-domain studies, and finally realistic city-domain simulations.

The purpose is to separate three aspects of the numerical method: whether the finite element implementation gives the expected convergence behaviour, how the two-level preconditioner affects GMRES convergence, and what practical limitations appear on large three-dimensional urban meshes. The diagnostic quantities used in the experiments are defined in Section 3.7.

4.1 Experiment design

The experiment design is guided by the research questions stated in Section 1.2. The tests are organized in three stages, shown in Figure 4.1, so that the numerical method is first checked on a problem with a known exact solution and then applied to increasingly realistic geometries.

The first stage is an exact plane-wave verification problem on the unit cube. Since an analytical solution is known, this test verifies that the finite element implementation gives the expected convergence rates under mesh refinement.

The second stage uses a rectangular box domain. This provides a controlled three-dimensional acoustic geometry for studying mesh refinement, frequency dependence, and preconditioner parameters. The box experiments are also used to compare the two-level method with simpler solver variants.

The third stage uses realistic city-domain meshes representing part of Gothenburg. These experiments test whether the same solver pipeline remains practical on geometrically complex urban domains and expose the main memory limitations of the current implementation.

4.1.1 Frequencies and wavelength resolution

The experiments use frequencies between 5 Hz and 30 Hz. With sound speed $c = 343\,\mathrm{m\,s^{-1}}$, the wavelength is

$$\lambda = \frac{c}{f}. \tag{4.1}$$

The corresponding wavelengths are shown in Table 4.1. As the frequency increases, the wavelength decreases, so the same mesh size gives fewer degrees of freedom

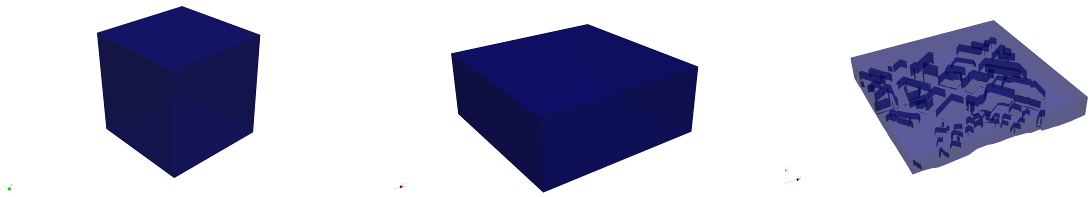

(a) Unit cube. (b) Box domain. (c) City domain.

Figure 4.1: Geometries used in the numerical experiments: (a) the unit cube used for exact plane-wave verification, (b) the rectangular box domain used for controlled parameter studies, and (c) the realistic city-domain mesh representing part of Gothenburg.

per wavelength. Higher-frequency cases are therefore more demanding both for the acoustic discretization and for the iterative solver.

Table 4.1: Frequencies used in the experiments and corresponding wavelengths for $c = 343\,\mathrm{m\,s^{-1}}$.

Frequency	Wavelength
5 Hz	68.6 m
10 Hz	34.3 m
20 Hz	17.15 m
30 Hz	11.43 m

4.1.2 Domains used in the experiments

Three types of computational domains are used. The first is the unit cube, which is used only for verification. The verification problem has a known exact plane-wave solution and is used to check that the finite element discretization gives the expected convergence rates.

The second is a rectangular box domain with dimensions approximately $200\,\mathrm{m} \times 200\,\mathrm{m} \times 80\,\mathrm{m}$. This domain is used as a controlled three-dimensional acoustic geometry. It is simple enough for systematic parameter studies, but large enough for the Helmholtz systems to show the difficult behaviour expected in acoustic simulations.

The third is a realistic city-domain mesh representing part of Gothenburg. This geometry contains ground, building surfaces, and open-air regions. The city domain is used to test whether the two-level preconditioner remains practical on realistic three-dimensional urban acoustic geometries.

The box and city domains are both represented by tetrahedral volume meshes. The city meshes were generated from urban geometry data using the DTCC platform and the `dtcc-sim` software package [5]. In all reported experiments, the acoustic medium is treated as homogeneous, and the source is represented by the Gaussian forcing term described in Section 3.2.

The box domain is therefore used mainly to understand the numerical method in a controlled setting, while the city domain is used to test the method in the intended application setting.

4.1.3 Solver comparisons

The main method studied in the thesis is the two-level preconditioned GMRES solver described in Chapter 3. In the box-domain experiments, it is compared with a one-level method where the global coarse correction is removed. This isolates the effect of the spectral coarse space.

Additional plain GMRES runs are used as a baseline. These solves use the same finite element systems but no preconditioner. They are included to show how many iterations are required when the Helmholtz system is solved without the two-level preconditioning strategy.

4.1.4 Main output quantities

The main output quantities are those defined in Section 3.7. Solver performance is measured using GMRES iteration counts, true relative residuals, coarse-space dimensions, and wall-clock time. Here, wall-clock time means the elapsed real time for the run, including the relevant setup and solve stages reported in the terminal output.

The computed acoustic field is assessed using RMS values, maximum amplitudes, and receiver comparisons. These quantities are not used as validation against measured sound levels. Instead, they are used as internal numerical diagnostics for studying how the computed solution changes under mesh refinement or changes in solver parameters.

4.1.5 Computational setup

The computations were run on a Linux workstation with an Intel Xeon Silver 4110 CPU at 2.10 GHz, using 16 logical CPU threads and 376 GiB of RAM. The operating system was Ubuntu Linux with kernel version 5.15.

The Python environment used Python 3.11.15, NumPy 1.26.4 [8], SciPy 1.17.1 [17], DOLFINx 0.9.0 [4], and petsc4py 3.21.1/PETSc [1]. The implementation code for this thesis was developed by the author. External software packages are used under their respective licenses.

4.2 Verification on an exact plane wave

Before studying the box and city geometries, the finite element implementation is checked on a problem with a known exact solution. This is a verification test: the purpose is to confirm that the implementation gives the expected finite element convergence rates when the mesh is refined.

The test is performed on the unit cube,

$$\Omega = (0,1)^3. \tag{4.2}$$

The exact solution is chosen as a plane wave,

$$u_{\mathrm{ex}}(\boldsymbol{x}) = \exp(\mathrm{i}k\boldsymbol{d}\cdot\boldsymbol{x})\,, \tag{4.3}$$

where $\boldsymbol{d} \in \mathbb{R}^3$ is a unit propagation direction. Since

$$-\Delta u_{\mathrm{ex}} = k^2 u_{\mathrm{ex}}, \tag{4.4}$$

the exact solution satisfies

$$-\Delta u_{\mathrm{ex}} - k^2 u_{\mathrm{ex}} = 0 \qquad \text{in } \Omega. \tag{4.5}$$

For this verification problem, the exact plane-wave field is enforced through matching Robin boundary data on the whole boundary of the unit cube. More precisely, the problem solved is

$$-\Delta u - k^2 u = 0 \qquad \text{in } \Omega, \tag{4.6}$$

with boundary condition

$$\frac{\partial u}{\partial n} - iku = \frac{\partial u_{\mathrm{ex}}}{\partial n} - iku_{\mathrm{ex}} \qquad \text{on } \partial\Omega, \tag{4.7}$$

so that the plane wave u_{ex} is the exact solution of the full boundary value problem. The numerical solution is then compared directly with this analytical reference under mesh refinement.

The test is performed at 5 Hz. The mesh is refined, and the numerical solution u_h is compared with u_{ex}. The relative L^2 error is computed as

$$e_{L^2} = \frac{\|u_{\mathrm{ex}} - u_h\|_{L^2(\Omega)}}{\|u_{\mathrm{ex}}\|_{L^2(\Omega)}}, \tag{4.8}$$

where

$$\|u_{\mathrm{ex}} - u_h\|_{L^2(\Omega)} = \left(\int_\Omega |u_{\mathrm{ex}} - u_h|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}. \tag{4.9}$$

The relative H^1-seminorm error is computed as

$$e_{H^1} = \frac{|u_{\mathrm{ex}} - u_h|_{H^1(\Omega)}}{|u_{\mathrm{ex}}|_{H^1(\Omega)}}, \tag{4.10}$$

with

$$|u_{\mathrm{ex}} - u_h|_{H^1(\Omega)} = \left(\int_\Omega |\nabla u_{\mathrm{ex}} - \nabla u_h|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}. \tag{4.11}$$

For first-order Lagrange finite elements, one expects approximately second-order convergence in the L^2 norm and first-order convergence in the H^1-seminorm, provided that the exact solution is smooth and the mesh is refined regularly [2]. The numerical results are shown in Table 4.2, and the corresponding convergence plot is shown in Figure 4.2.

The results in Table 4.2 and Figure 4.2 show the expected convergence behaviour. The relative L^2 error decreases faster than the relative H^1-seminorm error, which

Table 4.2: Exact plane-wave verification test at 5 Hz on the unit cube.

h	N	Relative L^2 error	Relative H^1-seminorm error
0.1250	729	8.78×10^{-6}	3.27×10^{-3}
0.0833	2197	3.96×10^{-6}	2.19×10^{-3}
0.0625	4913	2.24×10^{-6}	1.65×10^{-3}
0.0417	15625	9.99×10^{-7}	1.10×10^{-3}
0.0312	35937	5.63×10^{-7}	8.25×10^{-4}

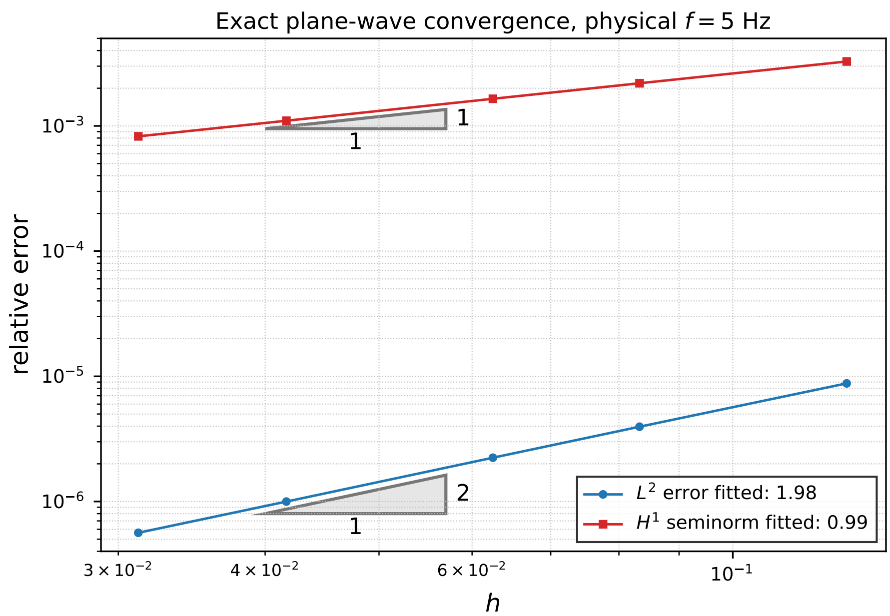

Figure 4.2: Convergence of the exact plane-wave verification problem at 5 Hz. The L^2 error decreases approximately quadratically, while the H^1-seminorm error decreases approximately linearly, as expected for first-order Lagrange finite elements.

is consistent with the theory for first-order finite elements. This confirms that the finite element implementation behaves correctly on a problem with a known exact solution.

This verification step is important before moving to the box and city geometries. In those later experiments, no exact solution is available, so the computed acoustic field is assessed through mesh refinement, receiver comparisons, and solver residuals rather than by direct comparison with an analytical solution.

4.3 Box-domain experiments

After the exact plane-wave verification, the solver is tested on a rectangular box domain. The box is used as a controlled intermediate geometry between the unit cube and the realistic city meshes. It uses the same type of Gaussian source and absorbing-boundary treatment as the city-domain experiments, but the geometry is simple enough that parameter studies are easier to interpret. The purpose of the box-domain experiments is to study mesh refinement, frequency dependence, and the influence of the main preconditioner parameters. The section also compares the two-level method with simpler solver variants, including a one-level method and plain GMRES without preconditioning.

