



Degree Project in Embedded Systems

Second cycle, 30 credits

Source and Sensor Placement Optimisation for Spatial Active Noise Control

FLORENT GRIMAU

Source and Sensor Placement Optimisation for Spatial Active Noise Control

FLORENT GRIMAU

Master's Programme, Embedded Systems, 120 credits

Date: November 10, 2025

Supervisor: Shoichi Koyama

National Institute of Informatics, Tokyo, Audio Processing Research Group

Examiner: Ana Rusu

School of Electrical Engineering and Computer Science

Host organization: National Institute of Informatics, Tokyo

Swedish title: Optimering av källa- och sensorsplacering för spatial aktiv bullerkontroll

Abstract

Spatial **Active noise control (ANC)** aims to manipulate sound fields within target regions through active control methods, with applications including noise cancellation and high-fidelity audio reproduction. The placement of secondary sources and sensors has a significant impact on system performance; however, traditional approaches optimise these placements independently, resulting in suboptimal results. This thesis addresses the joint source and sensor placement optimisation problem in spatial ANC systems. The challenge lies in formulating an effective cost function that simultaneously optimises secondary source placement for sound field synthesis and sensor placement for accurate spatial interpolation, while maintaining reasonable computational complexity. A novel joint cost function formulation was developed that explicitly incorporates interpolation error alongside synthesis error by comparing interpolated field estimates against the true desired field across the entire control zone. This approach addresses fundamental limitations in existing methods, including clustering behaviour and convergence issues, by properly accounting for estimation uncertainty inherent in limited sensor measurements. A comprehensive simulation environment was implemented to evaluate four placement algorithms: Random (baseline), Regular (geometric), Greedy (optimisation-based), and Matching Pursuit (signal processing-based). The algorithms were systematically compared using normalised mean square error and computational efficiency metrics across multiple resolution configurations. Results demonstrate that optimisation-based approaches provide meaningful performance improvements over simple placement strategies, with the Greedy algorithm achieving the best acoustic performance. However, significant computational trade-offs were revealed, with optimisation methods requiring substantially longer execution times. The novel cost function formulation successfully resolved convergence issues and demonstrated superior robustness across varying acoustic conditions.

Keywords

Active noise control, Spatial acoustics, Source placement optimisation, Sensor placement optimisation, Joint optimisation, Pressure matching, Cost function formulation, Greedy algorithms, Sound field synthesis, Acoustic control systems, Weighted pressure matching, Transfer function matrices, Spatial interpolation

Sammanfattning

Rumslig aktiv brusreducering (ANC) syftar till att styra ljudfält inom specifika målområden med hjälp av aktiva kontrollmetoder, med tillämpningar såsom brusreducering och högupplöst ljudåtergivning. Placeringen av sekundära ljudkällor och sensorer har stor inverkan på systemets prestanda, men traditionella metoder optimerar dessa placeringar var för sig, vilket ofta leder till suboptimala resultat. Denna avhandling behandlar en gemensam optimering av källa- och sensorplacering i rumsliga ANC-system. Utmaningen ligger i att formulera en effektiv kostnadsfunktion som samtidigt optimerar placeringen av sekundära källor för ljudfältsyntes och sensorer för noggrann rumslig interpolering, utan att den beräkningsmässiga komplexiteten blir orimlig. En ny formulering av en gemensam kostnadsfunktion har utvecklats, där både interpoleringsfel och syntesfel uttryckligen beaktas genom att interpolerade fältuppskattningar jämförs med det önskade ljudfältet över hela kontrollzonen. Detta angreppssätt hanterar grundläggande begränsningar i befintliga metoder – såsom klustereffekter och konvergensproblem – genom att korrekt beakta den osäkerhet som uppstår vid uppskattning från ett begränsat antal sensormätningar. En omfattande simuleringsmiljö har implementerats för att utvärdera fyra placeringsalgoritmer: slumpmässig (baslinje), regelbunden (geometrisk), greedy (optimeringsbaserad) och Matching Pursuit (signalbehandlingsbaserad). Algoritmerna har systematiskt jämförts utifrån normaliserat medelkvadratisk fel och beräkningseffektivitet över flera upplösningsskonfigurationer. Resultaten visar att optimeringsbaserade metoder ger tydliga prestandaförbättringar jämfört med enklare placeringsstrategier, där greedy-algoritmen uppnådde bäst akustisk prestanda. Dock identifierades tydliga beräkningsmässiga kompromisser, eftersom optimeringsmetoderna krävde avsevärt längre exekveringstid. Den nya formuleringen av kostnadsfunktionen löste tidigare konvergensproblem och visade överlägsen robusthet under varierande akustiska förhållanden.

Nyckelord

Aktiv brusreducering, Rumslig akustik, Optimering av källa placering, Optimering av sensorplacering, Gemensam optimering, Tryckmatchning, Formulering av kostnadsfunktion, Greedy-algoritmer, Ljudfältsyntes, Akustiska kontrollsystem, Viktad tryckmatchning, Överföringsfunktionsmatriser, Rumslig interpolering

Acknowledgments

I would like to express my gratitude to my examiner, Prof. Ana Rusu, for her valuable guidance and feedback throughout the development of this degree project. Her insights helped shape and strengthen this work.

I am deeply grateful to Pr. Koyama Shoichi from the National Institute of Informatics for supervising my work and making the internship, which formed the foundation of this thesis, possible. His expertise and thoughtful reflections throughout my time at NII created an ideal research environment. I especially appreciate how they facilitated my integration into the team and their welcoming approach during my time in Japan.

I would like to thank the National Institute of Informatics for providing me with the opportunity to conduct my research at their institute in central Tokyo. The convenient location, combined with regular seminars and talks by researchers from around the world, created a stimulating academic environment that greatly enriched my experience.

Special thanks to Claire and Gwendal for their friendship during my time in Japan. They helped me discover the country, shared their knowledge of places and culture, and made exploring together both fun and memorable. Their companionship provided welcome breaks from research and made Japan feel much less foreign.

To my parents, I am profoundly grateful for enabling my studies in Sweden through their financial support and unwavering belief in education. They have consistently encouraged me to pursue opportunities and challenged me to strive for more, making this thesis and this degree possible.

Finally, I would like to thank Emi, who has been an invaluable support for me both emotionally and with this work. Thank you for your help with the more mathematical aspects of this project, thank you for the regular calls while I was away and thank you for your continuous support over the last ten years.

Stockholm, November 2025

Florent Grimau

Contents

1	Introduction	1
1.1	Background	1
1.2	Problem	1
1.2.1	Original problem and definition	2
1.2.2	Scientific and engineering issues	3
1.3	Purpose	3
1.4	Goals	3
1.5	Research Methodology	4
1.6	Delimitations	5
1.7	Structure of the thesis	6
2	Background	7
2.1	Active Noise Control: Historical Development and Spatial Extension	7
2.2	From Single-Point to Spatial Control	8
2.2.1	Mathematical Foundations for Spatial Control	9
2.3	Pressure Matching and the Emergence of Placement optimisation	10
2.3.1	Early Pressure Matching Formulations	10
2.3.2	Weighted Pressure Matching	10
2.4	Mathematical Framework for Placement optimisation	11
2.4.1	Selection Matrix Formulation	11
2.4.2	Transfer Function Matrices and Spatial Coupling	11
2.5	Cost Function Evolution and Modern Formulations	12
2.5.1	Spatial Covariance and Prior Information	12
2.5.2	Joint optimisation Frameworks	13
2.5.3	Performance Metrics and Evaluation	13
2.6	Summary and Connection to Thesis Contributions	14

3	Methods	15
3.1	Research Process	15
3.1.1	Simulation Framework and Acoustic Modelling	15
3.1.2	Problem Formulation Framework	16
3.1.3	Candidate Position Definition and Constraints	17
3.2	Research Paradigm	18
3.2.1	Design Science Approach	18
3.2.2	Experimental Methodology	18
3.2.3	Research Philosophy	19
3.3	Simulation Design	19
3.3.1	Simulation Environment Setup	19
3.3.1.1	Room Configuration and Acoustic Properties	19
3.3.1.2	Spatial Discretisation Strategy	20
3.3.1.3	Frequency Domain Analysis	20
3.4	Simulation Model Validation	21
3.4.1	Greedy Algorithm Convergence Monitoring	21
3.4.2	Spatial Distribution Analysis	22
3.4.3	Deterministic Reproducibility	22
3.5	Performance Analysis Methods	23
3.5.1	Primary Performance Metrics	23
3.5.2	Multi-Angle Robustness Assessment	23
3.5.3	Baseline Establishment	23
3.5.4	Visual Field Analysis	25
4	Implementation and Algorithm Development	27
4.1	Joint Cost Function Development	27
4.1.1	Motivation and Problem Context	27
4.1.2	Parameter Definitions and Notation	27
4.1.3	Acoustic Transfer Function Formulation	28
4.1.4	Weighting Matrix Formulation	29
4.1.5	Initial Cost Function Development	30
4.1.6	Implementation and Observed Limitations	32
4.1.7	Novel Joint Cost Function Formulation	34
4.1.8	Mathematical Development	35
4.1.9	Formulation Properties	35
4.2	Software Architecture and Implementation	36
4.3	Placement optimisation Algorithms	37
4.3.1	Random Placement Baseline	37
4.3.1.1	Complexity Analysis	38

4.3.2	Regular Placement Algorithm	39
4.3.3	Greedy Pair-Selection Algorithm	41
4.3.4	Matching Pursuit Algorithm	43
5	Results and Analysis	47
5.1	Major Results	47
5.1.1	Default Configuration Performance	47
5.1.2	Cross-Configuration Analysis	48
5.1.3	Runtime Efficiency Analysis	49
5.1.4	Algorithm Convergence Behavior	49
5.1.5	Performance visualisation Results	50
5.2	Baseline Performance Validation	51
5.3	Algorithm Efficiency Analysis	52
5.4	Discussion	52
5.4.1	Algorithm Performance	52
5.4.2	Scalability and Resolution Dependencies	52
5.4.3	Limitation to narrowband	55
5.4.4	Practical Implementation Considerations	56
6	Conclusions and Future Work	57
6.1	Conclusions	57
6.1.1	Achievement of Research Goals	57
6.1.2	Key Research Insights	58
6.2	Limitations	58
6.2.1	Methodological Limitations	58
6.2.2	Algorithmic Limitations	59
6.2.3	Technical constraints	59
6.3	Future Work	59
6.3.1	Immediate Extensions	59
6.3.2	Algorithmic Development	60
6.4	Reflections	60
6.4.1	Research Impact and Broader Implications	60
6.4.2	Ethical and Environmental Considerations	60
6.4.3	Personal Reflections	61
	References	63
A	Software architecture details	69
A.1	Software Architecture Implementation	69
A.1.1	System Architecture	69

A.1.2	Data Flow and Execution	70
A.1.3	Caching Strategy	71
A.1.4	External Integration	71
A.1.5	Performance and Reproducibility	71
B	Simulation configurations	72
C	Full Sized Figures	76
C.1	Algorithm plots	76
C.1.1	Greedy Algorithm Plots	76
C.1.2	Regular Algorithm Plots	80
C.1.3	Matching Pursuit Plots	83

List of Figures

1.1	A typical setup for spatial ANC.	2
2.1	Reproduction of figure 1 from Paul Lueg’s 1936 patent [5], showing means of building destructive inference using a microphone and loudspeaker within a metallic tube.	8
3.1	Plot of the mirror image method. The green room is our main room, the purple rooms around are mirrors of our original room, and secondary sound paths are computed from the mirror sources in those rooms.	16
3.2	Candidate positions for sources and sensors in the simulation setup. The black outer boundary represents the room walls, secondary source candidates are shown around the perimeter, and sensor candidates are distributed within the control region.	17
3.3	Comparison of the different discretisation steps with a wavelength of a 500Hz sine wave, where each circle represents a control point.	21
3.4	Clustering examples	22
3.5	Plot of the evaluation angles where X is the angular discretisation step, y is the evaluation angular range, the bold arrow is the initial plane wave direction, and the other arrows represent the angle at which the NMSE is computed. For legibility, not all arrows are shown on the figure, but the dotted arrow is repeated every x angles, spanning a range of $2 \times y$	24
4.1	Undesirable clustering behaviour.	33
4.2	Cost yielded by the cost function over the amount of selected sources. This results from the greedy algorithm (See sec. 4.3.3) selecting 20 pairs of sources and sensors iteratively.	34

4.3	Cost history for the greedy algorithm (See Sec. 4.3.3). The cost function exhibits a monotonic decrease with the number of source/sensor pairs selected.	36
5.1	Normalised Mean-Square Error (NMSE) evaluation comparison for the four algorithms.	48
5.2	Algorithm performance visualisation for Greedy algorithm in default configuration (Larger figures can be found in appendix C).	51
5.3	Algorithm performance visualisation for the Regular algorithm in default configuration	53
5.4	Algorithm performance visualisation for Matching Pursuit algorithm in default configuration	54
A.1	Software architecture diagram	70
C.1	Greedy algorithm convergence in default configuration. 5.2a .	76
C.2	optimised source and sensor positions (Greedy, default). 5.2b .	77
C.3	Acoustic field visualisation (Greedy, default). 5.2c	78
C.4	NMSE evaluation (Greedy, default). 5.2d	79
C.5	optimised source and sensor positions (Regular, default). 5.3a .	80
C.6	Acoustic field visualisation (Regular, default). 5.3b	81
C.7	NMSE evaluation (Regular, default). 5.3c	82
C.8	optimised source and sensor positions (Matching Pursuit, default). 5.4a	83
C.9	Acoustic field visualisation (Matching Pursuit, default). 5.4b .	84
C.10	NMSE evaluation (Matching Pursuit, default). 5.4c	85

List of Tables

5.1	Default Configuration Performance Results	47
5.2	Cost Performance Across Configuration Profiles	49
5.3	Runtime Performance Across Configuration Profiles	49

List of acronyms and abbreviations

ANC	Active noise control
AR	Augmented Reality
EIM	Empirical Interpolation Method
MP	Matching Pursuit
NMSE	Normalised Mean-Square Error
RKHS	Reproducing Kernel Hilbert Space
RT60	Reverberation Time 60
SDG	Sustainable Development Goal
SDR	Signal-to-Distortion ratio
SFC	Sound Field Control
SSPOC	Sparse sensor placement optimisation for classification
UN	United Nations
VR	Virtual Reality

Chapter 1

Introduction

In this thesis, two cost function formulations are explored and developed for determining the optimal placement of secondary sources and sensors in the context of spatial **Active noise control (ANC)**.

1.1 Background

Spatial **ANC** aims to manipulate and reproduce a spatial sound field within target regions through active control methods. Its application range is wide and contains, among others, high-fidelity reproduction of recorded sound fields for **Augmented Reality (AR)/Virtual Reality (VR)** applications, **Sound Field Control (SFC)**, active noise cancelling projects [1, 2] or personal audio zones delivering content to specific listeners [3]. Spatial **ANC** uses an array of microphones (sensors) and loudspeakers (secondary sources) to estimate the continuous sound field within a region and actively control or reproduce desired acoustic conditions, respectively. Fig. 1.1 displays a typical setup for spatial **ANC**. The figure contains a primary noise source describing the pre-existing noise in the room, a control region where the sound fields need to be controlled, a set of candidate microphones within or near the control zone and a set of candidate "secondary" loudspeakers.

1.2 Problem

In spatial **ANC**, the placement of microphones and loudspeakers significantly impacts system performance [4]. Optimal microphone placement enables better sound field estimation within the control region, potentially achieving

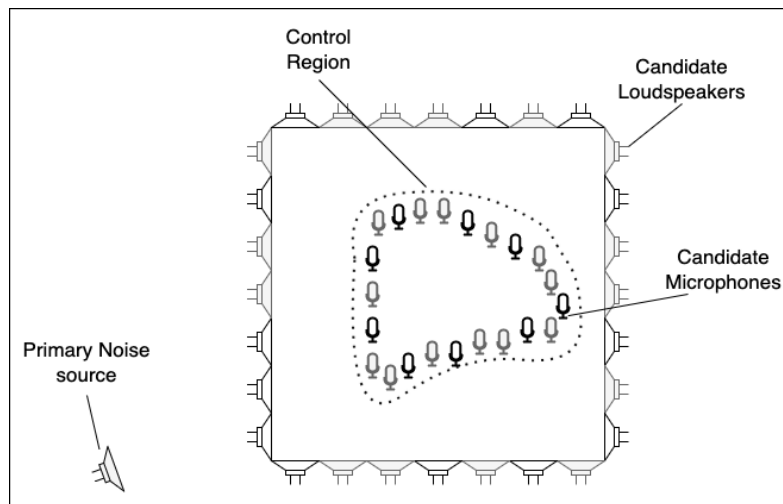


Figure 1.1: A typical setup for spatial ANC.

similar results with fewer sensors. Similarly, optimised loudspeaker placement improves sound field control capabilities. However, achieving such optimisation requires a deep understanding of the underlying estimation and reproduction algorithms.

1.2.1 Original problem and definition

The core research problem addressed in this thesis is: "How can we formulate an effective cost function for optimising the placement of loudspeakers and microphones in the context of spatial ANC, with reasonable computational complexity?"

This problem originates from a requirement for physical implementations to contain as few microphones and sensors as possible. This would not only reduce the cost of the installation but also the complexity. The scope includes:

- Implementation of a development environment with existing spatial ANC algorithms.
- Development of an algorithm for joint microphone and loudspeaker positioning.

The problem excludes hardware design and specific ANC control algorithms.

1.2.2 Scientific and engineering issues

Several scientific and engineering challenges make this optimisation problem particularly complex:

- Multi-objective optimisation: Balancing performance metrics (noise reduction, spatial coverage) against practical constraints (computational cost, physical placement limitations)
- High-dimensional search space: The number of possible placement configurations grows exponentially with the number of transducers
- Frequency-dependent behaviour: Optimal placement may vary across frequency ranges

1.3 Purpose

The purpose of this thesis is to advance the state-of-the-art in spatial ANC by enabling more efficient and effective placement of acoustic sources and sensors, thereby improving noise control in real-world applications while reducing system complexity and cost.

From a sustainability perspective, the optimisation approach directly reduces material consumption by requiring fewer transducers and lowers computational power requirements in deployed systems.

However, several considerations must be addressed: the benefits may initially favour high-end applications, potentially affecting accessibility to advanced noise control solutions.

1.4 Goals

The goal of this project is to develop and validate an effective cost function formulation for source and sensor placement in spatial acoustic control systems that enable better optimisation performance in terms of control accuracy, reproduction fidelity, and computational efficiency across diverse application scenarios.

Several sub-goals have been defined as follows:

1. Literature study: establish mathematical foundations for pressure matching-based spatial ANC and conduct a review of existing placement optimisation algorithms.

2. Environment set-up: Implement a state-of-the-art algorithm for spatial **ANC** to set up an environment where a placement algorithm can be developed.
3. Formulation: Formulate a relevant cost function for joint source and sensor placement, allowing the implementation of a minimisation algorithm.
4. Simulation and validation: Run the algorithm in a two-dimensional acoustic environment reproducing real-life conditions for evaluation and validation against other existing methods and in multiple environments (reverberating room vs. free-field, ...).

1.5 Research Methodology

This project will be conducted on an empirical basis. The assumption is that optimal source and sensor placement can be determined through mathematical optimisation of quantifiable performance metrics. Such metrics are obtained through a simulation with a strong physical background governed by well-established laws (wave equations, Sabine's formula for reverberation time).

The primary research method of this project is simulation-based. This involves developing a custom simulation environment and involves mathematical modelling (cost function formulation), algorithm implementation (for spatial **ANC** methods) and systematic evaluation across various scenarios with quantitative, deterministic metrics. This method was chosen for its relative simplicity to set up, high versatility for acoustic environment selection and low hardware cost.

This research follows a quantitative experimental approach. The simulation environment serves as a black box with controlled inputs and outputs. The inputs consist of the shape of the room, the placement of sources and sensors. The outputs are limited to standardised performance metrics (Cost function value and **Normalised Mean-Square Error (NMSE)**). The placement algorithm has a minimal coupling with the simulation details, yielding only the physical placement of the transducers and reading the resulting metrics.

Alternative considered methods include:

- Physical Experimental Methodology; this method was rejected since it would require expensive microphone arrays and loudspeaker systems,

specialised anechoic or reverberant chambers and would limit the possible room configurations.

- Purely Theoretical/Analytical approach; While mathematical analysis could provide better results since the discretisation step could be avoided, the complexity of acoustic environments would make the closed-form solutions intractable. Here, simulation helps bridge the gap between theory and practice by leveraging the most computationally intensive parts of the project.

1.6 Delimitations

This thesis is constrained to focus on the development and validation of the joint source and placement algorithms within a computational simulation framework. The research will operate exclusively within a Python-based acoustic simulation environment, explicitly excluding physical experimental validation. While those simulations would help enhance the quality of the final placement, they would require specialised equipment and controlled acoustic environments exceeding the scope and available resources for the project.

The acoustic analysis and simulation will be mainly conducted in a two-dimensional environment. This approach excludes computationally heavy geometric simulations such as furniture modelling and surface irregularities. The two-dimensional focus provides sufficient algorithmic complexity while enabling clear visualisation of placement results, reproduced and estimated sound fields. Although an extension to a simplified three-dimensional environment would have been interesting, it could not be implemented due to time constraints. In a similar way, the simulation will mainly be conducted for single-frequency sound field reproduction.

The optimisation framework operated on pre-determined candidate positions for both secondary sources and sensors, excluding a continuous approach for the placement of transducers. This discrete approach aligns with real-world constraints of placement, where there is often the need for a specific mount point in order to place microphones or loudspeakers. These sources and sensors are idealised points in space, deliberately avoiding the complexity of modelling the characteristics of a real loudspeaker or microphone.

There are several external constraints on the project, such as the computational cost of simulating a complete room. In order to address this problem, the discretisation step can be rather large, effectively limiting the precision of the simulations. Another constraint is in the use of simplified

acoustic models, limiting the real-world accuracy of, for example, the sound reverberation in a closed room.

1.7 Structure of the thesis

Chapter 2 presents relevant background information about spatial ANC. Chapter 3 presents the methodology and methods used to develop the optimisation algorithms. Chapter 4 presents the developed simulation environment and its parameters. Chapter 5 presents and analyses the benchmark results for all tested algorithms, evaluates their performance trade-offs, and examines the technical challenges encountered throughout the project development. Finally, chapter 6 contains the conclusion and ideas for future work in the field.

Chapter 2

Background

This chapter establishes the mathematical foundations for pressure-matching-based spatial ANC and reviews existing placement optimisation algorithms. The chapter traces the historical development of spatial ANC from Lueg's foundational 1936 patent ([5]) through modern multichannel systems, with particular focus on the evolution from single-point control to spatial pressure matching methods. Sections 2.1 to 2.3 develop the mathematical framework underlying pressure matching formulations, including transfer function matrices, spatial coupling effects, and the emergence of weighted pressure matching. Sections 2.4 and 2.5 provide a systematic review of placement optimisation algorithms, covering experimental design methods, information-theoretic* approaches, and joint source-sensor optimisation frameworks. The chapter concludes by identifying key research gaps that motivate the novel joint cost function and optimisation algorithms developed in subsequent chapters.

2.1 Active Noise Control: Historical Development and Spatial Extension

ANC finds its roots in Paul Lueg's foundational patent published in 1936 [5]. In this patent, Lueg proposed a revolutionary approach to noise reduction by using destructive interference rather than passive absorption methods. His concept involved capturing unwanted sound oscillations with a microphone and reproducing them through a loudspeaker with opposite phase, as shown in figure 2.1, effectively cancelling the original noise through

*Related to information theory

acoustic interference. This marked the first documented attempt to actively control sound using electronic means rather than purely mechanical or passive approaches.




Figure 2.1: Reproduction of figure 1 from Paul Lueg's 1936 patent [5], showing means of building destructive interference using a microphone and loudspeaker within a metallic tube.

The core principle described by Lueg remains fundamental to modern acoustic control systems: when two sound waves of equal amplitude but opposite phase meet, they undergo destructive interference, resulting in reduced overall sound pressure. Lueg's patent specifically addressed the challenge of positioning the microphone between the noise source and the reproducing apparatus to ensure proper phase relationships for effective cancellation. While technologically limited by 1930s electronics, his work established the theoretical foundation for what would become a major field of acoustic engineering encompassing both noise control and sound field reproduction.

2.2 From Single-Point to Spatial Control

The practical implementation of ANC systems remained dormant for several decades until advances in electronics enabled the development of the filtered-x least mean squares (FxLMS) algorithm in the 1980s [6, 7]. Early systems focused on single-point control, creating localised quiet zones with spatial extent fundamentally limited by the wavelength-to-ten rule, restricting effective noise reduction to regions smaller than $\lambda/10$, with λ the wavelength

[8, 9]. This limitation became increasingly problematic as applications demanded control over extended spatial regions.

The transition to multichannel spatial ANC emerged in the 1990s through the pioneering work of researchers who recognised that multiple secondary sources and error sensors could overcome single-point limitations [10, 11, 12, 13]. This evolution introduced the fundamental placement optimisation problem: given finite resources, where should secondary sources and error sensors be positioned to achieve optimal spatial control? The coupling between control points introduced complex interdependencies that could not be addressed through intuitive geometric arrangements alone.

2.2.1 Mathematical Foundations for Spatial Control

The mathematical framework for spatial ANC was developed alongside advances in acoustic theory and sound field synthesis. Williams' seminal work on Fourier acoustics* [14] established theoretical foundations for representing and analysing sound fields using Fourier methods.

Building on these acoustic principles, spatial ANC systems leverage the superposition principle, where the total acoustic field at any point results from the linear combination of individual source contributions. This principle enables the strategic placement of secondary sources to destructively interfere with primary noise sources, forming the theoretical basis for multichannel ANC approaches [10].

In practical implementations, the continuous distribution of secondary sources must be approximated by a finite number of discrete loudspeakers. The synthesised sound field is then expressed, in the frequency domain, as a linear superposition of the individual loudspeaker contributions:

$$u_{\text{syn}}(\mathbf{r}) = \sum_{l=1}^L d_l g_l(\mathbf{r}) \quad (2.1)$$

where d_l represents the driving signal for the l -th loudspeaker, $g_l(\mathbf{r})$ is the acoustic transfer function from the l -th loudspeaker to position \mathbf{r} , and L is the total number of loudspeakers. This discrete formulation directly implements the superposition principle that underlies all sound field synthesis approaches.

*Fourier acoustics applies Fourier analysis to solve acoustic wave equations, enabling the decomposition of complex sound fields into harmonic components. This mathematical framework forms the foundation for modern acoustic holography and sound field reconstruction techniques.

2.3 Pressure Matching and the Emergence of Placement optimisation

2.3.1 Early Pressure Matching Formulations

The development of pressure-matching methods emerged from the practical need to discretise continuous spatial control problems. Early researchers recognised that the optimal control problem seeking to minimise regional error (Eq. 2.2) was computationally intractable for practical implementation [10].

$$\min Q(\mathbf{d}) = \int_{\Omega} |u_{\text{syn}} - u_{\text{des}}|^2 d\mathbf{r} \quad (2.2)$$

where $u_{\text{syn}}(\mathbf{r})$ and $u_{\text{des}}(\mathbf{r})$ represent the synthesised and desired sound fields, respectively, at position \mathbf{r} .

The pressure matching approach provided an elegant solution by discretising the control region into M control points $\{\mathbf{r}_m\}_{m=1}^M$, transforming the continuous optimisation into a tractable matrix formulation [15]:

$$\min_{\mathbf{d}} \|\mathbf{G}\mathbf{d} - \mathbf{u}_{\text{des}}\|^2 \quad (2.3)$$

where $\mathbf{G} \in \mathbb{C}^{M \times L}$ contains the transfer functions between all source-control point pairs. This discretisation immediately raised the placement question: which control points should be selected from a larger set of candidates to achieve optimal performance?

2.3.2 Weighted Pressure Matching

The evolution toward weighted pressure matching marked a significant advancement in spatial ANC theory. Researchers recognised that uniform weighting of control points was often suboptimal, leading to weighted formulations [16, 17]:

$$\min_{\mathbf{d}} (\mathbf{G}\mathbf{d} - \mathbf{u}_{\text{des}})^H \mathbf{W} (\mathbf{G}\mathbf{d} - \mathbf{u}_{\text{des}}) \quad (2.4)$$

where \mathbf{W} is a positive definite weighting matrix. This seemingly simple modification had profound implications for placement optimisation, as it enabled emphasis of control performance in specific spatial regions while relaxing requirements elsewhere.

Recent advances in weighted pressure matching include kernel interpolation techniques that improve spatial accuracy while reducing computational

complexity [17]. These developments provide the mathematical foundation for the advanced cost function formulations explored in this thesis. Section 4.1.4 explores the **Reproducing Kernel Hilbert Space (RKHS)** theory and its application in weighted pressure matching.

2.4 Mathematical Framework for Placement optimisation

2.4.1 Selection Matrix Formulation

The formal mathematical framework for placement optimisation emerged through the introduction of selection matrices [18]. The placement problem operates over finite candidate sets $P = \{\mathbf{r}_1, \dots, \mathbf{r}_p\}$ for sources and $Q = \{\mathbf{r}'_1, \dots, \mathbf{r}'_q\}$ for sensors, where positions are constrained by practical limitations including room geometry, accessibility, and hardware constraints [19].

Binary selection matrices $\Phi_L \in \{0, 1\}^{p \times l}$, where l represents the amount of sources selected and $\Phi_M \in \{0, 1\}^{q \times m}$, where m represents the amount of sensors selected, represent source and sensor selections, respectively. The selected transfer function matrix becomes:

$$\mathbf{G}_{\text{sel}} = \Phi_M^T \mathbf{G}_{\text{full}} \Phi_L \quad (2.5)$$

where $\mathbf{G}_{\text{full}} \in \mathbb{C}^{p \times q}$ contains transfer functions between all candidate positions. This framework enables systematic formulation of placement optimisation as a discrete combinatorial problem while maintaining a connection to the underlying acoustic physics. The combinatorial nature of the selection problem has led to the development of specialised algorithms that exploit the structure of acoustic transfer function matrices [17].

2.4.2 Transfer Function Matrices and Spatial Coupling

The acoustic transfer function matrix serves as the fundamental link between physical placement and system performance [20, 14]. Each element $G_{ij} = G(\mathbf{r}_i | \mathbf{r}_j)$ represents the complex acoustic coupling between source position $i \in P$ and control point $j \in Q$, incorporating amplitude attenuation, phase relationships, and frequency-dependent effects [21].