4.3.1 Box-domain setup

The box domain has dimensions

$$200\,\mathrm{m} \times 200\,\mathrm{m} \times 80\,\mathrm{m}. \tag{4.12}$$

The Gaussian source is placed near the lower part of the domain, at

$$\boldsymbol{x}_0 = (100,\ 100,\ 8)\,\mathrm{m}. \tag{4.13}$$

The source amplitude, source width, material coefficient, and boundary treatment are the same as described in Section 3.2.

The box meshes are tetrahedral volume meshes. The mesh family used in the box experiments is summarized in Table 4.3.

Table 4.3: Box-domain meshes used in the experiments.

Mesh size h [m]	Degrees of freedom N
5.0	28577
4.0	54621
3.0	129472
2.5	216513
2.0	418241

The frequencies considered on the box domain are 5 Hz, 10 Hz, 20 Hz, and 30 Hz. Unless otherwise stated, the threshold-based two-level preconditioner uses

$$M = 16, \qquad \text{one overlap layer}, \qquad \text{three oversampling layers}, \qquad \rho = 1.10. \tag{4.14}$$

Later fixed-mode experiments prescribe the number of retained local modes directly instead of relying only on the threshold parameter ρ.

4.3.2 Mesh-refinement studies at 5, 10, and 20 Hz

The first box-domain study considers mesh refinement at 5 Hz, 10 Hz, and 20 Hz. The baseline preconditioner parameters in (4.14) are used. The goal is to test

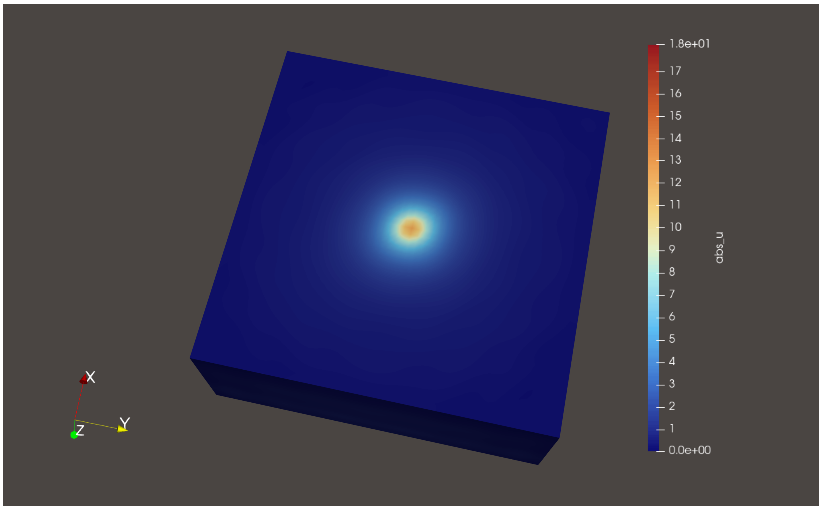

Figure 4.3: Example box-domain solution at 10 Hz on the $h = 4$ m mesh, showing the magnitude of the computed complex pressure amplitude. The largest amplitude occurs near the Gaussian source and the field decays away from it across the rectangular domain. The view is from below the domain, as indicated by the coordinate axes, so the visible face is the lower side of the box.

Table 4.4: Box-domain mesh-refinement study at 5 Hz.

h	N	n_c	Mean n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$
5.0	28577	32	2.00	27	2.85×10^{-6}	1.6402	20.7685
4.0	54621	35	2.19	32	2.60×10^{-6}	1.6571	21.6871
3.0	129472	30	1.88	36	5.85×10^{-6}	1.6730	21.0815
2.5	216513	30	1.88	40	8.93×10^{-6}	1.6799	21.9844

Table 4.5: Box-domain mesh-refinement study at 10 Hz.

h	N	n_c	Mean n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$
4.0	54621	218	13.62	30	3.50×10^{-6}	1.3214	18.4780
3.0	129472	238	14.88	32	5.06×10^{-6}	1.3605	18.2762
2.5	216513	247	15.44	35	5.88×10^{-6}	1.3772	18.9800
2.0	418241	247	15.44	40	6.98×10^{-6}	1.3933	19.3119

Table 4.6: Box-domain mesh-refinement study at 20 Hz.

h	N	n_c	Mean n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$
4.0	54621	915	57.19	35	1.78×10^{-6}	0.5095	11.4818
3.0	129472	1114	69.62	35	2.17×10^{-6}	0.5157	11.2833
2.5	216513	1219	76.19	35	3.59×10^{-6}	0.5186	12.0892
2.0	418241	1326	82.88	37	3.67×10^{-6}	0.5210	12.4252

whether the two-level preconditioner gives stable GMRES convergence as the mesh is refined and as the frequency increases.

The iteration counts remain moderate for all three frequencies. At 5 Hz, the count increases from 27 to 40 as the mesh is refined. At 10 Hz, it increases from 30 to 40. At 20 Hz, the iteration count remains between 35 and 37, but the coarse-space dimension becomes much larger. For example, on the $h = 4$ mesh, n_c increases from 35 at 5 Hz to 218 at 10 Hz and 915 at 20 Hz. This shows that the local spectral construction retains many more modes as the Helmholtz problem becomes more oscillatory.

The RMS values change smoothly under refinement at 5 Hz. At higher frequency, especially 20 Hz, global RMS values alone are not enough to judge acoustic convergence. Receiver-level comparisons are therefore considered next.

4.3.3 Receiver convergence on the box domain

Receiver convergence is measured using the relative receiver error defined in Section 3.7. For each frequency, the solution is evaluated at a fixed set of receiver points and compared with the finest available mesh used as reference. For 5 Hz, the $h = 2.5$ solution is used as reference. For 10 Hz and 20 Hz, the $h = 2$ solution is used as reference.

Table 4.7: Receiver convergence for the box-domain mesh-refinement studies.

Frequency [Hz]	h	Reference mesh	Relative receiver L^2 error
5	5.0	2.5	4.93%
5	4.0	2.5	3.12%
5	3.0	2.5	1.38%
10	4.0	2.0	23.94%
10	3.0	2.0	10.84%
10	2.5	2.0	4.89%
20	4.0	2.0	108.60%
20	3.0	2.0	63.98%
20	2.5	2.0	35.79%

The receiver errors decrease under mesh refinement for all three frequencies. The 5 Hz case gives the clearest receiver convergence, decreasing to 1.38% on the $h = 3$ mesh relative to the $h = 2.5$ reference. The 10 Hz case also improves clearly under refinement, but the errors are larger. At 20 Hz, the receiver error remains high even on the $h = 2.5$ mesh. Thus, the 20 Hz box results should be interpreted as a solver and mesh-sensitivity study rather than as fully receiver-converged acoustic predictions.

4.3.4 Effect of the spectral threshold

The spectral threshold ρ controls how many local eigenmodes are retained in the coarse space. The experiment is performed at 10 Hz on the $h = 2.5$ box mesh, with

$$M = 16, \qquad \text{one overlap layer}, \qquad \text{three oversampling layers.} \tag{4.15}$$

Only ρ is varied.

Table 4.8: Effect of the spectral threshold for the box domain at 10 Hz on the $h = 2.5$ mesh.

ρ	N	n_c	Mean n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$
0.80	216513	476	29.75	22	7.68×10^{-6}	1.3772	18.9800
1.00	216513	312	19.50	32	5.80×10^{-6}	1.3772	18.9800
1.10	216513	247	15.44	35	5.88×10^{-6}	1.3772	18.9800
1.25	216513	152	9.50	41	4.88×10^{-6}	1.3772	18.9800

A smaller threshold keeps more local modes and gives a stronger preconditioner. With $\rho = 0.80$, the coarse space has dimension $n_c = 476$, and GMRES converges in 22 iterations. With $\rho = 1.25$, the coarse space is reduced to $n_c = 152$, and the iteration count increases to 41. The RMS value and maximum amplitude are unchanged because all runs solve the same discrete problem; only the preconditioner changes.

4.3.5 Effect of oversampling

The next parameter study varies the number of oversampling layers used in the local spectral problems. The experiment is performed at 10 Hz on the $h = 2.5$ box mesh, with

$$M = 16, \qquad \text{one overlap layer}, \qquad \rho = 0.80. \tag{4.16}$$

The number of oversampling layers is varied from 1 to 4.

Table 4.9: Effect of oversampling for the box domain at 10 Hz on the $h = 2.5$ mesh.

Oversampling	N	n_c	Mean n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$
1	216513	599	37.44	29	9.09×10^{-6}	1.3772	18.9800
2	216513	430	26.88	29	7.47×10^{-6}	1.3772	18.9800
3	216513	476	29.75	22	7.68×10^{-6}	1.3772	18.9800
4	216513	455	28.44	22	5.83×10^{-6}	1.3772	18.9800

With one or two oversampling layers, GMRES requires 29 iterations. With three or four layers, the iteration count decreases to 22. The improvement is not explained only by the coarse-space dimension: the one-layer case has the largest coarse space in the table, but does not give the best convergence. This shows that the quality of the local spectral modes depends on the oversampled domains on which they are computed, not only on the number of retained modes.

4.3.6 One-level and two-level comparison

To isolate the effect of the spectral coarse correction, the two-level preconditioner is compared with a one-level variant. The one-level method uses the local restricted additive Schwarz correction but omits the global coarse correction.

Table 4.10: Comparison between the one-level and two-level methods for the box domain at 5 Hz.

Frequency [Hz]	h	Method	N	n_c	GMRES	Residual
5	5.0	one-level	28577	–	32	2.16×10^{-6}
5	5.0	two-level	28577	32	27	2.85×10^{-6}
5	4.0	one-level	54621	–	37	3.85×10^{-6}
5	4.0	two-level	54621	35	32	2.60×10^{-6}
5	3.0	one-level	129472	–	43	4.88×10^{-6}
5	3.0	two-level	129472	30	36	5.85×10^{-6}

At 5 Hz, the two-level method gives a moderate but consistent reduction in GMRES iterations. A stronger effect is observed at 10 Hz, as shown in Table 4.11.

Table 4.11: Comparison between the one-level and two-level methods for the box domain at 10 Hz on the $h = 2.5$ mesh.

Frequency [Hz]	h	Method	N	n_c	GMRES	Residual
10	2.5	one-level	216513	–	56	4.41×10^{-6}
10	2.5	two-level	216513	476	22	7.68×10^{-6}

At 10 Hz, the two-level method reduces the iteration count from 56 to 22. This confirms that the spectral coarse correction is important for reducing global wave-like error components that are not efficiently handled by local subdomain solves alone.

4.3.7 Plain GMRES baseline

As a naive algebraic baseline, selected box-domain systems were solved using plain GMRES without preconditioning. In these runs, GMRES is applied directly to

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b}. \tag{4.17}$$

The same finite element matrices and right-hand sides are used as in the preconditioned runs.

Plain GMRES converges on this controlled box mesh, but it requires several hundred iterations. For example, at 10 Hz on the same $h = 2.5$ mesh, plain GMRES requires 694 iterations, compared with 56 iterations for the one-level method and 22 iterations for the two-level method. This shows that the main improvement comes from the preconditioner, not from GMRES alone.

Table 4.12: Plain GMRES baseline without preconditioning on the box domain with $h = 2.5$.

Frequency [Hz]	h [m]	N	GMRES iterations	Residual
5	2.5	216513	659	9.88×10^{-7}
10	2.5	216513	694	9.93×10^{-7}
20	2.5	216513	618	9.96×10^{-7}
30	2.5	216513	541	9.83×10^{-7}

4.3.8 Fixed local mode count and number of subdomains

The threshold-based experiments show that the number of selected local modes can change when the frequency, mesh size, or subdomain size changes. To study the effect of the number of subdomains more directly, a fixed-mode strategy is used. The number of retained local modes is prescribed as

$$n_i = 100, \tag{4.18}$$

so that

$$n_c = Mn_i. \tag{4.19}$$

The experiment is performed at 20 Hz on the $h = 2.5$ box mesh, with one overlap layer and three oversampling layers.

Table 4.13: Effect of the number of subdomains for the box domain at 20 Hz and $h = 2.5$, using $n_i = 100$ retained local modes per subdomain.

M	N	n_c	n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \lvert u_h \rvert$	Time [s]
8	216513	800	100	27	1.99×10^{-6}	0.5186	12.0892	13524
16	216513	1600	100	31	2.91×10^{-6}	0.5186	12.0892	8121
32	216513	3200	100	19	4.34×10^{-6}	0.5186	12.0892	7690
64	216513	6400	100	12	2.27×10^{-6}	0.5186	12.0892	8001

With fixed n_i, increasing M increases the total coarse-space dimension. The iteration count decreases from 27 for $M = 8$ to 12 for $M = 64$, although the wall-clock time does not decrease monotonically. The $M = 64$ case has the lowest iteration count, but it also has the largest coarse space, so each preconditioner application is more expensive.

4.3.9 Effect of oversampling with fixed local mode count

The effect of oversampling is also tested with fixed local mode count. The experiment is performed at 20 Hz on the $h = 2.5$ box mesh, with

$$M = 64, \qquad n_i = 100, \qquad n_c = 6400. \tag{4.20}$$

Only the number of oversampling layers is varied.

Table 4.14: Effect of oversampling for the box domain at 20 Hz, $h = 2.5$, $M = 64$, and fixed $n_i = 100$.

Oversampling	N	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$	Time [s]
0	216513	67	6.33×10^{-6}	0.5186	12.0892	5875
1	216513	35	5.77×10^{-6}	0.5186	12.0892	6573
2	216513	23	3.90×10^{-6}	0.5186	12.0892	6268
3	216513	12	2.27×10^{-6}	0.5186	12.0892	8001

Since M, n_i, and n_c are fixed, the reduction in iteration count is not caused by a larger coarse space. Instead, it shows that local modes computed on larger oversampled regions give a more effective coarse basis. The three-layer case gives the lowest iteration count, while the two-layer case gives a useful compromise between iteration count and wall-clock time.

4.3.10 Fixed-mode study at 30 Hz

The final box-domain experiment increases the frequency to 30 Hz. The setup is fixed across the mesh sequence:

$$M = 64, \qquad \text{one overlap layer}, \qquad \text{two oversampling layers}, \qquad n_i = 100. \tag{4.21}$$

Thus,

$$n_c = M n_i = 6400. \tag{4.22}$$

Table 4.15: Box-domain mesh-refinement study at 30 Hz using a fixed-mode coarse space with $M = 64$, two oversampling layers, and $n_i = 100$.

h	n_i	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max \|u_h\|$	Δ_{rms}	Time [s]
4.0	54621	31	7.32×10^{-7}	0.1427	6.0831	–	1119
3.0	129472	43	1.03×10^{-6}	0.1248	5.5936	12.55%	2797
2.5	216513	52	1.21×10^{-6}	0.1166	6.1516	6.53%	6080

All three 30 Hz box-domain runs completed successfully. The iteration count increases from 31 on the $h = 4$ mesh to 52 on the $h = 2.5$ mesh. The RMS value also changes under refinement, although the relative change decreases from 12.55% to 6.53%. Since the wavelength at 30 Hz is only about 11.4 m, these results should be interpreted as a high-frequency feasibility and mesh-sensitivity study rather than as a fully mesh-converged acoustic prediction.

4.4 City-domain experiments

After the box-domain experiments, the method is tested on a realistic city geometry. The city-domain experiments are numerical solver tests on a more complex three-dimensional geometry. They are not intended as validation against measured sound levels. The purpose is to test whether the two-level preconditioner remains practical when the domain contains ground, buildings, and a strongly unstructured tetrahedral mesh.

The city experiments have three main purposes. First, they test whether the two-level preconditioner gives stable GMRES convergence on a real urban geometry. Second, they show how the behaviour changes as the frequency is increased from 5 Hz to 10 Hz, 20 Hz, and 30 Hz. Third, they identify the practical memory limitations caused by the explicit coarse basis on large three-dimensional city problems.

4.4.1 City-domain setup

The city-domain mesh represents a part of Gothenburg. The computational domain has approximate horizontal size $500\,\mathrm{m} \times 500\,\mathrm{m}$ and height $80\,\mathrm{m}$. The source is placed near the centre of the domain and close to the ground. As in the box-domain experiments, the coefficient is kept constant and the right-hand side is a localized Gaussian source.

The nominal city mesh sizes used in the experiments are

$$h \in \{5, 4, 3\}\,\mathrm{m}. \tag{4.23}$$

The corresponding finite element dimensions are shown in Table 4.16.

Table 4.16: City-domain meshes used in the experiments.

Mesh size h [m]	Degrees of freedom N
5.0	330073
4.0	634972
3.0	1466914

The city geometry is more demanding than the box geometry for two reasons. First, the mesh is geometrically more complicated, because the subdomains may contain buildings, ground surfaces, and open-air regions. Second, the number of degrees of freedom is much larger for the same nominal mesh size. The city experiments therefore test both the convergence behaviour of the preconditioner and the practical memory limits of the implementation.

4.4.2 Mesh-refinement studies at 5 and 10 Hz

The first city-domain studies are mesh-refinement tests at 5 Hz and 10 Hz. The corresponding wavelengths are

$$\lambda_5 = \frac{343}{5} \approx 68.6\,\mathrm{m}, \qquad \lambda_{10} = \frac{343}{10} \approx 34.3\,\mathrm{m}. \tag{4.24}$$

Thus, the same city meshes are significantly better resolved at 5 Hz than at 10 Hz.

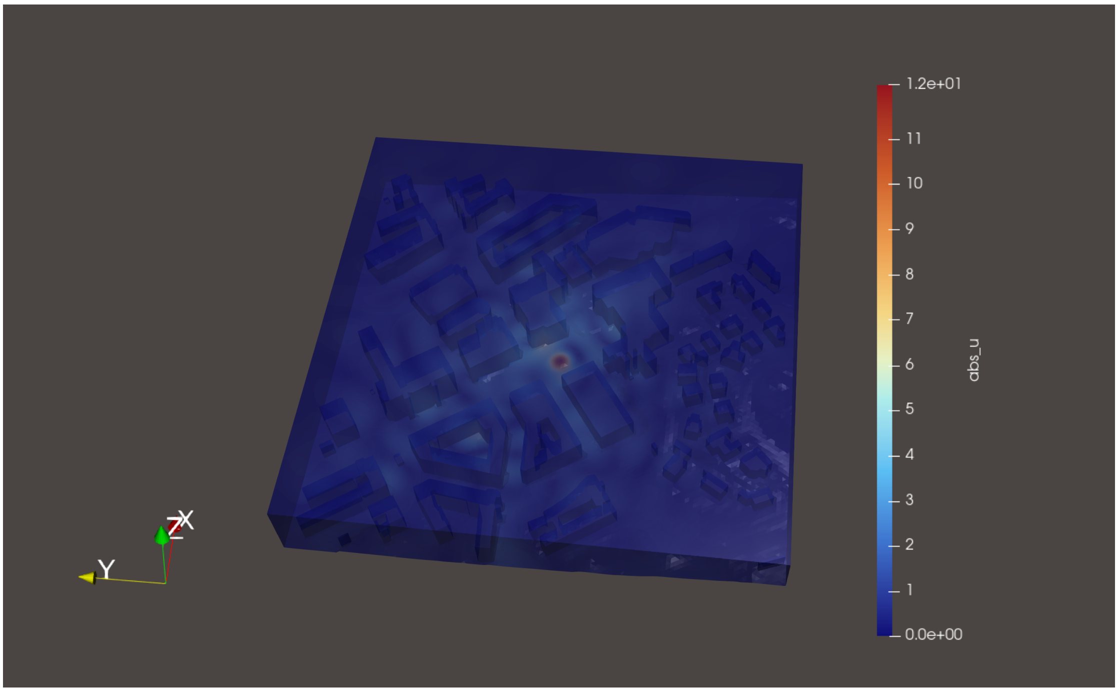

(a) View from one side of the city domain.

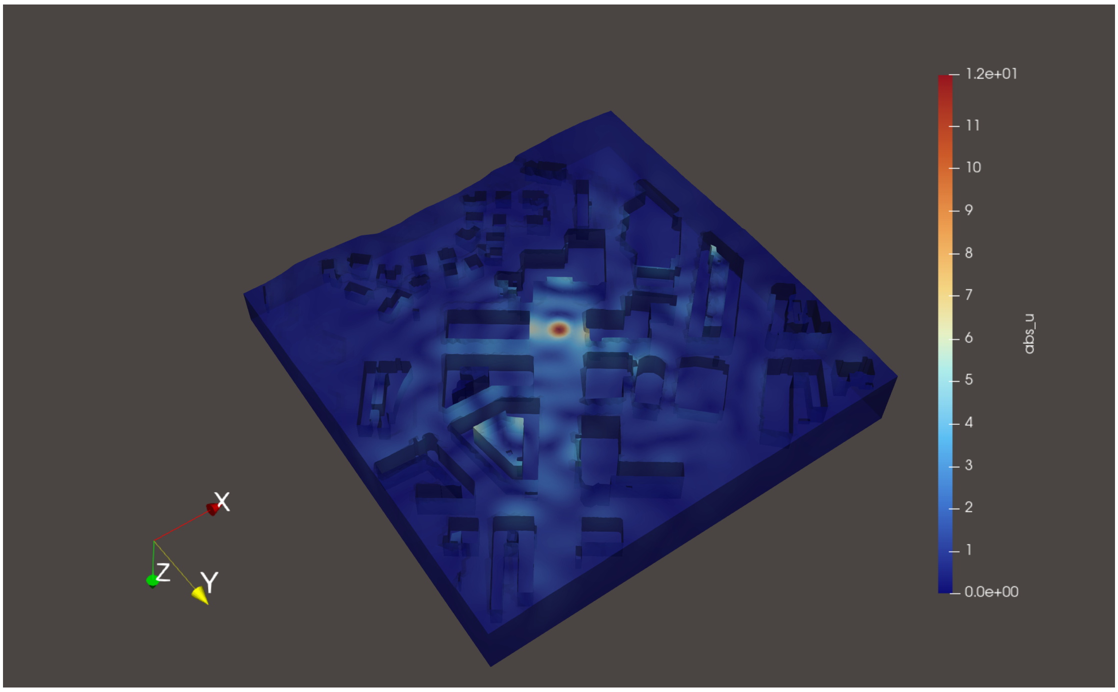

(b) View from the opposite side of the city domain.

Figure 4.4: Example city-domain solution at 5 Hz on the $h = 5\,\text{m}$ mesh, showing the magnitude of the computed complex pressure amplitude from two viewing directions. The transparent volume view makes the building geometry visible inside the computational domain. The largest amplitude is located near the Gaussian source, while the surrounding buildings and ground shape the spatial distribution of the field.

For the 5 Hz study, the preconditioner parameters are kept fixed across the mesh sequence:

$$M = 16, \qquad \text{overlap} = 1, \qquad \text{oversampling} = 3, \qquad \rho = 1.10. \tag{4.25}$$

The results are shown in Table 4.17.

Table 4.17: City-domain mesh-refinement study at 5 Hz with $\rho = 1.10$, $M = 16$, one overlap layer, and three oversampling layers.

h	N	n_c	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max\|u_h\|$	Δ_{rms}
5.0	330073	281	32	8.00×10^{-6}	0.6972	12.3144	–
4.0	634972	338	33	1.81×10^{-5}	0.6926	12.2302	0.67%
3.0	1466914	545	36	3.23×10^{-5}	0.6794	12.2239	1.90%

At 5 Hz, the GMRES iteration count remains stable under mesh refinement, increasing only from 32 iterations on the $h = 5$ m mesh to 36 iterations on the $h = 3$ m mesh. The RMS value changes by 0.67% from $h = 5$ m to $h = 4$ m, and by 1.90% from $h = 4$ m to $h = 3$ m. This makes the 5 Hz city case the clearest city-domain mesh-refinement result obtained in this work.

The 10 Hz study is more demanding because the wavelength is half as large. The first two 10 Hz runs use the threshold-based setup

$$M = 16, \qquad \text{overlap} = 1, \qquad \text{oversampling} = 3, \qquad \rho = 1.10. \tag{4.26}$$

The $h = 3$ m run was instead performed with a fixed-mode coarse space:

$$M = 64, \qquad \text{overlap} = 1, \qquad \text{oversampling} = 2, \qquad n_i = 50. \tag{4.27}$$

This gives

$$n_c = M n_i = 64 \cdot 50 = 3200. \tag{4.28}$$

The completed 10 Hz results are shown in Table 4.18. The $h = 5$ m and $h = 4$ m rows use the threshold-based setup in (4.26), while the $h = 3$ m row uses the fixed-mode setup in (4.27).