The spatial structure of transfer function matrices reveals important properties for placement optimisation [10]. Well-separated sources and

sensors tend to produce matrices with lower condition numbers*, while clustered arrangements often lead to numerical difficulties [18]. The frequency dependence of acoustic coupling means that optimal placements may vary across the frequency spectrum, requiring broadband optimisation strategies.

2.5 Cost Function Evolution and Modern Formulations

2.5.1 Spatial Covariance and Prior Information

Modern placement optimisation incorporates statistical frameworks that treat desired sound fields as random variables with known covariance structures [22]. For plane-wave fields with directional constraints, the spatial covariance matrix elements are given by:

$$\Sigma_{j,j'} \in \mathbb{C}^N = \int_{\theta_1}^{\theta_2} e^{jk(\theta) \cdot (\mathbf{r}_j - \mathbf{r}_{j'})} d\theta \quad (2.6)$$

Where the integration limits $[\theta_1, \theta_2]$ define the angular sector of interest, and N is the number of control points within the control zone. This framework enables optimisation for specific acoustic scenarios, including directional sound reproduction and spatially localised control zones. The mathematical foundations for spatial covariance modelling in acoustic applications draw from classical Fourier acoustics and the Jacobi-Anger expansion for plane-wave decomposition [20].

The incorporation of spatial covariance represents a significant advancement over deterministic optimisation approaches, as it enables robust placement design that performs well across multiple realisations of the desired sound field [23, 18]. This statistical perspective forms the foundation for the advanced cost function formulations developed in this thesis and connects to broader work in statistical signal processing for acoustic arrays [24].

Using prior information on the sound field with this statistical framework leads to the expected cost function:

$$\mathbb{E}[F(\hat{\mathbf{d}})] = \text{trace}(\mathbf{D}\Sigma) + \boldsymbol{\mu}^H \mathbf{D}\boldsymbol{\mu} \quad (2.7)$$

*The condition number quantifies numerical stability: low values indicate well-conditioned matrices with stable solutions, while high values suggest ill-conditioning where small input changes cause large solution errors.

Where Σ and μ represent the prior covariance matrix and mean of the desired field, and \mathbf{D} depends on the placement configuration [22]. Recent advances have incorporated prior information about sound field characteristics into optimisation frameworks [25], enabling more robust placement design approaches [23].

2.5.2 Joint optimisation Frameworks

The recognition that optimal sensor placement depends on source configuration, and vice versa, led to joint optimisation frameworks that address these coupling effects simultaneously [26]. Joint source and sensor optimisation significantly increases combinatorial complexity, with placement decisions involving $\binom{p}{l} \times \binom{q}{m}$ potential configurations [27].

Modern approaches employ alternating optimisation strategies and develop convergence guarantees for practical implementation. The comparison with methods such as the **Empirical Interpolation Method (EIM)** [28, 26] reveals trade-offs between computational efficiency and solution quality across different acoustic scenarios [18, 29]. **EIM** is similar to the **Matching Pursuit (MP)** algorithm explored in subsection 4.3.4. Recent work has also explored the integration of joint optimisation with advanced control methods and robust design principles [30].

2.5.3 Performance Metrics and Evaluation

The systematic evaluation of placement optimisation algorithms requires comprehensive performance metrics that capture both spatial accuracy and practical implementation constraints [6]. The Normalised Mean Square Error (NMSE) serves as the primary spatial accuracy measure, defined as:

$$\text{NMSE} = 10 \log_{10} \left(\frac{|\mathbf{u}_{\text{des}} - \mathbf{u}_{\text{syn}}|^2}{|\mathbf{u}_{\text{des}}|^2} \right) \quad (2.8)$$

where \mathbf{u}_{des} represents the desired sound field in the control zone and \mathbf{u}_{syn} represents the synthesised sound field within the control zone [23]. This frequency-dependent metric enables detailed analysis of reproduction accuracy across the operational bandwidth.

The assessment of robustness under uncertainty has emerged as a critical evaluation component. Statistical performance frameworks evaluate placement algorithm stability against positional errors, model uncertainties, and environmental variations [31]. These robustness metrics ensure that optimised

placements maintain acceptable performance under practical implementation conditions, addressing the gap between theoretical optimisation and real-world deployment.

2.6 Summary and Connection to Thesis Contributions

This chapter has traced the evolution of spatial ANC from Lueg's original single-point concept through modern placement optimisation frameworks. Key developments include:

- **Historical foundation:** From Lueg's 1936 patent through the development of multichannel spatial ANC systems
- **Pressure matching evolution:** The transition from continuous formulations to discrete matrix representations
- **Mathematical framework:** Selection matrices, transfer function analysis, and spatial covariance incorporation
- **Cost function development:** From deterministic least-squares to statistical frameworks with prior information

The progression reveals how practical placement optimisation emerged as a critical component of spatial ANC system design, with mathematical formulations evolving to address increasingly sophisticated design requirements.

The mathematical foundations established in this chapter provide the necessary background for the advanced cost function formulations and optimisation algorithms developed in the subsequent chapters of this thesis. The emphasis on statistical frameworks and prior information incorporation directly motivates the novel approaches to placement optimisation that constitute the primary contributions of this work.

Chapter 3

Methods

The purpose of this chapter is to provide an overview of the research methodology used in this thesis. Section 3.1 describes the simulation framework and acoustic modelling approach, including the problem formulation framework and candidate position constraints used to investigate joint source and sensor placement optimisation. Section 3.2 details the design science research paradigm that guides the development and validation of the novel cost function, including the experimental methodology and research philosophy adopted for this computational study. Section 3.3 covers the simulation environment setup, including room configuration, spatial discretisation strategies, and frequency domain analysis parameters used throughout the research. Section 3.4 explains the validation techniques used to verify algorithm behaviour and integration testing, including convergence monitoring, spatial distribution analysis, and reproducibility assessment. Finally, section 3.5 describes the performance metrics and statistical analysis methods used for algorithm comparison, including NMSE computation, multi-angle robustness assessment, baseline establishment, and visual field analysis techniques.

3.1 Research Process

3.1.1 Simulation Framework and Acoustic Modelling

This thesis investigates joint source and sensor placement optimisation through computational simulation using Python. The simulation environment leverages several key libraries: NumPy and SciPy for numerical computations, Matplotlib for visualisation, and Pyroomacoustics for acoustic modelling and

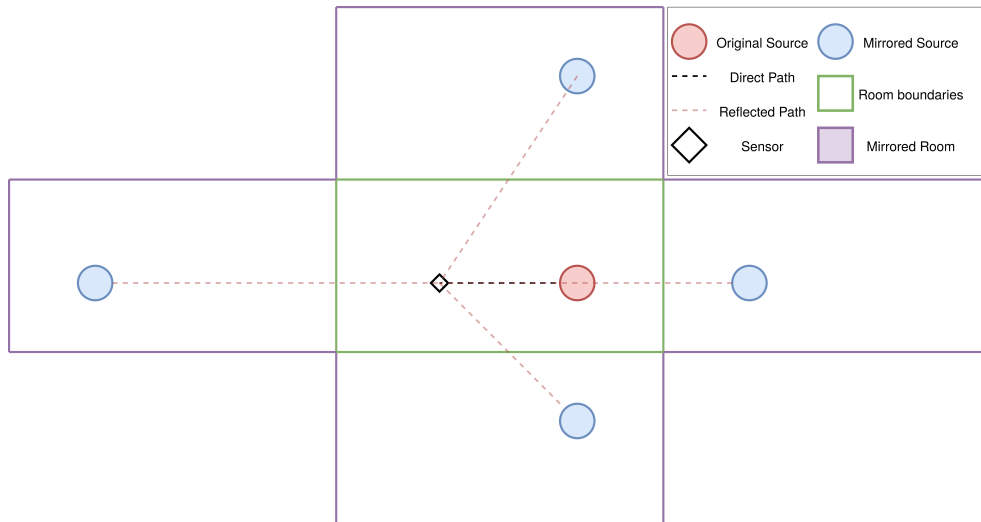


Figure 3.1: Plot of the mirror image method. The green room is our main room, the purple rooms around are mirrors of our original room, and secondary sound paths are computed from the mirror sources in those rooms.

room simulation.

The acoustic modelling employs Green’s functions with simulated room reflections using the image source method implemented in Pyroomacoustics. This approach models the reverberant environment by computing the sum of direct path and reflected path contributions for each source-receiver pair, as shown in figure 3.1. The transfer function matrix \mathbf{G} (Section 4.1.3) represents the acoustic coupling between all candidate source positions and control points, where each element g_{ij} describes the transfer function from the j -th candidate source to the i -th control point, incorporating both direct and reflected propagation paths.

3.1.2 Problem Formulation Framework

The research addresses the joint optimisation problem of selecting optimal positions for both secondary sources and sensors from predefined candidate sets. The objective is to formulate and compare cost functions that balance sound field reproduction fidelity with estimation accuracy from sensors.

Given the computational constraints and the focus on algorithm comparison, the analysis is restricted to narrowband operation at a single frequency. The desired sound field is represented by a plane wave propagating in a specified direction, providing a controlled and well-defined target for reproduction assessment.

3.1.3 Candidate Position Definition and Constraints

To manage computational complexity while maintaining realistic constraints, both secondary sources and sensors are restricted to predefined candidate positions arranged on discretised grids [32].

For sensors, the candidate positions \mathcal{S} form a subset of the discretised control points \mathcal{K} within the target region Ω , such that $\mathcal{S} \subset \mathcal{K}$. This subset constraint ensures that the central area of the control region remains free of sensors, which is essential for practical applications where users need unobstructed access to the controlled zone. The choice to remove those candidate positions, whilst having a negative impact on the interpolation quality, lowers the selection algorithm complexity by removing a big portion of candidate sensors. How much the interpolation accuracy is affected has not been investigated in this work. Figure 3.2 illustrates a typical candidate position layout used in the simulations.

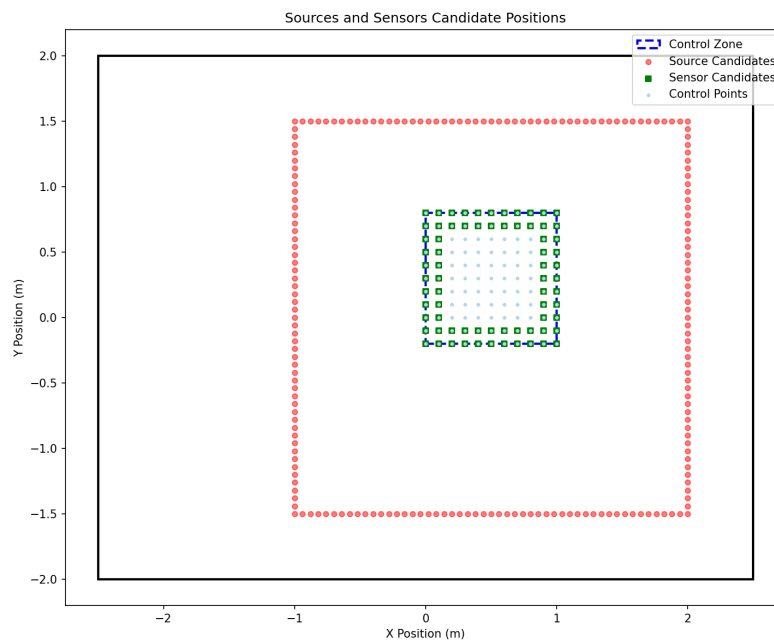


Figure 3.2: Candidate positions for sources and sensors in the simulation setup. The black outer boundary represents the room walls, secondary source candidates are shown around the perimeter, and sensor candidates are distributed within the control region.

Secondary source candidates are positioned around the boundary of the simulated room, following practical constraints for loudspeaker placement in real environments.

3.2 Research Paradigm

This thesis follows a **design science research methodology** combined with **experimental validation** to address the joint source and sensor placement problem in spatial ANC. The research paradigm integrates theoretical formulation with computational experimentation to develop and validate a novel optimisation approach.

3.2.1 Design Science Approach

The core research follows a systematic design process consisting of five phases:

1. **Formulate:** Development of a novel cost function that simultaneously incorporates sound field synthesis capabilities (through secondary source placement) and estimation accuracy (through sensor placement).
2. **Implement:** Computational implementation of the proposed cost function within the Python simulation framework, including the development of optimisation algorithms for joint placement selection.
3. **Test:** Verification of the cost function behaviour and optimisation algorithms through controlled simulation scenarios using the established simulation framework.
4. **Evaluate:** Performance assessment of placements generated by the proposed cost function using standardised metrics including **NMSE** and **Signal-to-Distortion ratio (SDR)**.
5. **Compare:** Benchmarking the proposed approach against existing placement methods to validate performance improvements and establish the contribution of the new formulation.

3.2.2 Experimental Methodology

The research employs a **controlled experimental approach** using computational simulation to test the research hypothesis. The experimental design follows established principles for algorithm evaluation in acoustic signal processing:

Research Hypothesis: The proposed joint cost function, which simultaneously optimises synthesis and approximation capabilities, will outperform

existing placement optimisation methods in terms of sound field reproduction accuracy and noise reduction performance.

Experimental Variables:

- *Independent Variable:* Placement optimisation algorithm (proposed vs. existing methods)
- *Dependent Variables:* **NMSE** and SDR of the resulting **ANC** system
- *Controlled Variables:* Room characteristics, candidate position grids, number of sources and sensors, frequency of operation

Validation Approach: The research employs *simulation-based validation* where the acoustic environment, source characteristics, and measurement conditions are precisely controlled. This computational approach enables systematic comparison across multiple placement algorithms while eliminating experimental uncertainties that would arise in physical implementations.

3.2.3 Research Philosophy

This work adopts a **pragmatic research philosophy**, focusing on developing practical solutions to real-world engineering problems. The research is driven by the practical need for improved spatial ANC systems rather than purely theoretical exploration. The validation through simulation serves as a necessary step toward eventual real-world implementation, following established practices in acoustic engineering research where computational validation precedes experimental verification.

The approach emphasises **quantitative evaluation** using well-established performance metrics from the acoustic signal processing community, ensuring reproducibility and comparability with existing literature. The research contributes both a novel theoretical formulation and empirical evidence of its effectiveness, bridging the gap between mathematical optimisation theory and practical acoustic engineering applications.

3.3 Simulation Design

3.3.1 Simulation Environment Setup

3.3.1.1 Room Configuration and Acoustic Properties

The simulation employs a standardised rectangular room with dimensions of 5.0m × 4.0m, positioned with its corners at (-2.5, -2.0) and (2.5,

2.0) in the coordinate system, as shown in figure 3.2. This single-room configuration ensures consistent comparison across all placement algorithms, representing a typical small-to-medium indoor environment. The room acoustics are modelled using the mirror image method implemented in Pyroomacoustics, with wall surfaces characterised by an energy absorption coefficient of 0.2, corresponding to a reflection coefficient of approximately 0.894. This represents moderately reflective walls similar to painted concrete or plasterboard surfaces. The mirror image method computes reflections up to order 10, capturing the dominant reverberation characteristics while maintaining computational efficiency. Under these conditions, the room exhibits an estimated **Reverberation Time 60 (RT60)*** of 0.3-0.5 seconds, typical for small indoor spaces with moderate acoustic treatment.