Table 4.18: Completed city-domain mesh-refinement study at 10 Hz. The $h = 5$ m and $h = 4$ m runs use the threshold-based setup with $M = 16$, while the $h = 3$ m run uses the fixed-mode setup with $M = 64$ and $n_i = 50$ retained modes per subdomain.

h	N	n_c	GMRES	Residual	$\|u_h\|_{\text{rms}}$	$\max\|u_h\|$	Δ_{rms}
5.0	330073	696	30	9.24×10^{-6}	0.5447	10.6307	–
4.0	634972	732	33	1.04×10^{-5}	0.5329	10.7529	2.16%
3.0	1466914	3200	55	2.68×10^{-5}	0.5437	11.0465	2.02%

Compared with the 5 Hz city study, the coarse spaces are much larger at 10 Hz. For example, on the $h = 5$ m mesh, the coarse-space dimension increases from $n_c = 281$ at 5 Hz to $n_c = 696$ at 10 Hz. This is expected, because the Helmholtz problem becomes more oscillatory as the frequency increases. The $h = 3$ m solve completes with the fixed-mode coarse space, but the GMRES iteration count increases to 55.

4.4.3 Receiver convergence on the city domain

The city solutions were also compared at the receiver points. The receiver error is computed using the relative receiver error defined in Section 3.7. For both frequencies, the $h = 3\,\mathrm{m}$ solution is used as the reference solution.

The 5 Hz receiver results are shown in Table 4.19.

Table 4.19: Receiver convergence for the city-domain 5 Hz study using the $h = 3\,\mathrm{m}$ solution as reference.

Mesh	h	Relative receiver L^2 error
$h5$	5.0	9.14%
$h4$	4.0	5.11%

The receiver error decreases from 9.14% on the $h = 5\,\mathrm{m}$ mesh to 5.11% on the $h = 4\,\mathrm{m}$ mesh. This confirms that the local receiver values improve under mesh refinement at 5 Hz.

The corresponding 10 Hz receiver results are shown in Table 4.20.

Table 4.20: Receiver convergence for the city-domain 10 Hz study using the $h = 3\,\mathrm{m}$ solution as reference.

Mesh	h	Relative receiver L^2 error
$h5$	5.0	36.34%
$h4$	4.0	19.19%

The 10 Hz receiver errors also decrease under mesh refinement, but they are much larger than in the 5 Hz study. This indicates that the receiver values are not yet fully mesh-converged at 10 Hz on the coarser city meshes. The $h = 3\,\mathrm{m}$ solution provides a useful reference, but a still finer mesh would be needed for a stronger receiver-level convergence conclusion.

4.4.4 Effect of retained local modes at 10 Hz

The completed $h = 3\,\mathrm{m}$, 10 Hz city run used a fixed-mode coarse space with $M = 64$, one overlap layer, and two oversampling layers. This was motivated by the memory cost of the threshold-based coarse space. When the frequency and mesh resolution increase, the local spectral problems may select many basis functions, and the global coarse basis can become very large. The fixed-mode strategy gives direct control of the coarse-space dimension:

$$n_c = M n_i. \tag{4.29}$$

To study this tradeoff, the same offline data from the $n_i = 100$ computation was reduced to smaller fixed local mode counts. The mesh, frequency, number of subdomains, overlap, and oversampling were kept fixed, while only the number of retained local modes was varied:

$$f = 10\,\mathrm{Hz}, \qquad h = 3\,\mathrm{m}, \quad M = 64,$$
$$\text{overlap} = 1, \qquad \text{oversampling} = 2, \qquad n_i \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}. \tag{4.30}$$

The results are shown in Table 4.21. The solution difference is measured relative to the $n_i = 50$ solution.

Table 4.21: Effect of the number of retained local modes for the city-domain problem at 10 Hz on the $h = 3\,\mathrm{m}$ mesh. The setup uses $M = 64$, one overlap layer, and two oversampling layers.

n_i	n_c	GMRES	Residual	Time (s)	Relative difference vs. $n_i = 50$
5	320	243	6.97×10^{-6}	4417	3.68×10^{-6}
10	640	197	1.28×10^{-5}	4363	4.86×10^{-6}
15	960	179	1.65×10^{-5}	4580	6.04×10^{-6}
20	1280	169	1.88×10^{-5}	4978	5.35×10^{-6}
25	1600	151	2.14×10^{-5}	5265	4.91×10^{-6}
30	1920	112	2.91×10^{-5}	5062	5.63×10^{-6}
35	2240	85	2.64×10^{-5}	4888	4.09×10^{-6}
40	2560	72	2.27×10^{-5}	5044	2.78×10^{-6}
45	2880	63	2.48×10^{-5}	4840	1.98×10^{-6}
50	3200	55	2.68×10^{-5}	5070	0.00

The iteration count decreases strongly as more local modes are retained. With only $n_i = 5$, GMRES requires 243 iterations, while with $n_i = 50$, only 55 iterations are needed. This trend is shown in Figure 4.5. Thus, increasing the local spectral content improves the quality of the preconditioner substantially.

At the same time, the computed solutions are almost unchanged. The relative difference to the $n_i = 50$ solution is of order 10^{-6} for all smaller mode counts, as shown in Figure 4.6. This indicates that the retained mode count mainly affects solver efficiency, not the final converged solution, provided that GMRES reaches the prescribed residual tolerance.

The wall-clock time is not monotone in n_i. Although more modes reduce the number of GMRES iterations, they also increase the cost of applying the preconditioner and solving the coarse problem. The best choice is therefore not determined by the iteration count alone. In this experiment, $n_i = 50$ gives the smallest iteration count among the successful runs, while still remaining small enough to fit in memory.

Larger retained-mode counts were also attempted. The original offline run with $n_i = 100$ completed and produced the full coarse basis with $n_c = 6400$, but the corresponding online solve was killed during the setup stage because of the memory required to load and use the coarse basis. A run with $n_i = 55$, corresponding to $n_c = 3520$, also did not reach the GMRES iteration stage. Thus, for this city problem and implementation, the practical limit for the explicit coarse basis lies

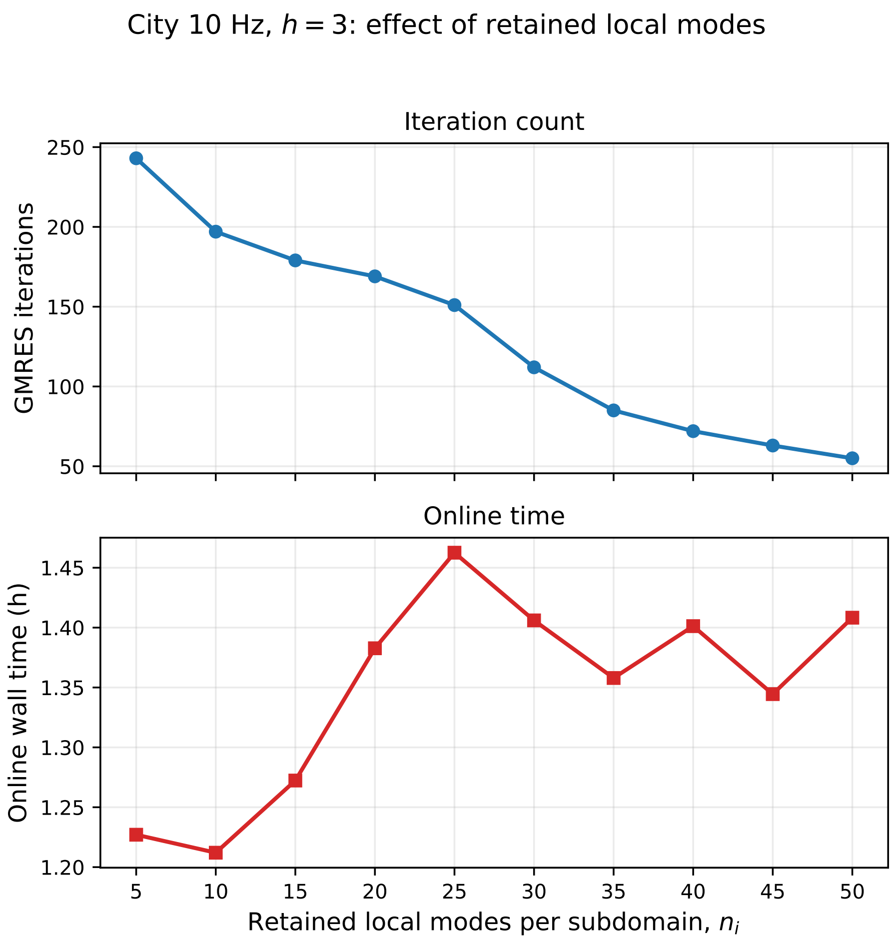

Figure 4.5: Effect of the retained number of local modes on the city-domain 10 Hz, $h = 3\,\mathrm{m}$ problem. Increasing n_i strongly reduces the GMRES iteration count, while the online wall time is not monotone because the cost of the coarse correction also increases.

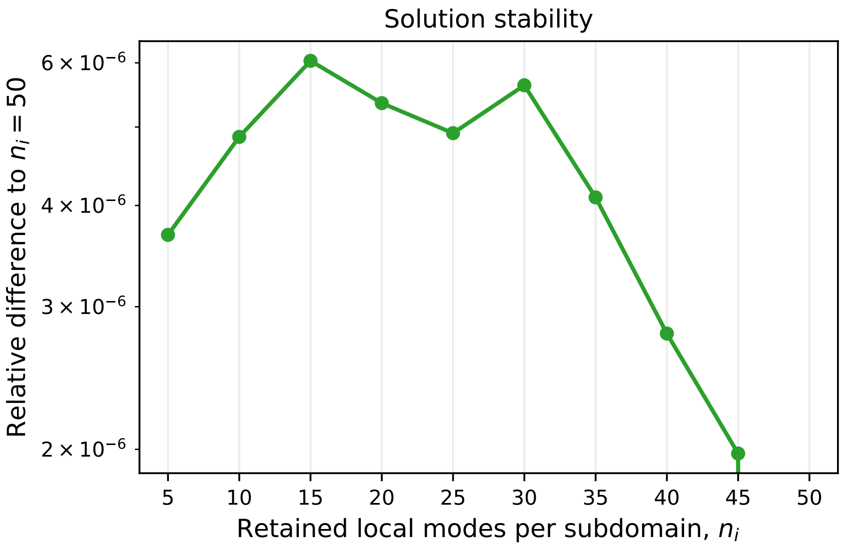

Figure 4.6: Relative solution difference for the city-domain 10 Hz, $h = 3$ m retained-mode study, measured against the $n_i = 50$ solution. The differences remain of order 10^{-6}, showing that the retained mode count mainly affects solver convergence rather than the final converged solution.

between $n_i = 50$ and $n_i = 55$. This memory limitation is discussed separately in Section 4.4.5.

4.4.5 Memory limitations

The city experiments revealed a practical memory limitation of the current implementation. The global coarse basis is stored explicitly as a complex matrix

$$\boldsymbol{E}_0 \in \mathbb{C}^{N \times n_c}. \tag{4.31}$$

Since the real and imaginary parts are stored separately in double precision, the storage required for the coarse basis alone is approximately

$$2 \cdot 8 \cdot N \cdot n_c \quad \text{B}. \tag{4.32}$$

This estimate does not include the sparse finite element matrix, local operators, local factorizations, the coarse matrix, or temporary arrays used during setup and application of the preconditioner.

The individual columns of $\boldsymbol{E}_0$ are not arbitrary global vectors. As described in Section 3.4, each coarse basis function is constructed from a local eigenmode on one oversampled patch, multiplied by the corresponding partition-of-unity weight:

$$\boldsymbol{\psi}_{i,j} = I_h(\chi_i \boldsymbol{\phi}_{i,j}). \tag{4.33}$$

Thus, each column of $\boldsymbol{E}_0$ has a patch-local origin. In the present implementation, these patch-local functions are nevertheless inserted into one explicit global matrix. This storage choice is simple, but it becomes expensive for large N and large n_c.

The limitation became visible in the 10 Hz city experiment on the $h = 3\,\mathrm{m}$ mesh, where

$$N = 1\,466\,914. \tag{4.34}$$

With $M = 64$ subdomains and $n_i = 100$ retained modes per subdomain, the coarse-space dimension is

$$n_c = 64 \cdot 100 = 6400. \tag{4.35}$$

The corresponding explicit coarse basis requires approximately

$$2 \cdot 8 \cdot 1\,466\,914 \cdot 6400 \approx 1.50 \times 10^{11}\,\mathrm{B} \approx 150\,\mathrm{GB} \tag{4.36}$$

for $\boldsymbol{E}_0$ alone. In practice, the online solve with this coarse space was killed during the setup stage.

Reducing the retained mode count to $n_i = 50$ gives

$$n_c = 64 \cdot 50 = 3200, \tag{4.37}$$

and reduces the storage estimate for the explicit coarse basis to approximately

$$2 \cdot 8 \cdot 1\,466\,914 \cdot 3200 \approx 7.51 \times 10^{10}\,\mathrm{B} \approx 75\,\mathrm{GB}. \tag{4.38}$$

This reduced setup completed successfully and gave the $h = 3\,\mathrm{m}$, 10 Hz result reported in Table 4.18.

A slightly larger reduced coarse space with $n_i = 55$, corresponding to $n_c = 3520$, did not reach the GMRES iteration stage. For this case, the coarse basis alone would require approximately

$$2 \cdot 8 \cdot 1\,466\,914 \cdot 3520 \approx 8.26 \times 10^{10}\,\mathrm{B} \approx 83\,\mathrm{GB}. \tag{4.39}$$

Thus, for this city problem and implementation, the practical memory limit of the explicit coarse basis lies between $n_i = 50$ and $n_i = 55$, or equivalently between $n_c = 3200$ and $n_c = 3520$.

This result shows that the coarse space cannot be chosen based only on GMRES iteration counts. Increasing n_i can improve the quality of the preconditioner and reduce the number of iterations, but it also increases the memory required to store and apply the coarse correction. For large three-dimensional city problems, memory is therefore a central part of the preconditioner design.

4.4.6 City-domain study at 20 Hz

The city-domain problem was also tested at 20 Hz. This frequency is more demanding than the 5 Hz and 10 Hz cases because the wavelength is shorter:

$$\lambda = \frac{343}{20} \approx 17.2\,\mathrm{m}. \tag{4.40}$$

The purpose of this experiment is to test how far the fixed-mode two-level preconditioner can be pushed on the city geometry.

The runs used $M = 64$ subdomains and $n_i = 40$ retained local modes per subdomain, giving

$$n_c = Mn_i = 2560. \tag{4.41}$$

For the $h = 5\,\mathrm{m}$ and $h = 4\,\mathrm{m}$ runs, two oversampling layers were used. For the $h = 3\,\mathrm{m}$ run, the oversampling was reduced to zero in order to keep the memory cost manageable. Therefore, the $h = 3\,\mathrm{m}$ result should be interpreted as a practical feasibility refinement rather than a perfectly controlled mesh-convergence point.

Table 4.22: City-domain runs at 20 Hz using a fixed number of retained local modes. The $h = 3\,\mathrm{m}$ run uses zero oversampling layers to reduce memory cost.

h	N	os	GMRES	Residual	$\|u_h\|_{\mathrm{rms}}$	$\max\|u_h\|$	Δ_{rms}
5.0	330073	2	131	5.78×10^{-6}	0.1963	6.9577	
4.0	634972	2	165	5.48×10^{-6}	0.1968	7.2175	0.27%
3.0	1466914	0	269	9.91×10^{-6}	0.1865	7.4910	5.21%

The $h = 5\,\mathrm{m}$ and $h = 4\,\mathrm{m}$ runs give very similar RMS values, with a relative change of only 0.27%. However, the finer $h = 3\,\mathrm{m}$ run changes the RMS value by 5.21% relative to the $h = 4\,\mathrm{m}$ result. This indicates that the 20 Hz city solution is not yet fully mesh-converged at $h = 4\,\mathrm{m}$. At the same time, the iteration count increases from 131 on the $h = 5\,\mathrm{m}$ mesh to 269 on the $h = 3\,\mathrm{m}$ mesh, showing that the linear systems become substantially more difficult as the mesh is refined.

The $h = 3\,\mathrm{m}$ run required about 8.92×10^4 seconds, or approximately 24.8 h, and the preconditioner data file reached about 60 GB before being deleted after the successful solve. This makes the run an important feasibility result: the method can solve a 20 Hz city-domain problem on the $h = 3\,\mathrm{m}$ mesh, but the computational cost is already very high.

4.4.7 City-domain feasibility study at 30 Hz

A final city-domain diagnostic was performed at 30 Hz. At this frequency, the wavelength is

$$\lambda = \frac{343}{30} \approx 11.4\,\mathrm{m}. \tag{4.42}$$

This makes the problem considerably harder than the lower-frequency city cases. The aim of this study is therefore not to claim a fully mesh-converged 30 Hz city solution, but to test the practical limits of the fixed-mode two-level preconditioner on the largest city meshes.

The first set of runs used $M = 64$ subdomains, one overlap layer, one oversampling layer, and $n_i = 60$ retained local modes per subdomain. This gives

$$n_c = Mn_i = 3840. \tag{4.43}$$

The results are shown in Table 4.23.

The $h = 5\,\mathrm{m}$ and $h = 4\,\mathrm{m}$ runs completed successfully. The $h = 3\,\mathrm{m}$ run completed the offline construction of the preconditioner, but the online solve was

Table 4.23: City-domain feasibility runs at 30 Hz using $M = 64$, one overlap layer, one oversampling layer, and $n_i = 60$ retained local modes per subdomain.

h	N	n_c	GMRES iterations	Residual	Time (s)	Status
5.0	330073	3840	155	2.66×10^{-6}	11816	done
4.0	634972	3840	178	2.01×10^{-6}	34718	done
3.0	1466914	3840	–	–	124023	killed

killed after loading the coarse basis. The saved preconditioner data for this case was about 95 GB, showing that the explicit coarse basis becomes the limiting factor at this frequency and mesh size.

To investigate whether the $h = 3$ m problem could still be solved with a smaller coarse space, the saved $n_i = 60$ offline data were reduced by retaining fewer local modes per subdomain. The reduced cases used the same mesh, frequency, partition, overlap, and oversampling as the killed $n_i = 60$ run. Only the number of retained local modes was changed. The results are shown in Table 4.24.

Table 4.24: Reduced retained-mode online solves for the city-domain 30 Hz, $h = 3$ m case. The reduced preconditioners were obtained from the saved $n_i = 60$ offline data.

n_i	n_c	GMRES iterations	Residual	Time (s)	Status
20	1280	102	2.22×10^{-6}	3685	done
30	1920	113	2.39×10^{-6}	4618	done
40	2560	122	2.88×10^{-6}	5819	done
50	3200	–	–	2151	killed

The reduced $h = 3$ m runs show that the 30 Hz city problem can be solved when the coarse space is made smaller. However, increasing the number of retained modes does not monotonically reduce the GMRES iteration count in this case. This is not unexpected for a nonsymmetric, indefinite preconditioned Helmholtz system: enriching the coarse space changes the full preconditioned operator, but it does not guarantee a monotone decrease in GMRES iterations. The main conclusion from this diagnostic is practical. For the 30 Hz city problem on the $h = 3$ m mesh, the implementation can solve reduced coarse spaces up to $n_i = 40$, while $n_i = 50$ and the original $n_i = 60$ case exceed the available memory during the online phase.

4.4.8 Plain GMRES baseline on the city mesh

As an additional reference, plain GMRES without preconditioning was tested on the coarsest city mesh. In this baseline, the same finite element matrix and right-hand side are used as in the preconditioned city experiments, but no local Schwarz correction and no coarse correction are applied. The purpose is to measure the algebraic difficulty of the unpreconditioned linear systems, not to propose plain GMRES as a practical solver for city-scale simulations.

The results for the city mesh with $h = 5\,\mathrm{m}$ are shown in Table 4.25. All cases reached the prescribed relative residual tolerance, but only after thousands of GMRES iterations.

Table 4.25: Plain GMRES baseline without preconditioning on the city mesh with $h = 5\,\mathrm{m}$ and $N = 330073$ degrees of freedom.

Frequency [Hz]	GMRES iterations	Final residual	Time [s]
5	8478	1.00×10^{-6}	1004
10	18324	1.00×10^{-6}	2269
20	24902	1.00×10^{-6}	2931
30	17440	1.00×10^{-6}	2142

The baseline shows that the unpreconditioned systems can be solved algebraically on the coarsest city mesh, but the cost is very high. This should be interpreted only as linear solver convergence: a small residual means that the discrete system has been solved, not that the mesh is fine enough to resolve the physical acoustic field at all frequencies. Compared with the two-level preconditioned city runs, where the iteration counts are on the order of tens to hundreds, plain GMRES is not practical for the intended large-scale simulations.

5

Discussion

The purpose of this thesis was to evaluate how a two-level spectral domain decomposition preconditioner performs for finite element Helmholtz problems arising in large-scale urban acoustic simulation. The results show that the method can strongly improve GMRES convergence, but also that linear solver convergence, acoustic mesh convergence, and memory cost must be interpreted separately. This chapter discusses the main observations and relates them to the research questions stated in Chapter 1.

5.1 Verification and role of the plane-wave test

The exact plane-wave test verifies the finite element implementation. Since the exact solution is known, the error can be measured directly in the relative L^2 norm and relative H^1-seminorm. The observed rates are approximately second order in L^2 and first order in the H^1-seminorm, which is the expected behaviour for first-order Lagrange finite elements on smooth solutions.

This result is important because the later box-domain and city-domain experiments do not have known exact solutions. In those cases, the computed solution can only be assessed through mesh refinement, receiver comparisons, residuals, and consistency between runs. The plane-wave verification therefore checks the finite element discretization before the experiments are used to study preconditioner performance and acoustic resolution.

5.2 Solver convergence and acoustic convergence

A basic observation is that GMRES convergence and acoustic mesh convergence are different properties. GMRES convergence means that the discrete linear system

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b} \tag{5.1}$$

has been solved to the prescribed residual tolerance. It does not by itself imply that the finite element solution is mesh-converged as an approximation of the continuous acoustic field.

This distinction is clearest in the box receiver convergence study. At 5 Hz, the receiver errors decrease from 4.93% on the $h = 5\,\mathrm{m}$ mesh to 1.38% on the $h = 3\,\mathrm{m}$ mesh when compared with the $h = 2.5\,\mathrm{m}$ reference solution. At 10 Hz, the receiver errors are larger, decreasing from 23.94% on the $h = 4\,\mathrm{m}$ mesh to 4.89% on the

$h = 2.5\,\mathrm{m}$ mesh relative to the $h = 2\,\mathrm{m}$ reference. At 20 Hz, the receiver errors remain much larger: 108.60%, 63.98%, and 35.79% for $h = 4\,\mathrm{m}$, $h = 3\,\mathrm{m}$, and $h = 2.5\,\mathrm{m}$, respectively. Thus, the higher-frequency solutions are more sensitive to mesh resolution, even when the linear systems are solved to small residuals.

The same pattern appears in the city-domain experiments. The 5 Hz city case gives the clearest mesh-refinement behaviour, with RMS values changing from 0.6972 on the $h = 5\,\mathrm{m}$ mesh to 0.6794 on the $h = 3\,\mathrm{m}$ mesh. The 10 Hz city case remains solvable, but requires a much larger coarse space on the finest mesh. The 20 Hz and 30 Hz city cases are therefore best interpreted as feasibility tests of the solver and preconditioner, not as fully mesh-converged acoustic simulations.

5.3 Effect of the two-level preconditioner

The first research question asks how the two-level preconditioner affects GMRES convergence on controlled and realistic geometries. The experiments show that the coarse correction is beneficial, especially as the problem becomes more difficult. Across the successful two-level runs, the GMRES iteration counts generally remain moderate and stable compared with the unpreconditioned baseline, even as the mesh and geometry become more demanding.

On the box domain, the one-level and two-level comparison shows that the effect is moderate at 5 Hz but substantial at 10 Hz. For example, at 5 Hz, the $h = 5\,\mathrm{m}$ case changes from 32 GMRES iterations with the one-level method to 27 iterations with the two-level method. On the $h = 3\,\mathrm{m}$ mesh, the iteration count changes from 43 to 36. The improvement is more pronounced for the 10 Hz, $h = 2.5\,\mathrm{m}$ case, where the one-level method requires 56 GMRES iterations while the two-level method requires 22 iterations.