3.3.1.2 Spatial Discretisation Strategy

The simulation framework implements three discretisation configurations to balance computational efficiency with spatial accuracy requirements. Figure 3.3 shows different configurations and their corresponding number of control points per wavelength. The default configuration (2) employs a 0.05m (5cm) grid spacing, providing approximately 13.6 grid points per wavelength at 500 Hz, which ensures adequate spatial sampling for accurate transfer function computation while maintaining reasonable computational demands. A high-resolution configuration with 0.025m (2.5cm) spacing is available primarily for visualisation purposes and detailed analysis of specific cases, offering 27.2 points per wavelength. Additionally, a fast-test configuration using 0.1m (10cm) spacing serves as a debugging tool, enabling rapid algorithm development and parameter tuning with reduced computational overhead. The default 0.05m resolution represents the primary analysis configuration, chosen to align with established practices in the acoustic simulation literature and facilitate comparison with existing placement optimisation studies.

3.3.1.3 Frequency Domain Analysis

The simulation operates in the narrowband domain at a single frequency of 500 Hz, selected as a representative mid-frequency that provides favourable conditions for algorithm development and validation. At this frequency, the wavelength of 0.68m relative to the 0.05m grid spacing ensures smooth spatial

***RT60** represents the reverberation time, defined as the time required for the sound pressure level in a reverberant environment to decay by 60 dB after the source stops




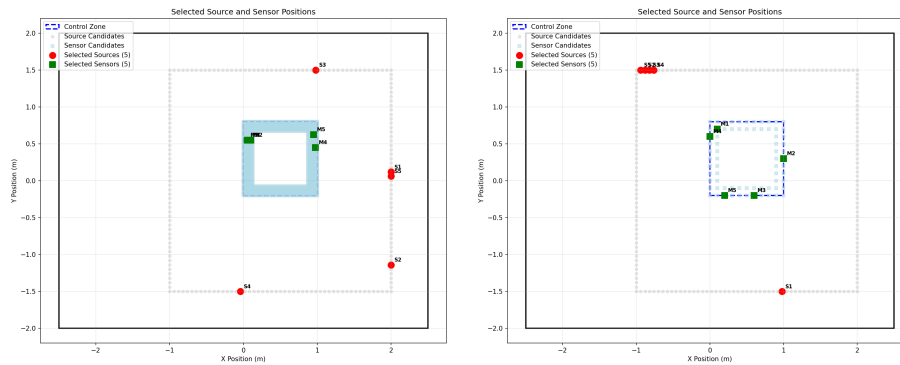
Figure 3.3: Comparison of the different discretisation steps with a wavelength of a 500Hz sine wave, where each circle represents a control point.

variation of the acoustic field without abrupt phase changes between adjacent grid points, facilitating reliable optimisation convergence. The choice of 500 Hz also happens to represent a frequency of practical importance for speech intelligibility [33]. While the placement optimisation results are inherently frequency-dependent due to spatial aliasing effects and wavelength-dependent interference patterns, the 500 Hz analysis serves as a solid foundation for algorithm comparison and validation. Extension to multiple frequencies and broadband optimisation represents a natural progression of this work, though implementation was deferred due to project timeline constraints. The narrowband approach enables thorough investigation of the fundamental optimisation principles while maintaining computational tractability for extensive algorithm comparison studies. A more extensive discussion on broadband extension can be found in section 5.4.3.

3.4 Simulation Model Validation

3.4.1 Greedy Algorithm Convergence Monitoring

The primary validation approach involves tracking the cost function evolution throughout the greedy selection process. For each algorithm iteration, the joint cost function value is recorded and plotted against the number of selected source-sensor pairs. Proper algorithm behaviour is verified by confirming monotonic cost function improvement, where each greedy selection step results in a reduction in the objective function value. This convergence monitoring serves as both a validation tool and a diagnostic method for identifying potential algorithmic issues. Across all tested configurations, the



(a) Example of sensor clustering behaviour exhibited by the greedy algorithm under high-quality test configuration

(b) Example of source clustering behaviour exhibited by the matching pursuit algorithm under fast test configuration

Figure 3.4: Clustering examples

greedy algorithms consistently demonstrate the expected monotonic decrease in cost function values, confirming correct implementation of the optimisation logic.

3.4.2 Spatial Distribution Analysis

During algorithm development, occasional clustering behaviour has been observed where selected sources and sensors concentrate in small spatial regions rather than distributing across the available candidate positions (Fig. 3.4). This clustering phenomenon may indicate regularisation issues or conditioning problems in the transfer function matrix \mathbf{G} , particularly when strong acoustic coupling between nearby positions creates preferential selection zones. Such cases are identified through visualisation of the selected positions and represent areas for future algorithmic refinement, potentially through explicit spatial diversity constraints or improved regularisation techniques.

3.4.3 Deterministic Reproducibility

All algorithms except the random placement baseline exhibit deterministic behaviour, producing identical results across multiple executions with the same input parameters. This deterministic property enables reliable comparison studies and ensures that performance differences between algorithms reflect genuine algorithmic characteristics rather than stochastic

variations. The random placement algorithm reproducibility is validated through Monte Carlo analysis with multiple independent runs using different random seeds.

3.5 Performance Analysis Methods

3.5.1 Primary Performance Metrics

Algorithm performance is evaluated using the Normalised Mean Square Error (NMSE) between the desired and reproduced sound fields within the control region. The NMSE serves as the primary quantitative metric for comparing placement optimisation algorithms, computed as:

$$\text{NMSE} = \frac{\|\mathbf{u}_{\text{reproduced}} - \mathbf{u}_{\text{desired}}\|^2}{\|\mathbf{u}_{\text{desired}}\|^2} \quad (3.1)$$

Where $\mathbf{u}_{\text{reproduced}}$ is the reproduced, or synthesised, sound field and $\mathbf{u}_{\text{desired}}$ is the desired sound field. The NMSE represents the ratio of reproduction error energy to desired field energy across all discretised control points.

3.5.2 Multi-Angle Robustness Assessment

To evaluate algorithm robustness beyond the nominal design conditions, each optimised placement is tested across multiple incident plane wave angles. The evaluation spans a range of angles centred around the design direction, with NMSE computed at each test angle. Specifically, performance is evaluated at angles $\alpha - y, \alpha - y + x, \dots, \alpha + y$, where x represents the evaluation step size and y defines the angular evaluation range. Figure 3.5 shows the angular evaluation strategy.

The overall robustness metric is calculated as the average NMSE across all evaluated angles, providing a comprehensive measure that accounts for both nominal performance and sensitivity to directional variations. The expected robustness profile exhibits an inverted bell-shaped curve when plotting NMSE versus incident angle, with minimum error at the design angle and increasing error for off-design directions.

3.5.3 Baseline Establishment

For the random placement algorithm, statistical characterisation is achieved through Monte Carlo simulation with 1000 independent runs using different




Figure 3.5: Plot of the evaluation angles where X is the angular discretisation step, y is the evaluation angular range, the bold arrow is the initial plane wave direction, and the other arrows represent the angle at which the NMSE is computed. For legibility, not all arrows are shown on the figure, but the dotted arrow is repeated every x angles, spanning a range of $2 \times y$.

random seeds. The best-performing random selection from this ensemble serves as the baseline reference, representing the upper bound of achievable performance through purely stochastic placement. All deterministic algorithms are evaluated through single-run executions, as their deterministic nature ensures reproducible results under identical input conditions.

3.5.4 Visual Field Analysis

Performance evaluation is complemented by visual comparison of synthesised sound fields through spatial interpolation and field visualisation techniques. These visualisations enable qualitative assessment of reproduction accuracy and identification of spatial regions with superior or degraded performance characteristics.

Chapter 4

Implementation and Algorithm Development

4.1 Joint Cost Function Development

4.1.1 Motivation and Problem Context

Traditional spatial ANC systems optimise secondary source placement and sensor placement independently, leading to suboptimal overall system performance. While some recent work has begun to address joint placement optimisation, existing cost functions suffer from fundamental limitations in how they balance synthesis accuracy and spatial estimation quality across the control region.

The core challenge lies in formulating a cost function that simultaneously optimises secondary source placement for effective sound field synthesis and sensor placement for accurate spatial interpolation, while properly accounting for the estimation uncertainty inherent in limited sensor measurements.

4.1.2 Parameter Definitions and Notation

To ensure consistency throughout the mathematical development, the following notation is established:

Spatial Configuration:

- P, Q : Sets of **candidate** secondary source and sensor positions, respectively, with cardinalities $|P| = p$ and $|Q| = q$.
- $L \subset P, M \subset Q$: Sets of **selected** secondary source and sensor indices, respectively, with cardinalities $|L| = l$ and $|M| = m$.

- $\Phi_L \in \{0, 1\}^{p \times l}$: Binary selection matrix for secondary sources
- $\Phi_M \in \{0, 1\}^{q \times m}$: Binary selection matrix for sensors

Sound Fields:

- $\mathbf{u}_{\text{des}} \in \mathbb{C}^m$: Desired sound field at **selected** sensor locations
- $\mathbf{u}_\epsilon \in \mathbb{C}^q$: Desired sound field at all the control points within the control zone.
- $\mathbf{d} \in \mathbb{C}^l$: Complex driving signals for selected secondary sources
- $\boldsymbol{\mu} \in \mathbb{C}^m$: Mean of the desired sound field distribution

regularisation Parameters:

- $\eta \in \mathbb{R}^+$: regularisation parameter for driving signal constraint
- $\lambda \in \mathbb{R}^+$: regularisation parameter for matrix inversion stability

Regularisation parameter selection follows an adaptive approach based on matrix conditioning. For each matrix requiring an inversion, the regularisation parameter is set to $\lambda = 10^{-4} \times k(\mathbf{A})$, where $k(\cdot)$ denotes the condition number. This choice provides an adaptive scaling; the more ill-conditioned the problem is, the more regularisation it will provide and numerical stability with 10^{-4} being both big enough to prevent numerical issues while remaining small enough to preserve some accuracy. The choice for this coefficient and technique comes from previous implementation ([22], [17]).

4.1.3 Acoustic Transfer Function Formulation

The transfer function matrix \mathbf{G} forms the core acoustic coupling between source and sensor candidates. Each element G_{ij} represents the complex transfer function from the j -th source candidate to the i -th sensor candidate at the analysis frequency.

Mathematically, \mathbf{G} is constructed from individual transfer function vectors:

$$\mathbf{G} = \begin{bmatrix} \mathbf{g}^T(\mathbf{r}_1) \\ \mathbf{g}^T(\mathbf{r}_2) \\ \vdots \\ \mathbf{g}^T(\mathbf{r}_q) \end{bmatrix} \in \mathbb{C}^{q \times p} \quad (4.1)$$

where $\mathbf{g}(\mathbf{r}_i) = [g(\mathbf{r}_i, \mathbf{s}_1), g(\mathbf{r}_i, \mathbf{s}_2), \dots, g(\mathbf{r}_i, \mathbf{s}_p)]^T$ is the vector of transfer functions from all p source candidates to sensor candidate at position \mathbf{r}_i .

The individual transfer function $g(\mathbf{r}_i, \mathbf{s}_j)$ represents the acoustic coupling from source position \mathbf{s}_j to receiver position \mathbf{r}_i , and takes different forms depending on the acoustic environment:

Free-field condition: In anechoic environments, that is, environments where reflections are absent and propagation is unrestricted, the transfer function is given by the 2D Green's function:

$$g(\mathbf{r}_i, \mathbf{s}_j) = \frac{j}{4} H_0^{(1)}(k|\mathbf{r}_i - \mathbf{s}_j|) \quad (4.2)$$

where $H_0^{(1)}$ is the zeroth-order Hankel function of the first kind, $k = 2\pi f/c$ is the wave-number, f is the frequency, and c is the speed of sound.

Reverberant environment: In rooms with reflecting boundaries, the transfer function includes contributions from both direct and reflected paths using the mirror image method:

$$g(\mathbf{r}_i, \mathbf{s}_j) = \sum_{n=0}^{N_{\max}} \beta^n \frac{j}{4} H_0^{(1)}(k|\mathbf{r}_i - \mathbf{s}_j^{(n)}|) \quad (4.3)$$

where $\mathbf{s}_j^{(n)}$ represents the n -th order mirror image of source j , β is the wall reflection coefficient, and N_{\max} is the maximum reflection order considered for computational tractability.

4.1.4 Weighting Matrix Formulation

The weighting matrix $\mathbf{W} \in \mathbb{C}^{m \times m}$ incorporates spatial interpolation characteristics into the cost function formulation. Every sensor position in the control area will not have the same *weight* in the sound field interpolation, and this matrix serves as the weighting factor enabling finer interpolation.

Following the weighted pressure matching framework presented in [23], the weighting matrix is defined as:

$$\mathbf{W} = \mathbf{P}_M^H \Phi_M^T \left[\int_{\Omega} \boldsymbol{\kappa}^*(r) \boldsymbol{\kappa}^T(r) dr \right] \Phi_M \mathbf{P}_M \quad (4.4)$$

Where:

- $\mathbf{P}_M = (\Phi_M^T \mathbf{K} \Phi_M + \lambda \mathbf{I}_Q)^{-1}$ is the regularised interpolation matrix
- Φ_M is the sensor selection matrix

- $\boldsymbol{\kappa}(r) = [\kappa(r, r_1), \dots, \kappa(r, r_j)]^T$ is the kernel function vector with elements defined by the 2D Bessel function kernel:

$$\kappa(r_1, r_2) = J_0(k\|r_1 - r_2\|_2) \quad (4.5)$$

where J_0 is the zeroth-order Bessel function of the first kind and k is the wave-number

- $\mathbf{K} \in \mathbb{C}^{q \times q}$ is the Gram matrix of kernel functions between all candidate sensor positions, with elements $\mathbf{K}_{jj'} = \kappa(r_j, r_{j'})$
- λ is the regularisation parameter for numerical stability

The choice of the 2D Bessel function kernel is theoretically grounded in **RKHS** theory for acoustic field interpolation [34]. The kernel is derived from the spatial correlation structure of a diffuse sound field:

$$\kappa(r_1, r_2) = \frac{1}{2\pi} \int_{S^1} e^{-jk^T(r_1 - r_2)} d\xi = J_0(k\|r_1 - r_2\|_2) \quad (4.6)$$

This kernel satisfies the reproducing property for functions satisfying the Helmholtz equation, ensuring that the interpolated field respects the underlying acoustic physics [35, 36]. The resulting Gram matrix \mathbf{K} corresponds to the normalised spatial covariance matrix of a diffuse sound field, providing optimal interpolation in the **RKHS** sense [34].

This weighting matrix formulation accounts for the spatial interpolation properties of the selected sensor configuration. The integral $\int_{\Omega} \boldsymbol{\kappa}^*(r) \boldsymbol{\kappa}^T(r) dr$ represents the spatial correlation structure of the kernel interpolation across the target region Ω , while the matrices \mathbf{P}_M and $\boldsymbol{\Phi}_M$ project this structure onto the selected sensor subset.

4.1.5 Initial Cost Function Development

Baseline Synthesis Formulation

The starting point for joint placement optimisation considers the standard least-squares sound field synthesis problem [22]. For a given source selection L , the synthesis error is minimised by*:

*In the original paper ([22]), the author works with mode matching rather than pressure matching. The equations in this section have been adapted to pressure matching and therefore do not completely reflect the source.

$$F(\mathbf{d}) = (\mathbf{G}\Phi_L\mathbf{d} - \mathbf{u}_{\text{des}})^H \mathbf{W}(\mathbf{G}\Phi_L\mathbf{d} - \mathbf{u}_{\text{des}}) + \eta\|\mathbf{d}\|^2 \quad (4.7)$$

where $\mathbf{G}\Phi_L\mathbf{d}$ represents the sound field synthesised by the selected sources, and the regularisation term $\eta\|\mathbf{d}\|^2$ prevents excessive driving signal amplitudes.

Optimal Driving Signal Solution

The optimal driving signal vector $\hat{\mathbf{d}}$ that minimises Equation (4.7) is obtained by setting $\nabla_{\mathbf{d}}F(\mathbf{d}) = 0$:

$$\hat{\mathbf{d}} = (\Phi_L^T \mathbf{G}^H \mathbf{W} \mathbf{G} \Phi_L + \eta \mathbf{I})^{-1} \Phi_L^T \mathbf{G}^H \mathbf{W} \mathbf{u}_{\text{des}} \quad (4.8)$$

Define the synthesis matrix:

$$\mathbf{A} = (\Phi_L^T \mathbf{G}^H \mathbf{W} \mathbf{G} \Phi_L + \eta \mathbf{I})^{-1} \quad (4.9)$$

such that $\hat{\mathbf{d}} = \mathbf{A} \Phi_L^T \mathbf{G}^H \mathbf{W} \mathbf{u}_{\text{des}}$.

Extension to Joint Source and Sensor Selection

The baseline formulation in Equation (4.7) does not account for sensor location optimisation, implicitly assuming that measurements are available at all control points within the control zone. To enable joint optimisation of both source and sensor placement, sensor selection matrices are introduced into the formulation to explicitly model the subset of control points where sensors are actually placed.

For joint source and sensor selections L and M , the cost function becomes:

$$F(\mathbf{d}) = (\mathbf{G}\Phi_L\mathbf{d} - \mathbf{u}_{\text{des}})^H \Phi_M^T \mathbf{W} \Phi_M (\mathbf{G}\Phi_L\mathbf{d} - \mathbf{u}_{\text{des}}) + \eta\|\mathbf{d}\|^2 \quad (4.10)$$

Where Φ_M selects the active sensor subset from all available sensor candidates, and \mathbf{W} now operates in the selected sensor space.

The optimal driving signal solution becomes:

$$\hat{\mathbf{d}} = (\Phi_L^T \mathbf{G}^H \Phi_M^T \mathbf{W} \Phi_M \mathbf{G} \Phi_L + \eta \mathbf{I})^{-1} \Phi_L^T \mathbf{G}^H \Phi_M^T \mathbf{W} \Phi_M \mathbf{u}_{\text{des}} \quad (4.11)$$

Derivation of the D Matrix

Substituting the optimal driving signal solution into the expected cost and expanding terms yields:

$$F(\hat{\mathbf{d}}) = \mathbf{u}_{des}^H (\Phi_M \mathbf{W} \Phi_M^T - \Phi_M \mathbf{W} \Phi_M^T \mathbf{G} \Phi_L \mathbf{A} \Phi_L^T \mathbf{G}^H \Phi_M \mathbf{W} \Phi_M^T) \mathbf{u}_{des} \quad (4.12)$$

$$\mathbf{A} = (\Phi_L^T \mathbf{G}^H \Phi_M \mathbf{W} \Phi_M^T \mathbf{G} \Phi_L + \lambda \mathbf{I})^{-1} \quad (4.13)$$

$$\mathbf{D} = \Phi_M \mathbf{W} \Phi_M^H - \Phi_M \mathbf{W} \Phi_M^H \mathbf{G} \Phi_L \mathbf{A} \Phi_L^T \mathbf{G}^H \Phi_M \mathbf{W} \Phi_M^H \quad (4.14)$$

This formulation represents the residual error matrix after optimal synthesis.

The expected cost becomes:

$$\mathcal{J}(L, M) = \mathbb{E}[F(\hat{\mathbf{d}})] = \text{Tr}[\mathbf{D}] \quad (4.15)$$

Stochastic Extension for Robustness

To improve placement robustness across varying acoustic conditions, the formulation incorporates prior knowledge about the statistical distribution of desired sound fields. Assuming the desired field follows a Gaussian distribution $\mathbf{u}_{des} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ represents the mean field and $\boldsymbol{\Sigma}$ [23] captures spatial correlations, the expected cost becomes:

$$\mathcal{J}(L, M) = \mathbb{E}[F(\hat{\mathbf{d}})] = \text{Tr}[\mathbf{D}\boldsymbol{\Sigma}] + \boldsymbol{\mu}^H \mathbf{D} \boldsymbol{\mu} \quad (4.16)$$

This stochastic formulation enables placement optimisation that performs well across a range of expected sound field scenarios rather than a single deterministic case.

4.1.6 Implementation and Observed Limitations

The initial cost function was implemented using the trace formulation $\mathcal{J} = \text{Tr}[\mathbf{D}\boldsymbol{\Sigma}]$, with the spatial covariance matrix $\boldsymbol{\Sigma}$ modelling expected field correlations. However, practical implementation revealed several critical limitations:

Sensor Clustering: The optimisation consistently produced sensor configurations concentrated in small spatial regions rather than distributed across the control zone. Figure 4.1 demonstrates the undesirable clustering

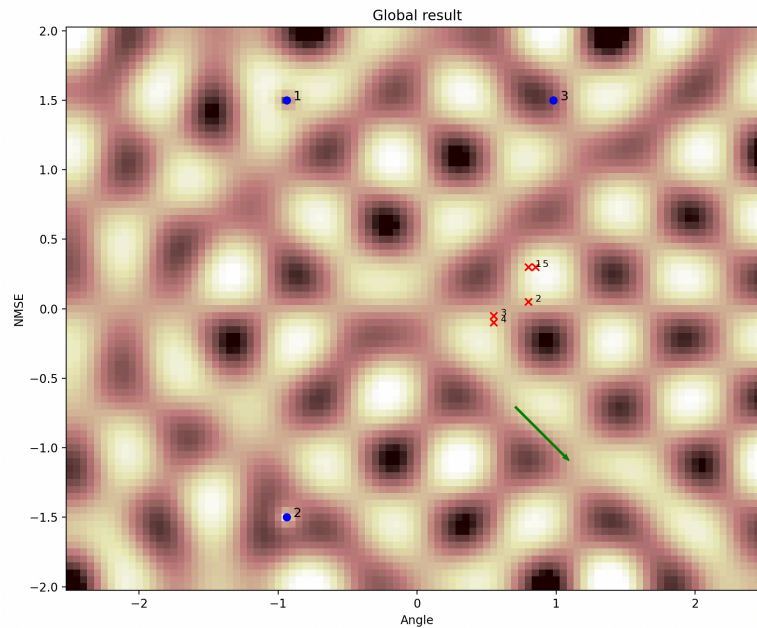


Figure 4.1: Undesirable clustering behaviour.

behaviour where newly selected sensors (numbered 4 and 5, red crosses) are placed very close to existing sensors rather than distributed across the control region. Blue dots represent selected secondary sources, and the green arrow indicates the desired plane wave direction. This clustering shows that the cost function prioritises local reproduction accuracy over spatial interpolation capability, leading to poor coverage of the control zone. The greedy algorithm (Sec. 4.3.3) initially selected three source/sensor pairs, then attempted to add two additional sensors, resulting in the clustered configuration shown. This limitation motivated the development of the novel joint cost function presented in Section 4.1.7..

Convergence Issues: Figure 4.2 shows that cost values exhibited monotonically increasing behaviour during greedy optimisation, indicating deteriorating rather than improving performance.

These limitations indicate an issue with the formulated cost function. It seems to show that the function favours synthesis accuracy at the selected sensors, discarding the interpolation accuracy.

These limitations motivated the development of an alternative cost function that explicitly accounts for spatial interpolation uncertainty and full-field performance requirements.

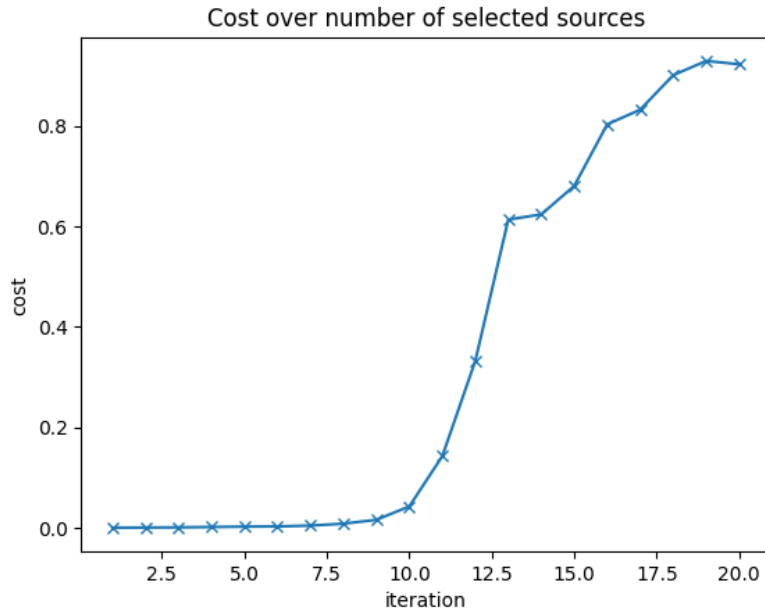


Figure 4.2: Cost yielded by the cost function over the amount of selected sources. This results from the greedy algorithm (See sec. 4.3.3) selecting 20 pairs of sources and sensors iteratively.

4.1.7 Novel Joint Cost Function Formulation

To address the identified limitations, a new joint cost function was developed that explicitly incorporates interpolation error alongside synthesis error by comparing the interpolated field estimate directly against the true desired field \mathbf{u}_e across the entire control zone, Ω , where $\mathbf{K}(\mathbf{r})$ represents the spatial interpolation kernel evaluated at position \mathbf{r} :

$$Q(\mathbf{d}) = \int_{\Omega} |\mathbf{K}(\mathbf{r}) (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{G} \mathbf{d} - \mathbf{u}_e|^2 d\mathbf{r}, \quad (4.17)$$

Here, $\mathbf{K}(\mathbf{r})(\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{G} \mathbf{d}$ can be interpreted as the synthesised sound field ($\approx \mathbf{u}_{\text{syn}}$). Discretising the spatial domain and representing the system in matrix form yields:

$$Q(\mathbf{d}) = \|\mathbf{K}_e (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{G} \mathbf{d} - \mathbf{u}_e\|^2, \quad (4.18)$$

where \mathbf{K}_e is the Gram matrix between all control points and the candidate sensor positions. When replacing \mathbf{d} with the optimal driving signal $\hat{\mathbf{d}}$, similarly to what has been done in Eq. 4.8, the equation becomes :

$$Q(\hat{\mathbf{d}}) = \|\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{G}(\mathbf{G}^H\mathbf{W}\mathbf{G} + \eta\mathbf{I})^{-1}\mathbf{G}^H\mathbf{u}_{des} - \mathbf{u}_\epsilon\|^2. \quad (4.19)$$

Finally, when inserting the selection matrices Φ_L and Φ_M for sources and sensors, respectively, the equation becomes:

$$Q(\hat{\mathbf{d}}) = \|\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{G}\Phi_L (\Phi_L^T\mathbf{G}^H\Phi_M\mathbf{W}\Phi_M^T\mathbf{G}\Phi_L + \eta\mathbf{I})^{-1}\Phi_L^T\mathbf{G}^H\Phi_M\mathbf{u}_{des} - \mathbf{u}_\epsilon\|^2. \quad (4.20)$$

4.1.8 Mathematical Development

Recognising that $\mathbf{u}_{des} = \Phi_{C'}^T\mathbf{u}_\epsilon$, where $\Phi_{C'}^T$ extracts the desired field at sensor positions from the full field ($\neq \Phi_M$), the cost function becomes:

$$\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{G}\Phi_L\mathbf{A}\Phi_L^T\mathbf{G}^H\Phi_M\Phi_{C'}^T\mathbf{u}_\epsilon - \mathbf{u}_\epsilon, \quad (4.21)$$

where $\mathbf{A} = (\Phi_L^T\mathbf{G}^H\Phi_M\mathbf{W}\Phi_M^T\mathbf{G}\Phi_L + \eta\mathbf{I})^{-1}$ is the regularised synthesis matrix.

This yields the quadratic form $\mathbf{u}_\epsilon^H\mathbf{D}\mathbf{u}_\epsilon$ with error covariance matrix:

$$\mathbf{D} = (\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{G}\Phi_L\mathbf{A}\Phi_L^T\mathbf{G}^H\Phi_M\Phi_{C'}^T - \mathbf{I})^H (\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{G}\Phi_L\mathbf{A}\Phi_L^T\mathbf{G}^H\Phi_M\Phi_{C'}^T - \mathbf{I}). \quad (4.22)$$

The final placement cost function becomes:

$$Q(L, M) = \text{trace}(\mathbf{D}\Sigma) + \mu^H\mathbf{D}\mu \quad (4.23)$$

4.1.9 Formulation Properties

Reconstruction Error Quantification: The matrix \mathbf{D} measures deviation from perfect field reconstruction, with the identity matrix \mathbf{I} corresponding to perfect interpolation.

Combined Interpolation-Synthesis Terms: The formulation includes both synthesis accuracy (through \mathbf{A}) and estimation accuracy (through $\mathbf{K}_\epsilon(\mathbf{K} + \lambda\mathbf{I})^{-1}$) in a single expression.

Convergence Behaviour: This formulation produces monotonically decreasing cost values during greedy optimisation, as shown in Fig. 4.3.

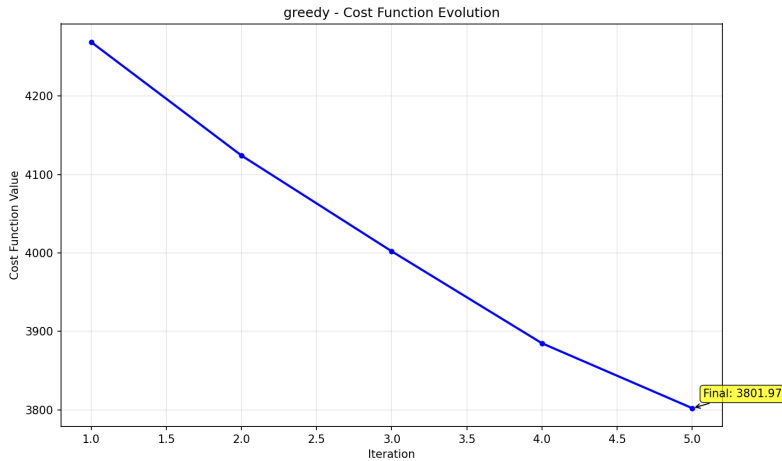


Figure 4.3: Cost history for the greedy algorithm (See Sec. 4.3.3). The cost function exhibits a monotonic decrease with the number of source/sensor pairs selected.

4.2 Software Architecture and Implementation

The simulation framework implements a *pragmatic layered architecture* designed for reproducible algorithm comparison. The system employs four layers with strategic cross-layer communication: configuration management, domain modelling (Room, ControlZone), algorithm implementations, and matrix computations. Detailed software architecture can be found in Appendix A.

All system behaviour is controlled through hierarchical TOML files with strongly-typed data-class validation. Configuration feeds directly into all layers simultaneously, eliminating hardcoded parameters and ensuring reproducible results across algorithm comparisons.

All optimisation algorithms implement the same interface: `algorithm(room, num_sources, num_sensors) → (sources, sensors, cost)`. Dynamic loading via `__getattr__` enables plug-and-play comparison, while algorithms access standardised Room properties (matrices G , K , W) and directly call matrix computations for performance optimisation.