This behaviour is consistent with the role of the spectral coarse space. The local Schwarz corrections reduce local error components, while the coarse correction represents global or slowly decaying error components that are difficult to remove using local solves alone. For Helmholtz problems, these global components become more important as the frequency and problem size increase. The results therefore support the use of a two-level method rather than a purely one-level Schwarz preconditioner.

The plain GMRES baseline gives an additional reference point. On the box-domain mesh with $h = 2.5\,\mathrm{m}$, unpreconditioned GMRES required between 541 and 694 iterations for the tested frequencies. This is much larger than the corresponding preconditioned iteration counts. In particular, at 10 Hz, plain GMRES required 694 iterations, compared with 56 iterations for the one-level method and 22 iterations for the two-level method. This comparison shows that the local Schwarz correction already gives a large improvement over solving the unpreconditioned Helmholtz system, and that the spectral coarse correction gives a further improvement.

The city plain-GMRES baseline reinforces the same conclusion on the realistic geometry. Even on the coarsest city mesh, with $h = 5\,\mathrm{m}$ and $N = 330073$ degrees of freedom, solving without preconditioning required between 8478 and 24902 GMRES iterations for the tested frequencies. In contrast, the two-level preconditioned city runs required only tens to hundreds of iterations. Thus, the city baseline confirms that the preconditioner is essential for making the realistic city-domain solves practical.

5.4 Influence of frequency and mesh resolution

The second research question concerns how frequency and mesh resolution influence solution behaviour, receiver convergence, and solver performance. The results show a clear frequency dependence.

At a fixed sound speed $c = 343\,\mathrm{m\,s^{-1}}$, the wavelength is

$$\lambda = \frac{c}{f}. \tag{5.2}$$

The wavelengths used in the thesis are approximately 68.6 m, 34.3 m, 17.2 m, and 11.4 m for 5 Hz, 10 Hz, 20 Hz, and 30 Hz, respectively. As the frequency increases, the wavelength decreases, and the same mesh size gives fewer degrees of freedom per wavelength. This affects both the acoustic solution and the linear solver.

The box-domain results show that the 5 Hz case is well resolved by the available meshes. At 10 Hz, the receiver errors are still reduced by refinement, but the coarser meshes are less accurate. At 20 Hz and 30 Hz, the available meshes are too coarse to claim full acoustic convergence. The 30 Hz box study shows that the linear systems can still be solved with the fixed-mode preconditioner, but the RMS value changes by 12.55% between $h = 4\,\mathrm{m}$ and $h = 3\,\mathrm{m}$, and by 6.53% between $h = 3\,\mathrm{m}$ and $h = 2.5\,\mathrm{m}$. This indicates improvement under refinement, but not full mesh convergence.

For the city domain, the same trend is stronger because the geometry is more complex and the number of degrees of freedom is much larger. The 20 Hz city run on the $h = 3\,\mathrm{m}$ mesh required 269 GMRES iterations and used no oversampling in order to keep the memory cost manageable. The 30 Hz, $h = 3\,\mathrm{m}$ run with $n_i = 60$ retained local modes per subdomain did not complete the online phase. Thus, the higher-frequency city cases reveal both acoustic resolution challenges and memory limitations in the preconditioner.

5.5 Influence of preconditioner parameters

The third research question concerns the effect of preconditioner parameters: the spectral threshold, oversampling size, number of subdomains, and number of retained local modes.

In the threshold-based box-domain experiment at 10 Hz, increasing the spectral threshold ρ reduces the coarse-space dimension but increases the GMRES iteration count. For example, changing ρ from 0.80 to 1.25 reduces n_c from 476 to 152, while the iteration count increases from 22 to 41. This shows the expected tradeoff: a richer coarse space is more expensive but gives a stronger preconditioner.

Oversampling also has a visible effect. In the box-domain 10 Hz oversampling study, increasing the oversampling from 1 to 4 layers changes the iteration count from 29 to 22 and improves the final residual. In the fixed-mode 20 Hz box study with $M = 64$ and $n_i = 100$, the effect is even clearer: the iteration count decreases from 67 with no oversampling to 12 with three oversampling layers. This indicates that oversampling improves the quality of the local spectral basis, although it also increases the cost of the local eigenproblems and local operators.

The number of subdomains also changes the balance between local and coarse costs. In the 20 Hz fixed-mode box study with $n_i = 100$, increasing M from 8 to 64 increases n_c from 800 to 6400 and reduces the GMRES iteration count from 27 to 12. However, the wall-clock time does not decrease monotonically, because the cost of the coarse correction and preconditioner application also changes. Thus, iteration count alone is not sufficient for selecting M.

The retained-mode city study at 10 Hz shows the clearest effect of n_i. On the $h = 3\,\mathrm{m}$ city mesh, increasing n_i from 5 to 50 reduces the GMRES iteration count from 243 to 55. The relative solution difference, measured against the $n_i = 50$ solution, also decreases as more modes are retained. This confirms that the retained local spectral modes are important for convergence on the city geometry.

However, the 30 Hz city diagnostic shows that this behaviour is not a simple monotone rule in every case. For the reduced $h = 3\,\mathrm{m}$, 30 Hz runs obtained from the saved $n_i = 60$ offline data, the cases $n_i = 20$, 30, and 40 all completed, with 102, 113, and 122 GMRES iterations, respectively. This does not contradict the usefulness of the coarse space. It shows that, for a nonsymmetric and indefinite preconditioned Helmholtz operator, increasing the retained mode count changes the whole preconditioned operator and does not guarantee a monotone decrease in GMRES iterations. In this case, the main result is practical: smaller coarse spaces made the finest 30 Hz city case solvable within the available memory, whereas larger coarse spaces did not.

5.6 Memory limitations of the explicit coarse basis

The fourth research question concerns the memory limitations caused by the spectral coarse space. The detailed estimates in Section 4.4.5 show that the dominant memory cost is the explicit storage of the global coarse basis $\boldsymbol{E}_0$. This cost scales with the product of the number of finite element degrees of freedom and the coarse-space dimension.

This explains why the largest city runs were not limited only by the mesh size N, but by the combined quantity Nn_c. Refining the mesh increases N, while increasing the number of retained local modes or the number of subdomains increases n_c. Both therefore increase the storage required for the coarse basis. In the finest city cases, this became the practical bottleneck before the linear iteration itself could be completed.

The main consequence is that the richest coarse space is not always the best practical choice. A larger coarse space can improve GMRES convergence, but only if it can be stored and applied within the available memory. This was visible in the city experiments: reducing the number of retained modes sometimes made a run possible even when a larger coarse space failed. The parameter choice must therefore balance three effects: convergence improvement, acoustic consistency, and memory feasibility.

The memory limitation is also connected to the current implementation strategy. The local spectral modes are computed on local oversampled patches, but after selection they are stored in one global matrix $\boldsymbol{E}_0$. Since these basis functions have local support before being inserted into the global coarse basis, this suggests that a more memory-efficient implementation should exploit the patch-local structure. This point is returned to in the outlook part of Chapter 6.

5.7 Practical guidelines for parameter choice

The experiments suggest practical guidelines for choosing the preconditioner parameters for a given frequency. These guidelines are not universal, because the best choice depends on the mesh, geometry, memory limit, and implementation. However, the results show several clear trends.

The first requirement is that the mesh must resolve the acoustic wavelength. For a frequency f, the wavelength is

$$\lambda = \frac{c}{f}. \tag{5.3}$$

As f increases, λ decreases, and both the finite element space and the coarse space must become richer. Therefore, higher frequencies generally require finer meshes, more robust preconditioning, and more careful memory control.

For the threshold-based coarse space, the spectral threshold ρ controls how many local modes are retained. A smaller value of ρ gives a larger coarse space and usually improves GMRES convergence, but increases memory use. In the box-domain 10 Hz experiment, decreasing the threshold from $\rho = 1.25$ to $\rho = 0.80$ increased the coarse-space dimension from 152 to 476, while reducing the GMRES iteration count from 41 to 22. Thus, ρ should be chosen as small as memory allows when robust convergence is the main objective.

For fixed-mode coarse spaces, the number of retained local modes n_i gives direct control of the coarse-space dimension,

$$n_c = M n_i. \tag{5.4}$$

The city 10 Hz retained-mode experiment shows that increasing n_i can strongly improve convergence: on the $h = 3$ m city mesh, the GMRES iteration count decreased from 243 for $n_i = 5$ to 55 for $n_i = 50$. However, the memory cost also increases linearly with n_i. Therefore, n_i should be increased only until the iteration count is acceptable and the solution is stable enough for the intended comparison.

Oversampling improves the quality of the local spectral basis, but also increases the size of the local problems. In the box-domain 20 Hz fixed-mode experiment, using $M = 64$ and $n_i = 100$, the GMRES iteration count decreased from 67 with no oversampling to 12 with three oversampling layers. This suggests that oversampling is especially important at higher frequencies. However, on the largest city meshes, oversampling may have to be reduced to keep the memory and setup cost feasible, as was done for the 20 Hz city $h = 3$ m run.

The number of subdomains M also affects the balance between local and coarse costs. Increasing M reduces the size of the local subdomain problems, but increases the number of coarse basis functions when n_i is fixed. For fixed n_i, the coarse-space dimension grows as

$$n_c = Mn_i. \tag{5.5}$$

Thus, increasing M can improve local scalability, but may also increase the cost of the coarse correction.

A practical selection strategy is therefore:

1. choose the finest mesh that is computationally feasible and gives a reasonable number of mesh points per wavelength;
2. start with a moderate number of subdomains, so that local problems are not too large;
3. use one or more oversampling layers when memory allows;
4. choose n_i or ρ so that the coarse-space dimension fits in memory;
5. increase the coarse-space richness only if GMRES convergence is too slow or if the solution changes significantly when more modes are retained.

For the city-domain experiments in this thesis, the results suggest the following interpretation. At 5 Hz, the method is relatively robust and moderate coarse spaces are sufficient. At 10 Hz, the finest city mesh requires a richer coarse space, and $n_i = 50$ was the largest retained-mode count that completed successfully in the fixed-mode study. At 20 Hz, the problem is more demanding, and the $h = 3$ m run required reducing oversampling to remain feasible. At 30 Hz, the city problem should be treated as a feasibility test: $h = 5$ m and $h = 4$ m completed with $n_i = 60$, but the $h = 3$ m case required reducing the retained-mode count, and larger coarse spaces exceeded the available memory.

Thus, the parameter choice should be frequency-dependent and memory-aware. The best practical choice is not necessarily the richest coarse space, but the smallest coarse space that gives stable GMRES convergence and acceptable solution consistency within the available memory.

6 Conclusion and Outlook

This thesis studied finite element solution of the time-harmonic Helmholtz equation for acoustic simulations in urban environments. The main focus was the implementation and evaluation of a two-level preconditioned GMRES solver based on restricted additive Schwarz local corrections and a spectral coarse space.

The finite element implementation was first verified on an exact plane-wave solution. The observed convergence rates agree with the expected behaviour for first-order Lagrange finite elements, with approximately second-order convergence in the relative L^2 norm and first-order convergence in the relative H^1-seminorm. This confirms that the finite element discretization and the basic Helmholtz implementation behave correctly on a problem with a known exact solution.

The box-domain experiments showed the expected frequency dependence of the Helmholtz problem. At 5 Hz, the solution was well resolved by the available meshes. At 10 Hz, the problem became more demanding, and at 20 Hz and 30 Hz, the available meshes were not sufficient to claim full acoustic mesh convergence. The two-level coarse correction improved GMRES convergence compared with the one-level method, especially for the more demanding cases. The experiments also showed that the spectral threshold, oversampling, number of subdomains, and retained local mode count all influence the balance between convergence and computational cost.

The city-domain experiments demonstrated that the method can be applied to realistic urban geometry, but also revealed the main practical limitation of the current implementation. The 5 Hz city case gave the clearest mesh-refinement behaviour, while the 10 Hz, 20 Hz, and 30 Hz cases became increasingly demanding. The largest runs showed that explicitly storing the global coarse basis can become memory-limiting. In particular, the 10 Hz and 30 Hz city experiments showed that small changes in the number of retained local modes can determine whether the online solve fits in memory.

Overall, the results show that a two-level spectral Schwarz preconditioner is a promising approach for large Helmholtz problems in urban acoustic simulations. However, they also show that solver convergence and acoustic mesh convergence are separate issues: a small GMRES residual does not guarantee that local acoustic quantities are mesh-converged. For higher frequencies and realistic city meshes, parameter choices must therefore be both frequency-aware and memory-aware.

Several concrete directions remain for future work. The first step toward more physically predictive simulations is source calibration. In this thesis, the source is represented by a numerical Gaussian forcing term with fixed amplitude and width. This is suitable for studying solver behaviour, but it is not yet a calibrated physical sound source. A natural next step is to calibrate the Gaussian source against an analytical Green's-function solution, for example a point source above hard ground using the image-source construction. This would make the computed amplitudes easier to interpret physically and would provide a better starting point for comparison with measured or reference sound levels.

Another important direction is to reduce the memory cost of the coarse correction. In the present implementation, the local spectral problems are computed patch by patch, but the selected modes are stored in one explicit global coarse basis matrix $\boldsymbol{E}_0$. This becomes the main bottleneck for the largest city problems. A more scalable implementation should exploit the patch-local structure of the modes, for example by storing the basis functions only on their local supports, applying $\boldsymbol{E}_0$ and $\boldsymbol{E}_0^*$ in a matrix-free way, or distributing the coarse-basis storage and coarse correction across several compute nodes. For very large problems, the coarse problem itself may also need to be solved iteratively or by another domain decomposition method instead of by a direct solver.

Further work is also needed on the physical model. The present experiments use a homogeneous acoustic medium, sound-hard building and ground surfaces, and absorbing outer boundary facets. More realistic urban sound propagation would require calibrated impedance boundary conditions, finite-impedance or porous ground models, atmospheric effects, and improved open-boundary treatments. These extensions should be introduced only after the source and simpler hard-ground cases have been carefully verified.

The problem sizes studied here are relevant for realistic city-scale simulation. The finest city mesh used in the thesis has approximately 1.47×10^6 finite element degrees of freedom, and the coarsest city mesh already has more than 3.3×10^5 degrees of freedom. These sizes are large enough to expose the computational bottlenecks that appear in three-dimensional urban Helmholtz problems, especially the memory cost of the spectral coarse basis. At the same time, the experiments show that even larger and more highly resolved meshes would be needed for fully mesh-converged simulations at higher frequencies.

The frequencies studied in the thesis are also relevant as low-frequency acoustic test cases. Frequencies between 5 Hz and 30 Hz correspond to wavelengths between approximately 68.6 m and 11.4 m in air. These wavelengths are large compared with many urban details, but they are useful for testing wave-based propagation effects on city-scale geometries. The results show that the method is most robust at the lower end of this range, while the higher-frequency cases already reveal the expected increase in mesh resolution, coarse-space size, and memory demand.

Finally, the parameter studies in this thesis could be extended into a more systematic parameter-selection strategy. The results show that the number of retained local modes, oversampling size, number of subdomains, mesh size, and frequency must be chosen together. A useful next step would be to develop frequency- and memory-aware rules for choosing these parameters, and to compare the resulting solver with other preconditioning strategies on the same box and city-domain test cases.

In the longer term, a scalable version of the solver could be useful as part of a city noise-modelling workflow. The present work does not yet provide a validated engineering tool for environmental noise prediction, but it demonstrates solver components that could support such a tool: finite element Helmholtz modelling on city meshes, two-level preconditioned GMRES, and diagnostics for convergence and memory cost. A natural implementation path would be to integrate a more scalable version of the solver into `dtcc-sim` as part of the DTCC platform, so that practitioners could use the method together with existing city-geometry and mesh-generation tools.

A

Additional diagnostics

This appendix collects diagnostic experiments that support the main numerical results but are not part of the main experiment sequence in Chapter 4. The diagnostics are used to check solver behaviour, inspect the local spectral coarse-space construction, and explain the source-amplitude normalization used for visualizing the computed acoustic fields.

A.1 Damped Richardson diagnostic

All main solves in this thesis use GMRES applied to the left-preconditioned system. As an additional diagnostic, a damped Richardson iteration was tested on one representative Helmholtz problem. The purpose is not to replace GMRES, but to illustrate why a Krylov method is needed for the indefinite systems considered in this work.

The diagnostic was performed on the box-domain 20 Hz case with $h = 4\,\mathrm{m}$, $M = 16$, one overlap layer, three oversampling layers, and the threshold-based coarse space with $\rho = 1.10$. The system size was

$$N = 54621, \tag{A.1}$$

and the coarse-space dimension was

$$n_c = 915. \tag{A.2}$$

For a fixed damping parameter ω, the preconditioned Richardson iteration is

$$\boldsymbol{u}^{(m+1)} = \boldsymbol{u}^{(m)} + \omega \boldsymbol{\mathcal{M}}^{-1} \left(\boldsymbol{b} - \boldsymbol{A}\boldsymbol{u}^{(m)} \right), \tag{A.3}$$

where $\boldsymbol{A}$ is the finite element Helmholtz matrix, $\boldsymbol{\mathcal{M}}^{-1}$ is the two-level preconditioner, and ω is a scalar relaxation parameter. The true relative residual monitored in the diagnostic is

$$\frac{\left\| \boldsymbol{b} - \boldsymbol{A}\boldsymbol{u}^{(m)} \right\|_2}{\left\| \boldsymbol{b} \right\|_2}. \tag{A.4}$$

Several damping parameters were tested for 30 Richardson iterations. The final residuals are shown in Table A.1. None of the tested values reached the GMRES tolerance 10^{-6}.

Table A.1: Damped Richardson diagnostic for the box-domain 20 Hz case with $h = 4\,\mathrm{m}$, $M = 16$, one overlap layer, three oversampling layers, and $\rho = 1.10$.

ω	Iterations	Final relative residual	Time (s)
0.5	30	5.52×10^{0}	41.12
0.2	30	5.87×10^{-2}	41.20
0.1	30	2.36×10^{-1}	41.49
0.05	30	4.71×10^{-1}	41.26
0.02	30	6.89×10^{-1}	41.37
0.01	30	7.96×10^{-1}	41.17
0.005	30	8.78×10^{-1}	41.15

The residual histories for selected damping parameters are shown in Figure A.1. The case $\omega = 0.5$ initially decreases the residual, but then diverges. Smaller damping parameters are more stable, but converge very slowly. The best tested value, $\omega = 0.2$, reduces the residual to about 5.9×10^{-2} after 30 iterations, still far from the GMRES tolerance.

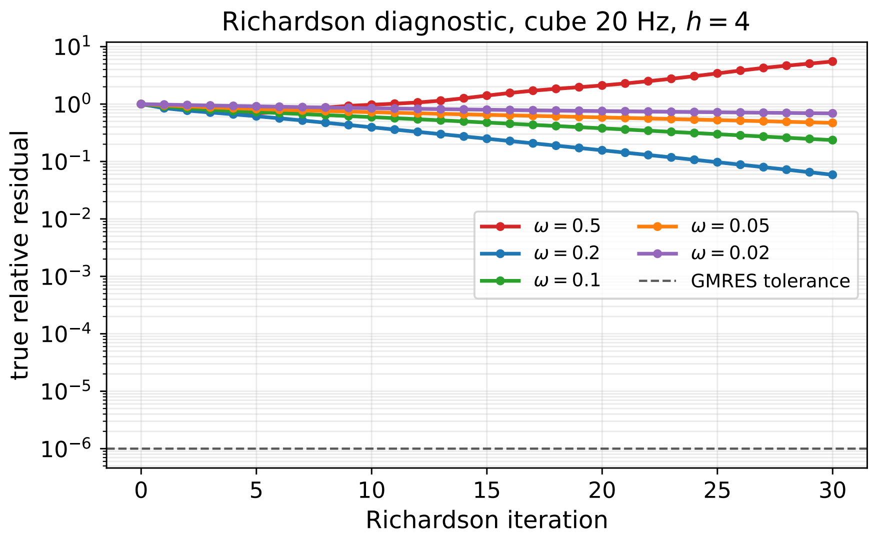

Figure A.1: Damped Richardson residual histories for the box-domain 20 Hz, $h = 4\,\mathrm{m}$ diagnostic problem. The dashed horizontal line marks the GMRES tolerance used in the main experiments.

This diagnostic confirms that the preconditioner improves the system but does not make a simple stationary iteration robust enough for the Helmholtz problems considered here. The indefinite and non-normal character of the preconditioned Helmholtz operator makes the choice of damping parameter delicate. GMRES is therefore used throughout the main experiments because it adapts the residual minimization over a Krylov subspace instead of relying on one fixed damping parameter.

A.2 Local eigenvalue diagnostics

A second diagnostic concerns the local eigenvalue problems used to construct the spectral coarse space. This diagnostic was motivated by correspondence with the authors of the RAS/MS-GFEM preconditioner. They pointed out that the local eigenvalues should be real and non-negative, up to numerical roundoff, and suggested inspecting the computed spectra directly. This is useful because GMRES iteration counts alone may not reveal small issues in the construction of the coarse basis.

The diagnostic shown here was performed for the box-domain problem at 20 Hz on the $h = 2\,\text{m}$ mesh. The setup used

$$M = 16, \qquad \text{overlap} = 1, \qquad \text{oversampling} = 3, \qquad \rho = 1.10. \tag{A.5}$$

For this run, the total coarse-space dimension was

$$n_c = 1326, \tag{A.6}$$

corresponding to an average of approximately 82.88 retained modes per subdomain.

Figure A.2 shows the raw local eigenvalue spectrum for patch 13. This patch retained 109 local modes. The eigenvalues are plotted on a semilogarithmic scale to make the structure of the spectrum visible. This plot checks that the local eigenproblem produces a reasonable ordered spectrum and that the retained modes come from the expected part of the spectrum.

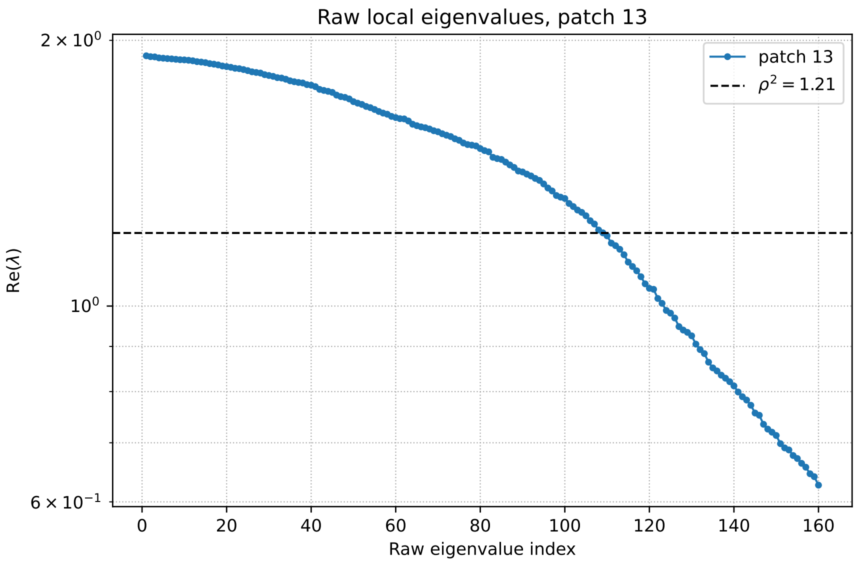

Figure A.2: Raw local eigenvalue spectrum for patch 13 in the box-domain 20 Hz, $h = 2\,\text{m}$ diagnostic run. The spectrum is plotted on a semilogarithmic scale to inspect the ordering and magnitude of the local eigenvalues entering the spectral coarse-space construction.

Figure A.3 shows how many local eigenmodes were retained on each subdomain by the threshold criterion. The retained mode count is not constant across patches, because each local eigenproblem depends on the local subdomain and its oversampled region. This explains why a fixed spectral threshold can produce different local coarse-space dimensions on different subdomains.