Pyroomacoustics integration occurs at the Room level, enabling seamless switching between free-field (2D Green's functions) and reverberant environments (mirror image method). Additional scientific libraries (numpy, scipy) are integrated at the matrix computation layer for mathematical operations.

4.3 Placement optimisation Algorithms

This section will introduce the algorithms tested with the cost function developed in Section 4.1.7.

4.3.1 Random Placement Baseline

This simple algorithm (Alg. 1) serves as a baseline for performance comparison with the other algorithms. It randomly selects subsets $L \subset P$ and $M \subset Q$ of the desired cardinalities $|L| = l$ and $|M| = m$, respectively, where l and m are simulation parameters determined by the user's requirements. To ensure the random baseline is meaningful, the algorithm performs multiple independent runs and selects the configuration that achieves the lowest cost function value. The number of runs is typically set to one thousand to provide adequate sampling of the solution space.

Algorithm 1 Random selection algorithm

Require: Candidate source positions P , candidate sensor positions Q , desired number of sources l , desired number of sensors m , number of runs N_{runs}

Ensure: Selection of l sources L and m sensors M

Initialise $best_cost \leftarrow \infty$

Initialise $L_{best} \leftarrow \emptyset, M_{best} \leftarrow \emptyset$

for $run = 1$ to N_{runs} **do**

$P_{available} \leftarrow P, Q_{available} \leftarrow Q$

\triangleright Randomly select l sources without replacement

for $i = 1$ to l **do**

$\hat{r}_i \leftarrow$ random choice from $P_{available}$

$L \leftarrow L \cup \{\hat{r}_i\}$

$P_{available} \leftarrow P_{available} \setminus \{\hat{r}_i\}$

end for

\triangleright Randomly select m sensors without replacement

for $j = 1$ to m **do**

$\hat{r}_j \leftarrow$ random choice from $Q_{available}$

$M \leftarrow M \cup \{\hat{r}_j\}$

$Q_{available} \leftarrow Q_{available} \setminus \{\hat{r}_j\}$

end for

$current_cost \leftarrow \text{ComputeCost}(L, M)$

if $current_cost < best_cost$ **then**

$best_cost \leftarrow current_cost$

$L_{best} \leftarrow L, M_{best} \leftarrow M$

end if

end for

return (L_{best}, M_{best})

4.3.1.1 Complexity Analysis

Time Complexity: The algorithm's computational cost is dominated by the cost function evaluation performed in each run. The cost function requires matrix inversion operations that scale with the cube of the number of selected sources. The overall time complexity is:

$$O(N_{runs} \times l^3) \quad (4.24)$$

where N_{runs} is the number of random trials and l is the number of sources to select. The algorithm scales linearly with the number of runs, but cubically with the number of sources/sensors to select.

4.3.2 Regular Placement Algorithm

A first naive approach for source and sensor placement is to arrange them in a regular geometric pattern around the control zone (Alg. 2). This deterministic method provides a second baseline that requires no optimisation or prior knowledge about the desired sound field, making it computationally efficient and reproducible. The regular placement strategy distributes sources and sensors uniformly around the control zone perimeter based on angular spacing, ensuring good spatial coverage without requiring cost function evaluations. While this approach may not achieve optimal acoustic performance, it serves as a practical starting point and provides insight into the benefits of more sophisticated optimisation methods. The algorithm selects l sources from the candidate set P and m sensors from the candidate set Q by computing angular positions relative to the control zone centre and selecting evenly spaced candidates in clockwise order.

Algorithm 2 Regular Placement Algorithm

Require: Control zone centre \mathbf{c} , candidate source positions P , candidate sensor positions Q , desired number of sources l , desired number of sensors m

Ensure: Selection of l sources L and m sensors M

Initialise $L = \emptyset, M = \emptyset$

▷ Regular sensor selection around control zone perimeter

for $i = 1$ to q **do**

$\theta_i \leftarrow \text{atan2}(q_{i,y} - c_y, q_{i,x} - c_x)$ ▷ Angle from centre to sensor i

end for

sorted_sensors \leftarrow sort(Q , by angle θ) ▷ Clockwise order

$step_s \leftarrow q/m$

for $j = 1$ to m **do**

$idx \leftarrow \lfloor (j - 1) \times step_s \rfloor \bmod q$

$M \leftarrow M \cup \{\text{sorted_sensors}[idx]\}$

end for

▷ Regular source selection around control zone perimeter

for $i = 1$ to p **do**

$\phi_i \leftarrow \text{atan2}(p_{i,y} - c_y, p_{i,x} - c_x)$ ▷ Angle from center to source i

end for

sorted_sources \leftarrow sort(P , by angle ϕ) ▷ Clockwise order

$step_{src} \leftarrow p/l$

for $k = 1$ to l **do**

$idx \leftarrow \lfloor (k - 1) \times step_{src} \rfloor \bmod p$

$L \leftarrow L \cup \{\text{sorted_sources}[idx]\}$

end for

return (L, M)

Complexity Analysis

Time Complexity: The algorithm's computational cost is dominated by sorting operations to arrange candidates by angular position around the control zone centre. The time complexity is:

$$O(q \log q + p \log p) \quad (4.25)$$

Where q and p are the total numbers of candidate sensors and sources, respectively. Unlike optimisation-based approaches, this algorithm performs no cost function evaluation, eliminating expensive matrix operations and resulting in significantly lower computational cost.

This complexity is implementation-dependent. The $O(q \log q)$ sorting

cost could be reduced to $O(m + l)$ if candidate positions are organised in a spatial data structure that supports efficient nearest-neighbour queries by angular position, eliminating the need for explicit sorting.

4.3.3 Greedy Pair-Selection Algorithm

The greedy pair-selection algorithm (Alg. 3) represents a joint optimisation approach that simultaneously selects sources and sensors by minimising the cost function at each iteration. This method directly optimises the acoustic control cost function $S(L, M)$, which incorporates prior statistical information about the desired sound field through the covariance matrix Σ and mean vector μ . At each iteration, the algorithm evaluates all possible source-sensor pairs from the remaining candidates and selects the pair that yields the greatest reduction in the cost function. This greedy approach ensures monotonic improvement in acoustic performance while building the solution incrementally.

The algorithm must be adapted for cases where the desired number of sources and sensors differ ($l \neq m$). In such cases, the algorithm works as described with pair selection until $\min(l, m)$ pairs have been selected. Subsequently, the algorithm continues by selecting individual sources or sensors (whichever type remains to be selected) that minimise the cost function, given the current selection state.

Algorithm 3 Greedy pair selection algorithm

Require: Candidate source positions P , candidate sensor positions Q , desired number of sources l , desired number of sensors m , mean μ and covariance Σ matrices of \mathbf{u}_{des}

Ensure: Selection of l sources L and m sensors M

Initialise $L = \emptyset, M = \emptyset$

Initialise $P_{available} = P, Q_{available} = Q$

Initialise selection matrices Φ_L, Φ_M

Initialise iteration counters $i = 1, j = 1$

▷ Phase 1: Pair selection until $\min(l, m)$ pairs selected

while $i \leq l$ and $j \leq m$ **do**

Find $(\hat{r}_i, \hat{r}_j) = \arg \min_{r_i \in P_{available}, r_j \in Q_{available}} S(L \cup \{r_i\}, M \cup \{r_j\})$

$L \leftarrow L \cup \{\hat{r}_i\}$ and update Φ_L

$M \leftarrow M \cup \{\hat{r}_j\}$ and update Φ_M

$P_{available} \leftarrow P_{available} \setminus \{\hat{r}_i\}$

$Q_{available} \leftarrow Q_{available} \setminus \{\hat{r}_j\}$

$i \leftarrow i + 1, j \leftarrow j + 1$

end while

▷ Phase 2: Individual selection for remaining sources (if $l > m$)

while $i \leq l$ **do**

Find $\hat{r}_i = \arg \min_{r_i \in P_{available}} S(L \cup \{r_i\}, M)$

$L \leftarrow L \cup \{\hat{r}_i\}$ and update Φ_L

$P_{available} \leftarrow P_{available} \setminus \{\hat{r}_i\}$

$i \leftarrow i + 1$

end while

▷ Phase 3: Individual selection for remaining sensors (if $m > l$)

while $j \leq m$ **do**

Find $\hat{r}_j = \arg \min_{r_j \in Q_{available}} S(L, M \cup \{r_j\})$

$M \leftarrow M \cup \{\hat{r}_j\}$ and update Φ_M

$Q_{available} \leftarrow Q_{available} \setminus \{\hat{r}_j\}$

$j \leftarrow j + 1$

end while

return (L, M)

Complexity Analysis

Time Complexity: The algorithm's computational cost is dominated by the pair selection phase, where each iteration evaluates all combinations of

remaining source-sensor pairs. The cost function complexity grows cubically with the number of selected sources. The overall time complexity is:

$$O(p \times q \times (\min(l, m))^4)$$

Where p and q are the total numbers of candidate sources and sensors, and l and m are the numbers to select. The algorithm exhibits quartic scaling with the smaller of the two selection sizes due to the combination of evaluating all source-sensor pairs and the cubic cost of matrix inversion operations.

4.3.4 Matching Pursuit Algorithm

The **MP** algorithm represents a fundamentally different approach to source and sensor selection compared to the previously described methods. Unlike the greedy pair-selection algorithm that directly optimises the acoustic control cost function $S(L, M)$ and incorporates prior statistical information through the covariance matrix Σ and mean vector μ , the Matching Pursuit algorithm operates purely on signal processing principles without using any cost function evaluation [37]. Instead, MP treats the acoustic control problem as a sparse signal decomposition task, where the desired sound field is decomposed into a sparse representation using an overcomplete dictionary of source-sensor pair responses. At each iteration, the algorithm selects the source-sensor pair whose normalised field pattern best matches the current residual (the remaining error between desired and synthesised fields) by maximising the inner product $|\langle \mathbf{residual}, \phi_{i,j} \rangle|$. This approach relies solely on the instantaneous geometric similarity between field patterns and does not incorporate any prior knowledge about typical sound field characteristics or statistical distributions. The algorithm continues this greedy pattern-matching process until convergence criteria are met, building a sparse solution incrementally without ever evaluating the acoustic performance through the cost function.

Algorithm 4 Matching Pursuit for Acoustic Source and Sensor Selection

Require: Desired sound field \mathbf{u}_{des} , candidate source positions P , candidate sensor positions Q , convergence threshold ϵ , maximum iterations K_{max}

Ensure: Selection of source-sensor pairs $\mathcal{D} = \{(p_k, q_k)\}_{k=1}^K$

Initialise **residual** $\leftarrow \mathbf{u}_{des}$

Initialise selected pairs $\mathcal{D} \leftarrow \emptyset$

Initialise available sources $P_{avail} \leftarrow P$

Initialise available sensors $Q_{avail} \leftarrow Q$

Initialise iteration counter $k \leftarrow 1$

while $k \leq K_{max}$ and $\|\mathbf{residual}\|_2 > \epsilon$ **do**

$max_inner_product \leftarrow 0$

$(best_source, best_sensor) \leftarrow (null, null)$

\triangleright Find best source-sensor pair

for $p_i \in P_{avail}$ **do**

for $q_j \in Q_{avail}$ **do**

$\phi_{i,j} \leftarrow \text{GetDictionaryAtom}(p_i, q_j)$ \triangleright Normalised field pattern

$inner_product \leftarrow |\langle \mathbf{residual}, \phi_{i,j} \rangle|$

if $inner_product > max_inner_product$ **then**

$max_inner_product \leftarrow inner_product$

$(best_source, best_sensor) \leftarrow (p_i, q_j)$

end if

end for

end for

\triangleright Check convergence condition

if $max_inner_product < \epsilon$ **then**

break

\triangleright No significant improvement possible

end if

\triangleright Update selection and residual

$\phi_{best} \leftarrow \text{GetDictionaryAtom}(best_source, best_sensor)$

$coefficient \leftarrow \langle \mathbf{residual}, \phi_{best} \rangle$

$\mathbf{residual} \leftarrow \mathbf{residual} - coefficient \times \phi_{best}$

$\mathcal{D} \leftarrow \mathcal{D} \cup \{(best_source, best_sensor)\}$

$P_{avail} \leftarrow P_{avail} \setminus \{best_source\}$

$Q_{avail} \leftarrow Q_{avail} \setminus \{best_sensor\}$

$k \leftarrow k + 1$

end while

return \mathcal{D}

Algorithm 5 GetDictionaryAtom Function

Require: Source position l_i , sensor position m_j , room transfer functions

Ensure: Normalised dictionary atom $\phi_{i,j}$

$\mathbf{u}_{syn} \leftarrow \text{ComputeSynthesizedField}(l_i, m_j) \triangleright$ Field from source-sensor pair

$\phi_{i,j} \leftarrow \frac{\mathbf{u}_{syn}}{\|\mathbf{u}_{syn}\|_2} \triangleright$ Normalise to unit norm

return $\phi_{i,j}$

Complexity Analysis

Time Complexity: The algorithm's computational cost is determined by evaluating all source-sensor pair combinations in each iteration and computing forward field synthesis for each pair. The time complexity is:

$$O(K_{max} \times p \times q \times N_{field}) \quad (4.26)$$

Where K_{max} is the maximum number of iterations, p and q are the total numbers of candidate sources and sensors, and N_{field} is the size of the field vector. Unlike optimisation-based approaches, this algorithm performs forward field computations rather than inverse matrix operations, avoiding expensive matrix inversions and resulting in linear scaling with field size.

Chapter 5

Results and Analysis

This chapter presents the benchmark results for four acoustic control algorithms tested across three configuration profiles. The evaluation focuses on cost optimisation performance and computational efficiency, examining how these metrics relate to practical implementation requirements.

5.1 Major Results

The benchmark tested four algorithms (Random, Regular, Greedy, and Matching Pursuit) on a Linux system with Python 3.11. The evaluation used three resolution configurations: Fast Test (0.1m), Default (0.05m), and High Quality (0.025m). The details of these different configurations can be found in Appendix B.

5.1.1 Default Configuration Performance

Table 5.1 presents the performance results for the default configuration (0.05m resolution), which represents standard production conditions.

Table 5.1: Default Configuration Performance Results

Algorithm	Cost Value	Runtime (s)	Improvement (%)
Greedy	3 801.97	6 101.74	8.2
Random	4 143.29	46.75	0.0
Regular	4 154.12	0.01	-0.3
Matching Pursuit	4 277.24	176.24	-3.2

The Greedy algorithm achieves the best cost performance with a value

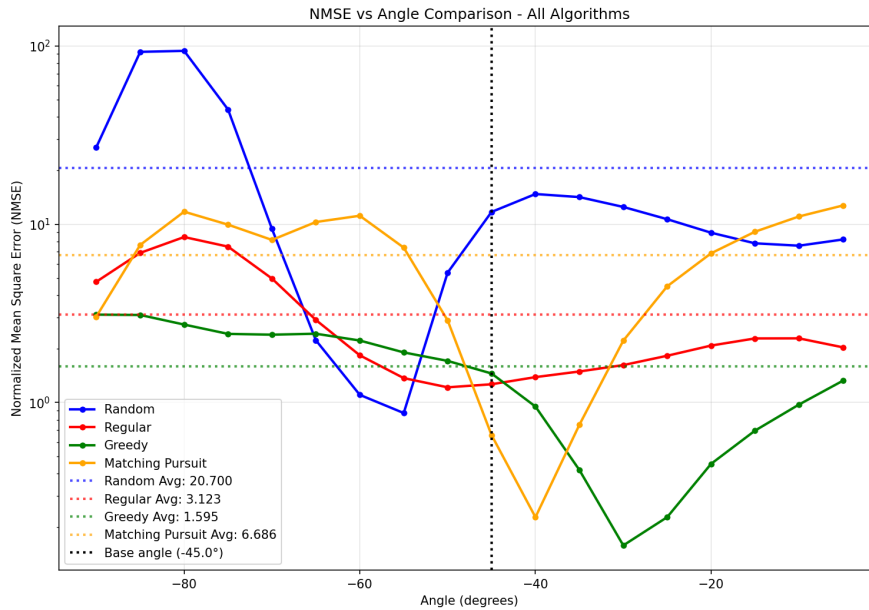


Figure 5.1: **NMSE** evaluation comparison for the four algorithms.

of 3,801.97, representing an 8.2% improvement over the Random baseline. However, this performance comes at a significant computational cost of 6,101.74 seconds, approximately 130 times longer than the Random algorithm's 46.75 seconds.

The **NMSE** evaluation for the greedy, random, regular placement and matching pursuit algorithms is shown in figure 5.1. The greedy algorithm, having the lowest average **NMSE**, shows that it is the only algorithm to truly consider the prior information on the sound field for placement. While the matching-pursuit algorithm performs better around the initial angle, it quickly becomes less precise when the desired plane wave angle changes more than ≈ 15 degrees.

5.1.2 Cross-Configuration Analysis

Table 5.2 shows the cost performance across all three configuration profiles for algorithms that completed successfully.

Table 5.2: Cost Performance Across Configuration Profiles

Algorithm	Fast Test (0.1m)	Default (0.05m)	High Quality (0.025m)
Random	1 121.21	4 143.29	16 238.90
Regular	1 298.84	4 154.12	16150.72
Greedy	907.87	3801.97	-
Matching Pursuit	1 160.17	4 277.24	16 500.84

The cost scaling follows a predictable pattern as resolution increases. From Fast Test to Default configuration, costs increase by factors of 3.7× (Random), 3.2× (Regular), 4.2× (Greedy), and 3.7× (Matching Pursuit). The Greedy algorithm maintains its performance advantage in the Fast Test configuration but did not complete the High Quality run for reasons explained in Sec 5.1.3.

5.1.3 Runtime Efficiency Analysis

Table 5.3 presents the runtime performance across all configurations.

Table 5.3: Runtime Performance Across Configuration Profiles

Algorithm	Fast Test (s)	Default (s)	High Quality (s)
Random	0.44	46.75	9 360.29
Regular	0.01	0.01	0.33
Greedy	181.85	6 101.74	-
Matching Pursuit	44.61	176.24	5 750.65

The Regular algorithm demonstrates exceptional runtime efficiency, completing all configurations in under one second. The Random algorithm shows reasonable scalability, while the Greedy algorithm's computational requirements increase dramatically with resolution, preventing completion of the High Quality configuration. Based on the scaling patterns observed across other algorithms, the Greedy algorithm would have required approximately 55 hours to complete the High Quality configuration, making it computationally prohibitive for practical applications at this resolution level.

5.1.4 Algorithm Convergence Behavior

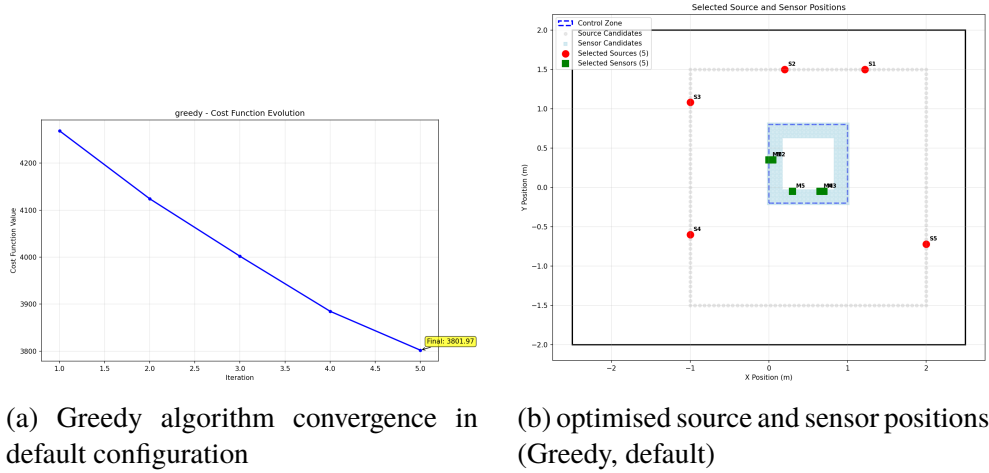
The Greedy algorithm demonstrates systematic improvement through iterative optimisation. Starting from an initial cost of 4,268.99, it converges to 3,801.97 through five iterations, showing a total improvement of 10.9%. The

cost progression follows: 4,268.99 → 4,124.30 → 4,002.07 → 3,884.76 → 3,801.97.

The Matching Pursuit algorithm shows much slower convergence. Over five iterations, it achieves only minimal improvement from 4,317.74 to 4,277.24 (0.9% total improvement). The algorithm did not reach convergence within the iteration limit, indicating limitations in its optimisation approach.

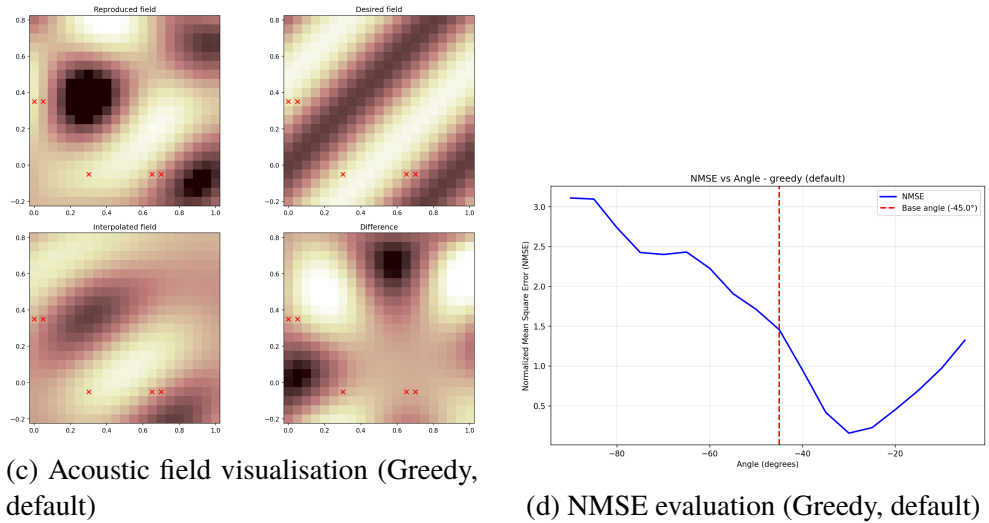
5.1.5 Performance visualisation Results

Figure 5.2 presents the key visualisation results from the benchmark, showing different aspects of algorithm performance.



(a) Greedy algorithm convergence in default configuration

(b) optimised source and sensor positions (Greedy, default)



(c) Acoustic field visualisation (Greedy, default)

(d) NMSE evaluation (Greedy, default)

Figure 5.2: Algorithm performance visualisation for Greedy algorithm in default configuration (Larger figures can be found in appendix C).

5.2 Baseline Performance Validation

The Random algorithm establishes a reliable baseline with consistent performance across all test configurations. Cost measurements demonstrate precision at two decimal places, with the smallest observed performance difference being 10.83 cost units (Regular vs Random), allowing for clear algorithm differentiation. The baseline algorithm exhibits predictable computational scaling behaviour with resolution: 0.44s (Fast Test), 46.75s (Default), and 9,360.29s (High Quality), confirming expected quadratic scaling with problem size.

5.3 Algorithm Efficiency Analysis

The benchmark quantifies the computational cost of optimisation improvements relative to the random baseline:

$$\text{Greedy Efficiency} = \frac{341.32 \text{ cost improvement}}{6054.99 \text{ additional seconds}} = 0.056 \text{ cost units/second}$$

This efficiency metric reveals that the Greedy algorithm's 8.2% performance improvement over the Random baseline requires substantial additional computation time, providing a quantitative framework for evaluating optimisation algorithm trade-offs in practical deployment scenarios.

5.4 Discussion


These results reveal fundamental trade-offs and limitations in the chosen algorithms for source and sensor placement optimisation.

5.4.1 Algorithm Performance


While the greedy algorithm consistently shows a better resulting selection, its poor scalability shows an important relationship between computational investment and optimisation quality. This finding aligns with recent literature, where it is mentioned that even in a 2D environment, the computational cost for such algorithms becomes quickly very challenging [18, 38, 39].

5.4.2 Scalability and Resolution Dependencies


All the algorithms scale relatively similarly across the different resolutions in which they were run. The simple regular placement algorithm, while being quick to run, demonstrates low performance (See Fig. 5.3). A need then arises for a good middle ground between performance and computational cost, where heuristic algorithms usually lie. The matching pursuit algorithm, although good for direct reproduction, has a lower robustness to changes in the desired sound field (See Fig. 5.4).



(a) optimised source and sensor positions (Regular, default)




(b) Acoustic field visualisation (Regular, default)




(c) NMSE evaluation (Regular, default)


Figure 5.3: Algorithm performance visualisation for the Regular algorithm in default configuration



(a) optimised source and sensor positions (Matching Pursuit, default)



(b) Acoustic field visualisation (Matching Pursuit, default)



(c) NMSE evaluation (Matching Pursuit, default)

Figure 5.4: Algorithm performance visualisation for Matching Pursuit algorithm in default configuration

5.4.3 Limitation to narrowband

This thesis is limited to testing narrowband (single-frequency) scenarios. In real-world applications, such occurrences are small to non-existent. Most practical noise sources exhibit broadband characteristics with energy distributed across multiple frequencies, requiring algorithms to handle complex spectral content simultaneously.

The narrowband limitation affects algorithm evaluation in several ways. Broadband scenarios introduce frequency-dependent optimisation challenges where algorithms must balance performance across the entire spectrum rather than optimising for a single frequency. The computational scaling observed for different algorithms may change significantly when extended to multiple frequency bins, potentially altering the performance hierarchy established in this study.

While the narrowband results provide valuable insights into algorithm behaviour and computational characteristics, extending these findings to broadband applications requires careful validation. The relative advantages of the Greedy algorithm's optimisation approach versus simpler methods like Regular placement may shift when managing the increased complexity of multi-frequency control, necessitating future work to evaluate algorithm performance under realistic broadband conditions.

A typical approach for extending placement algorithms to broadband scenarios involves running the optimization across multiple frequency bins, where each bin represents a discrete frequency range. For the 500 Hz to 1000 Hz range, for example, one might use five configurations with plane waves at the centre frequency of each bin (e.g., 550 Hz, 650 Hz, 750 Hz, 850 Hz, and 950 Hz). The cost function would then aggregate performance across all bins—either through uniform averaging or weighted summation emphasizing critical frequency ranges. However, this multi-frequency extension introduces several open questions relevant to this work: How does computational cost scale with the number of frequency bins for each algorithm? Does the Greedy algorithm's optimisation advantage over simpler placement methods strengthen or weaken in broadband scenarios? What weighting strategies best balance broadband performance with computational efficiency? These questions represent important directions for future research building upon the narrowband framework established in this thesis.

5.4.4 Practical Implementation Considerations

While the produced plots and metrics may help compare several algorithms, the simulation-based product is not a pure reflection of reality. Several factors distinguish real-world acoustic control implementations from idealised simulation conditions, potentially affecting algorithm performance and selection criteria.

The simulated environment assumes perfect knowledge of acoustic transfer functions, noise-free measurements, and ideal actuator responses. In practice, acoustic systems face measurement uncertainties, environmental variations, and hardware limitations that can significantly impact control effectiveness. Microphone noise and actuator non-linearities introduce errors not captured in the simulation framework.

Additionally, the simulation results depend on several assumptions that may not hold in practice. The acoustic modelling assumes linear superposition and ignores non-linear effects at high sound pressure levels. Potential implementation errors in the algorithm code, approximations in numerical methods, or incorrect parameter settings could bias the comparative analysis. The controlled simulation environment cannot replicate acoustic complexity, including reverberation, acoustic coupling, and moving obstacles that may alter relative performance rankings, potentially favouring robust algorithms over those optimised for idealised conditions.

Finally, purely mathematical comparison of pressure field reconstruction does not capture the complexity of human auditory perception. The cost metrics used throughout this analysis assume that minimising pressure differences directly translates to improved sound quality, yet human hearing involves sophisticated non-linear processing that may render small perturbations imperceptible. Consequently, algorithms showing inferior performance in pressure-based metrics might achieve perceptually equivalent results, while subtle differences between high-performing algorithms may be acoustically irrelevant to human listeners. Future evaluation frameworks should incorporate perceptual metrics to better assess the practical significance of algorithmic improvements beyond pure mathematical optimisation.

Chapter 6

Conclusions and Future Work

This chapter summarises the main contributions and findings of this thesis, discusses the limitations encountered during the research and outlines potential directions for future work in joint source and sensor placement optimisation for spatial ANC.

6.1 Conclusions

This thesis has successfully developed and evaluated a novel cost function formulation for simultaneous source and sensor placement optimisation. The research has made contributions to the field while confirming the difficulty of overcoming the fundamental trade-offs in acoustic control system optimisation.

6.1.1 Achievement of Research Goals

The primary research goals established in chapter 1 have been successfully achieved:

Literature Study and Mathematical Foundation

A review of spatial ANC development from Paul Lueg's 1936 patent through modern methods for sound field reproduction was completed. The mathematical framework for pressure-matching-based spatial ANC was established, including the emergence of placement optimisation algorithms.

Environment Setup and Implementation

A complete Python-based environment was implemented for simulating reverberant and free field environments, incorporating realistic acoustic modelling through Sabine's mirror image method [40]. The framework implements state-of-the-art methods for sound field interpolation and synthesis, enabling the implementation of optimisation algorithms for source and sensor placement.

Novel Cost Function Formulation

The core contribution of this thesis is the formulation of a new joint cost function that tries to explicitly address the limitations of existing approaches. The formulation directly compares interpolated and synthesised fields against the true desired field, in an effort to balance interpolation accuracy with synthesis precision.

Validation and Algorithm Comparison

Four distinct placement algorithms were implemented and systematically evaluated. The benchmarking revealed a performance hierarchy and highlighted the trade-offs between performance and computational complexity across multiple resolution configurations.

6.1.2 Key Research Insights

The cost function used with a greedy pair-selection algorithm revealed to produce better performance metrics over the random placement baseline. This came at the cost of computational complexity. The systematic comparison across a range of angles also demonstrated that the cost function correctly incorporated with the statistical framework, enabling lowered performance loss with small to medium system changes.

6.2 Limitations

6.2.1 Methodological Limitations

The research was constrained to single-frequency operation at 500Hz, which limits the direct applicability of results to broadband acoustic control scenarios. Real-world sound fields are expected to exhibit many different

frequencies, requiring the use of broadband optimisation techniques. The thesis was also limited to a two-dimensional environment, excluding computationally expensive three-dimensional geometric effects, but effectively limiting the use of such methods in real-life situations. Finally, the research was limited to a simulation-based verification and validation. This involved simplifications both in the environment configuration and physical sound propagation.