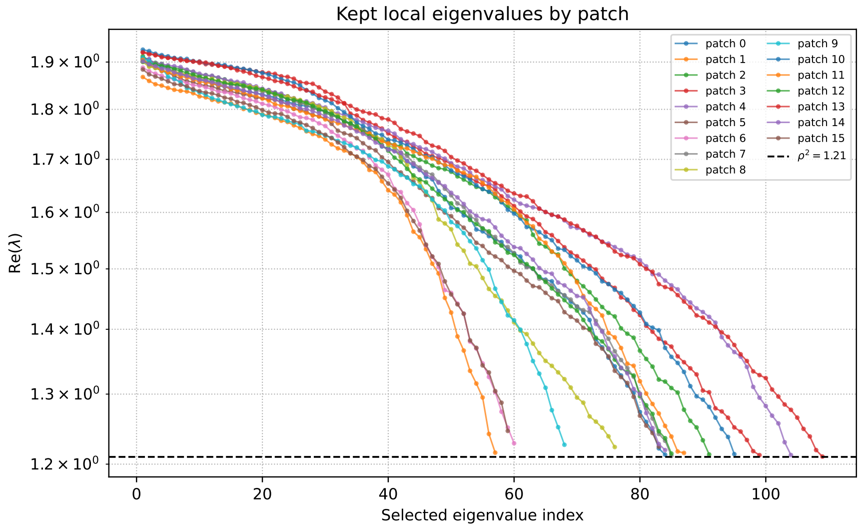

Figure A.3: Number of retained local eigenmodes for the box-domain 20 Hz, $h = 2\,\mathrm{m}$ diagnostic run with $M = 16$, one overlap layer, three oversampling layers, and $\rho = 1.10$.

The diagnostic supports two points. First, the local spectra can be inspected directly, instead of relying only on the final GMRES iteration count. Second, it shows why the number of retained modes can become large at higher frequency and finer mesh resolution. In the main city-domain experiments, this motivated the fixed-mode strategy, where n_i is prescribed directly in order to control the coarse-space dimension and the memory required by the explicit coarse basis.

A.3 Source-amplitude calibration diagnostics

The numerical experiments in this thesis use a Gaussian right-hand side as a smooth approximation of a point source. This is convenient numerically, because a true Dirac delta source is singular and is not represented directly in the finite element space. However, it also means that the source amplitude used in the code is not automatically a calibrated physical sound-pressure level. The computed solution is therefore initially a relative complex pressure amplitude.

This issue was discussed with Jens Forssén. A physically meaningful acoustic calibration should relate the numerical source strength to an analytical reference solution, such as a point source in free space or a point source above a reflecting ground plane. The purpose of the diagnostic in this appendix is to explain how the source amplitude was interpreted and why the final thesis uses a simple free-field normalization.

Free-space point-source reference

For the Helmholtz equation in free space, the Green's function for a point source at $\boldsymbol{x}_0$ is

$$G(\boldsymbol{x}, \boldsymbol{x}_0) = \frac{e^{ik|\boldsymbol{x}-\boldsymbol{x}_0|}}{4\pi|\boldsymbol{x}-\boldsymbol{x}_0|}, \tag{A.7}$$

up to the sign convention used for the time dependence. Since this thesis uses the convention $e^{-i\omega t}$, the sign in the exponential is consistent with outgoing waves. The magnitude of the pressure from a point source with complex strength Q is therefore

$$|u(r)| = \frac{|Q|}{4\pi r}, \qquad r = |\boldsymbol{x}-\boldsymbol{x}_0|. \tag{A.8}$$

The numerical source used in the experiments is not a Dirac delta source, but a Gaussian source of the form

$$q(\boldsymbol{x}) = A_s \exp\left(-\frac{\|\boldsymbol{x}-\boldsymbol{x}_0\|^2}{2\sigma^2}\right), \tag{A.9}$$

where A_s is the numerical source amplitude and σ is the Gaussian width. The effective free-space point-source strength associated with this Gaussian source is approximated by

$$Q_{\text{eff}} = A_s(2\pi)^{3/2}\sigma^3 \exp\left(-\frac{k^2\sigma^2}{2}\right). \tag{A.10}$$

This expression is the Fourier response of the Gaussian source at wavenumber k. It gives a frequency-dependent effective source strength.

To normalize the computed solution so that the equivalent free-space point source gives 1 Pa at $r = 1\,\text{m}$, the numerical solution is multiplied by

$$\alpha = \frac{4\pi}{|Q_{\text{eff}}|}. \tag{A.11}$$

The resulting scale factors for the frequencies used in the thesis are shown in Table A.2.

Table A.2: Free-field source normalization for the Gaussian source with $\sigma = 5\,\text{m}$ and $A_s = 1$. The scale factor α multiplies the numerical solution so that the equivalent free-space source gives 1 Pa at $r = 1\,\text{m}$.

f (Hz)	k	Q_{eff}	α
5	9.159×10^{-2}	1.773×10^{3}	7.089×10^{-3}
10	1.832×10^{-1}	1.294×10^{3}	9.709×10^{-3}
20	3.664×10^{-1}	3.677×10^{2}	3.417×10^{-2}
30	5.495×10^{-1}	4.515×10^{1}	2.783×10^{-1}

The scale factor increases with frequency because the Gaussian source has a frequency-dependent effective strength. For the same numerical amplitude $A_s = 1$, the factor $\exp(-k^2\sigma^2/2)$ becomes smaller as k increases. Thus, the effective point-source strength decreases at higher frequency, and a larger scaling factor is needed to obtain the same free-field reference pressure.

Hard-ground image-source diagnostic

In the discussion with Jens Forssén, a second calibration idea was also considered. For a point source above an ideal perfectly reflecting ground plane, the analytical solution can be written as the sum of a direct source and an image source mirrored below the ground. If the source is located at $\boldsymbol{x}_0$, the image source is located at $\boldsymbol{x}_0^{\text{img}}$. The hard-ground reference field is then

$$G_{\text{hg}}(\boldsymbol{x}) = \frac{e^{ik|\boldsymbol{x}-\boldsymbol{x}_0|}}{4\pi|\boldsymbol{x}-\boldsymbol{x}_0|} + \frac{e^{ik|\boldsymbol{x}-\boldsymbol{x}_0^{\text{img}}|}}{4\pi|\boldsymbol{x}-\boldsymbol{x}_0^{\text{img}}|}. \tag{A.12}$$

This model is more appropriate than the free-space Green's function when the ground is treated as a perfectly reflecting boundary.

As a diagnostic, the numerical solutions were fitted to this hard-ground Green's-function model over several radial sampling shells. If the numerical source and computational domain behaved like an ideal point source above an infinite reflecting plane, the fitted source strength would be nearly constant across the shells and the relative fit error would be small.

For the 20 Hz box-domain case on the $h = 2$ m mesh, the source was located at

$$\boldsymbol{x}_0 = (100, 100, 8), \tag{A.13}$$

and the image source was therefore

$$\boldsymbol{x}_0^{\text{img}} = (100, 100, -8). \tag{A.14}$$

The resulting fitted source strengths are shown in Table A.3.

Table A.3: Hard-ground Green's-function calibration diagnostic for the 20 Hz, $h = 2$ m box-domain case. The fitted source strength should be approximately constant if the numerical field behaves like an ideal point source above a perfectly reflecting infinite plane.

Shell (m)	Points	$\lvert Q\rvert$	Scale to 1 Pa at 1 m	Relative fit error
12–20	2318	263.12	0.0478	0.656
15–25	4313	240.41	0.0523	0.703
20–35	11505	213.46	0.0589	0.761
25–45	23670	194.69	0.0645	0.798
30–60	53214	173.41	0.0725	0.844

The fitted source strength is not constant. It decreases from approximately 263 to 173 as the sampling shell moves farther away from the source. The relative fit error is also large, increasing from about 0.66 to 0.84. This means that the finite-domain numerical solution does not match the ideal hard-ground Green's function closely enough to use this fit as a reliable physical source calibration.

There are several reasons for this. First, the numerical source is a Gaussian volume source rather than an exact point source. Second, the computational domain is finite, so the side and top boundaries affect the field. Third, the hard-ground model assumes an infinite flat reflecting plane, while the numerical problem also contains artificial truncation boundaries. The same issue becomes even stronger in the city geometry, where buildings and terrain introduce additional reflections and scattering.

For this reason, the hard-ground Green's-function fit is used only as a diagnostic. The main calibration used in the thesis is instead the free-field Gaussian-source normalization in Table A.2. This gives a consistent reference scaling for the numerical amplitudes, while avoiding the misleading impression that the finite-domain simulations are fully calibrated outdoor sound-pressure-level predictions.

Bibliography

[1] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. Curfman McInnes, R. Tran Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang, *PETSc/TAO Users Manual*. Argonne National Laboratory, 2024. Available: `https://petsc.org/`.

[2] S. C. Brenner and L. R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd ed. New York: Springer, 2008. doi: `10.1007/978-0-387-75934-0`.

[3] V. Dolean, P. Jolivet, and F. Nataf, *An Introduction to Domain Decomposition Methods: Algorithms, Theory, and Parallel Implementation*. Philadelphia: Society for Industrial and Applied Mathematics, 2015. doi: `10.1137/1.9781611974065`.

[4] DOLFINx developers, *DOLFINx documentation, version 0.9.0*. Available: `https://docs.fenicsproject.org/dolfinx/v0.9.0/python/`.

[5] DTCC Platform, *DTCC Sim*. GitHub repository, 2026. Available: `https://github.com/dtcc-platform/dtcc-sim`.

[6] O. G. Ernst and M. J. Gander, "Why it is difficult to solve Helmholtz problems with classical iterative methods," in *Numerical Analysis of Multiscale Problems*, I. G. Graham, T. Y. Hou, O. Lakkis, and R. Scheichl, Eds. Berlin: Springer, 2012, pp. 325–363. doi: `10.1007/978-3-642-22061-6_10`.

[7] M. J. Gander and H. Zhang, "A class of iterative solvers for the Helmholtz equation: factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods," *SIAM Review*, vol. 61, no. 1, pp. 3–76, 2019. doi: `10.1137/16M109781X`.

[8] C. R. Harris et al., "Array programming with NumPy," *Nature*, vol. 585, pp. 357–362, 2020. doi: `10.1038/s41586-020-2649-2`.

[9] F. Ihlenburg, *Finite Element Analysis of Acoustic Scattering*. New York: Springer, 1998. doi: `10.1007/b98828`.

[10] International Organization for Standardization, *ISO 9613-2: Acoustics – Attenuation of sound during propagation outdoors – Part 2: General method of calculation*. Geneva: International Organization for Standardization, 1996.

[11] G. Karypis and V. Kumar, "A fast and high quality multilevel scheme for partitioning irregular graphs," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 359–392, 1998. doi: 10.1137/S1064827595287997.

[12] C. Ma, C. Alber, R. Scheichl, and Y. Zhang, "Two-level restricted additive Schwarz preconditioner based on multiscale spectral generalized FEM for heterogeneous Helmholtz problems," *arXiv preprint arXiv:2409.06533v2*, 2025. doi: 10.48550/arXiv.2409.06533.

[13] A. D. Pierce, *Acoustics: An Introduction to Its Physical Principles and Applications*, 3rd ed. Cham: Springer, 2019. doi: 10.1007/978-3-030-11214-1.

[14] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 3, pp. 856–869, 1986. doi: 10.1137/0907058.

[15] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. Philadelphia: Society for Industrial and Applied Mathematics, 2003. doi: 10.1137/1.9780898718003.

[16] E. M. Salomons, *Computational Atmospheric Acoustics*. Dordrecht: Kluwer Academic Publishers, 2001. doi: 10.1007/978-94-010-0660-6.

[17] P. Virtanen et al., "SciPy 1.0: Fundamental algorithms for scientific computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020. doi: 10.1038/s41592-019-0686-2.

[18] A. Toselli and O. Widlund, *Domain Decomposition Methods – Algorithms and Theory*. Berlin: Springer, 2005. doi: 10.1007/b137868.

[19] World Health Organization Regional Office for Europe, *Environmental Noise Guidelines for the European Region*. Copenhagen: World Health Organization Regional Office for Europe, 2018. Available: https://www.who.int/publications/i/item/9789289053563.

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

AVANCEZ 1829

CHALMERS
UNIVERSITY OF TECHNOLOGY