6.2.2 Algorithmic Limitations

While the four implemented algorithms represent different optimisation philosophies, they do not encompass every possible way of solving the problem. Additional algorithms, including hybrid approaches, heuristics such as the **Sparse sensor placement optimisation for classification (SSPOC)** [41], heuristics using the novel cost function, or even machine learning algorithms, could provide different results, with different computational cost vs. performance. The scalability issues are also limited to the extent to which the cost function can be evaluated with the greedy algorithm.

6.2.3 Technical constraints

While multiple regularisation methods have been tested throughout the development, no systematic evaluation of the performance increase was performed. This could mean that the chosen regularisation method may not be the most appropriate, and a deeper focus on this topic could enhance the cost function performances. Many parameters, such as the number of sources and sensors to select, the discretisation step, and initial candidate positions, could have benefited from a more systematic evaluation through simulation, but time constraints greatly limited the extent to which such parameters could be tested.

6.3 Future Work

6.3.1 Immediate Extensions

The first immediate extension to this thesis would be the extension to broadband scenarios. Initial approaches could focus on small frequency ranges with some discrete steps, as previous literature on placement optimisation has done. This would greatly enhance the value of this work, allowing to

verify if the cost function described in this thesis could apply to scenarios closer to reality. The second evident extension would be the extension to three-dimensional simulation. This would significantly enhance the practical relevance of the results, bringing the study one more step closer to a real-life scenario. Finally, physical experimental validation within a control environment could help determine the validity of the cost function in a real scenario.

6.3.2 Algorithmic Development

Having only been really applied to the selection in the greedy algorithm, the cost function could benefit from extended algorithmic testing. These tests could potentially lift or lighten some of the trade-offs discussed in the thesis.

6.4 Reflections

6.4.1 Research Impact and Broader Implications

This research contributes to the advancement of spatial acoustic control technology, which has significant implications for improving the quality of life in built environments. The optimisation framework developed in this thesis has the potential to reduce the number of transducers required for effective noise control, thereby lowering system costs and complexity while maintaining or improving acoustic performance. From a sustainability perspective, the work directly addresses environmental concerns by enabling more efficient acoustic control systems that require fewer material resources and lower energy consumption during operation. The reduction in required hardware components contributes to decreased manufacturing impact and electronic waste generation.

6.4.2 Ethical and Environmental Considerations

The work contributes to several [United Nations \(UN\) Sustainable Development Goal \(SDG\)](#), particularly [SDG 3](#) (Good Health and Well-being) through improved acoustic environments, as seen in [1, 2] and [SDG 9](#) (Industry, Innovation and Infrastructure) through technological advancement in acoustic engineering.

6.4.3 Personal Reflections

The interdisciplinary nature of this research, from acoustics to optimisation theory, software engineering and signal processing, provided valuable insights into the complexity of bridging the gap between physical phenomena and simulated reproductions. The computational challenges encountered in the development of this work highlighted the importance of considering implementation constraints from the early stage of development.

The experience of developing and systematically comparing different algorithms reinforced the value of methodical benchmarking in research and the importance of clearly stating the key metrics of a system to build a solid comparison framework.

References

- [1] H. Ohnishi, K. Uesaka, K. Ohnishi, M. Nishimura, and S. Teranishi, “Development of the noise barrier using actively controlled acoustical soft edge-part 2 : Field test using a loud speaker and a high speed running truck,” 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:201864238> [Pages 1 and 60.]
- [2] INC Engineering Co., Ltd., “Construction sites got silence via a novel noise reduction technique: Construction sites and active noise control,” *IHI Engineering Review*, vol. 48, no. 2, pp. 4–5, 2015, accessed 2025-10-20. [Online]. Available: https://www.ihico.jp/en/technology/techinfo/contents_no/_icsFiles/afieldfile/2023/06/17/ba0b5650119f945ee9832e807f17c7e7.pdf [Pages 1 and 60.]
- [3] Q. Xu, S. Wang, J. Tao, H. Zou, and X. Qiu, “Creating personal sound zones in car cabins with active noise control,” *Applied Acoustics*, vol. 235, p. 110670, May 2025. doi: 10.1016/j.apacoust.2025.110670 [Page 1.]
- [4] G. Chardon, S. Koyama, and L. Daudet, “Numerical evaluation of source and sensor placement methods for sound field control,” December 2020. doi: 10.48465/FA.2020.0748 HAL Id: hal-03053122. [Page 1.]
- [5] P. Lueg, “Process of silencing sound oscillations,” U.S. Patent, Jun. 1936, filed March 8, 1934. [Pages xi, 7, and 8.]
- [6] S. J. Elliott, *Signal processing for active control*, transferred to digital printing 2005 ed., ser. Signal processing and its applications, S. J. Elliott, Ed. San Diego, Calif. [u.a.]: Academic Press, 2005. ISBN 0122370856 Literaturverz. S. [489] - 506. [Online]. Available: <https://www.sciencedirect.com/book/9780122370854/signal-processing-for-active-control> [Pages 8 and 13.]

- [7] S. Kuo and D. Morgan, “Active noise control: a tutorial review,” *Proceedings of the IEEE*, vol. 87, no. 6, pp. 943–975, Jun. 1999. doi: 10.1109/5.763310 [Page 8.]
- [8] S. Elliott, P. Joseph, A. Bullmore, and P. Nelson, “Active cancellation at a point in a pure tone diffuse sound field,” *Journal of Sound and Vibration*, vol. 120, no. 1, pp. 183–189, Jan. 1988. doi: 10.1016/0022-460x(88)90343-4 [Page 9.]
- [9] P. Joseph, S. Elliott, and P. Nelson, “Near field zones of quiet,” *Journal of Sound and Vibration*, vol. 172, no. 5, pp. 605–627, May 1994. doi: 10.1006/jsvi.1994.1202 [Page 9.]
- [10] P. A. Nelson and S. J. Elliott, *Active control of sound*, paperback ed., 4. print., [nachdr.] ed., P. A. Nelson, Ed. San Diego [u.a.]: Academic Press, Nov. 2000. ISBN 0125154259 Literaturverz. S. 421 - 432. [Pages 9, 10, and 11.]
- [11] J. Daniel, “Further investigations of high-order ambisonics and wavefield synthesis for holophonic sound imaging,” in *Audio Engineering Society Convention 114*. Audio Engineering Society, 2003. [Page 9.]
- [12] M. A. Poletti, “Three-dimensional surround sound systems based on spherical harmonics,” *Journal of the Audio Engineering Society*, vol. 53, no. 11, pp. 1004–1025, 2005. [Page 9.]
- [13] B. Rafaely, *Fundamentals of Spherical Array Processing*. Berlin, Heidelberg: Springer International Publishing, 2019. ISBN 9783319995618 [Page 9.]
- [14] E. G. Williams and J. A. Mann, “Fourier acoustics: Sound radiation and nearfield acoustical holography,” *The Journal of the Acoustical Society of America*, vol. 108, no. 4, pp. 1373–1373, Oct. 2000. doi: 10.1121/1.1289662 [Pages 9 and 11.]
- [15] O. Kirkeby, P. A. Nelson, F. Orduna-Bustamante, and H. Hamada, “Local sound field reproduction using digital signal processing,” *The Journal of the Acoustical Society of America*, vol. 100, no. 3, pp. 1584–1593, Sep. 1996. doi: 10.1121/1.416060. [Online]. Available: <https://doi.org/10.1121/1.416060> [Page 10.]
- [16] N. Ueno, S. Koyama, and H. Saruwatari, “Three-dimensional sound field reproduction based on weighted mode-matching method,” *IEEE/ACM*

- Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 12, pp. 1852–1867, Dec. 2019. doi: 10.1109/taslp.2019.2934834 [Page 10.]
- [17] S. Koyama, K. Kimura, and N. Ueno, “Weighted pressure and mode matching for sound field reproduction: Theoretical and experimental comparisons,” *arXiv preprint arXiv:2303.13027*, Mar. 2023. doi: 10.48550/ARXIV.2303.13027. [Online]. Available: <https://arxiv.org/abs/2303.13027> [Pages 10, 11, and 28.]
- [18] S. Koyama, G. Chardon, and L. Daudet, “Optimizing source and sensor placement for sound field control: An overview,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 696–714, 2020. doi: 10.1109/taslp.2020.2964958. [Online]. Available: <https://ieeexplore.ieee.org/document/8952606> [Pages 11, 12, 13, and 52.]
- [19] S. Joshi and S. Boyd, “Sensor selection via convex optimization,” *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, Feb. 2009. doi: 10.1109/tsp.2008.2007095 [Page 11.]
- [20] P. M. Morse, K. U. Ingard, and R. T. Beyer, “Theoretical acoustics,” *Journal of Applied Mechanics*, vol. 36, no. 2, pp. 382–383, Jun. 1969. doi: 10.1115/1.3564682 [Pages 11 and 12.]
- [21] J. F. Allard and N. Atalla, *Propagation of Sound in Porous Media: Modelling Sound Absorbing Materials*, 2nd ed. Chichester, UK: Wiley, Oct. 2009. ISBN 9780470747339 [Page 11.]
- [22] K. Kimura, S. Koyama, N. Ueno, and H. Saruwatari, “Mean-square-error-based secondary source placement in sound field synthesis with prior information on desired field,” pp. 1552–1556, 2021. [Pages 12, 13, 28, and 30.]
- [23] S. Kozuka, S. Koyama, H. Ito, and N. Kamado, “Source and sensor placement for sound field control based on mean square error with prior spatial covariance,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2025, submitted. [Pages 12, 13, 29, and 32.]
- [24] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*, H. L. V. Trees, Ed. New York,

- NY: Wiley, Mar. 2002. ISBN 9780471221104 Includes bibliographical references and index. - Online resource; title from PDF title page (Wiley, viewed May 5, 2017). [Page 12.]
- [25] N. Ueno, S. Koyama, and H. Saruwatari, “Directionally weighted wave field estimation exploiting prior information on source direction,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 2383–2395, 2021. doi: 10.1109/tsp.2021.3070228 [Page 13.]
- [26] S. Koyama, G. Chardon, and L. Daudet, “Joint source and sensor placement for sound field control based on empirical interpolation method,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, Apr. 2018. doi: 10.1109/icassp.2018.8461597 pp. 501–505. [Online]. Available: <https://ieeexplore.ieee.org/document/8461597> [Page 13.]
- [27] S. A. Verburg, F. Elvander, T. van Waterschoot, and E. Fernandez-Grande, “Optimal sensor placement for the spatial reconstruction of sound fields,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2024, no. 1, Aug. 2024. doi: 10.1186/s13636-024-00364-4 [Page 13.]
- [28] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, “An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations,” *Comptes Rendus. Mathématique*, vol. 339, no. 9, pp. 667–672, Oct. 2004. doi: 10.1016/j.crma.2004.08.006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1631073X04004248> [Page 13.]
- [29] L. Lu, K.-L. Yin, R. C. de Lamare, Z. Zheng, Y. Yu, X. Yang, and B. Chen, “A survey on active noise control in the past decade—part ii: Nonlinear systems,” *Signal Processing*, vol. 181, p. 107929, Apr. 2021. doi: 10.1016/j.sigpro.2020.107929 [Page 13.]
- [30] X. Li, C. Lu, W. Chen, Y. Zhu, and C. Cheng, “Research on fast optimal reference sensor placement in active road noise control,” *Noise Control Engineering Journal*, vol. 71, no. 5, pp. 315–331, Sep. 2023. doi: 10.3397/1/377126 [Page 13.]
- [31] P. C. Hansen and L. Eldén, “Understanding the svd for solving linear least squares problems,” *SIAM Review*, vol. 39, no. 2, pp. 206–222, 1997. doi: <https://doi.org/10.1137/1.9781611971217> [Page 13.]

- [32] T. Nishida, N. Ueno, S. Koyama, and H. Saruwatari, “Region-restricted sensor placement based on gaussian process for sound field estimation,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 1718–1733, 2022. doi: 10.1109/tsp.2022.3156012 [Page 17.]
- [33] E. B. Holmberg, R. E. Hillman, and J. S. Perkell, “Glottal airflow and transglottal air pressure measurements for male and female speakers in low, normal, and high pitch,” *Journal of Voice*, vol. 3, no. 4, pp. 294–305, Dec. 1989. doi: 10.1016/s0892-1997(89)80051-7. [Online]. Available: [https://www.jvoice.org/article/S0892-1997\(89\)80051-7/abstract](https://www.jvoice.org/article/S0892-1997(89)80051-7/abstract) [Page 21.]
- [34] S. Koyama, J. Brunnstrom, H. Ito, N. Ueno, and H. Saruwatari, “Spatial active noise control based on kernel interpolation of sound field,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3052–3063, 2021. doi: 10.1109/taslp.2021.3107983 [Page 30.]
- [35] J. Ahrens and S. Spors, “An analytical approach to sound field reproduction using circular and spherical loudspeaker distributions,” *Acta Acustica united with Acustica*, vol. 94, no. 6, pp. 988–999, Nov. 2008. doi: 10.3813/aaa.918115 [Page 30.]
- [36] K. Kimura, S. Koyama, and H. Saruwatari, “Perceptual quality enhancement of sound field synthesis based on combination of pressure and amplitude matching,” 2023. [Page 30.]
- [37] K. Manohar, B. W. Brunton, J. N. Kutz, and S. L. Brunton, “Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns,” *IEEE Control Systems*, vol. 38, no. 3, pp. 63–86, Jun. 2018. doi: 10.1109/mcs.2018.2810460 [Page 43.]
- [38] B. Lam, W.-S. Gan, D. Shi, M. Nishimura, and S. Elliott, “Ten questions concerning active noise control in the built environment,” *Building and Environment*, vol. 200, p. 107928, Aug. 2021. doi: 10.1016/j.buildenv.2021.107928. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0360132321003322> [Page 52.]
- [39] M. Bouchard, “Multichannel affine and fast affine projection algorithms for active noise control and acoustic equalization systems,” *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 1, pp. 54–60, Jan. 2003. doi: 10.1109/tsa.2002.805642 [Page 52.]

- [40] W. C. Sabine, *Collected papers on acoustics*, unabridged and unaltered republication ed. Los Altos, Calif.: Peninsula Publ., 1992. ISBN 0932146600 [Page 58.]
- [41] B. W. Brunton, S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Sparse sensor placement optimization for classification,” *SIAM Journal on Applied Mathematics*, vol. 76, no. 5, pp. 2099–2122, Jan. 2016. doi: 10.1137/15m1036713 [Page 59.]

Appendix A

Software architecture details

A.1 Software Architecture Implementation

This appendix documents the technical implementation of the simulation framework, including component specifications and implementation details. The complete system architecture is shown in Fig. [A.1](#).

A.1.1 System Architecture

The simulation framework implements a four-layer architecture designed for efficient acoustic field control optimisation. The architecture enables both hierarchical communication through all layers and direct cross-layer access between the algorithm and matrix computation layers during iterative optimisation, bypassing the domain layer when performance is critical.

The **Configuration Management layer** handles system-wide parameters through a `ConfigManager` class that validates and distributes settings from TOML files across the entire system. Configuration sections cover acoustic parameters, room geometry, algorithm settings, matrix computation options, and visualisation preferences.

The **Domain Modelling layer** manages the acoustic environment state through three primary classes. The `Room` class maintains the acoustic environment and stores precomputed matrices in memory, while the `ControlZone` class handles candidate position management for sources and sensors. The `PlaneWaveField` class implements acoustic field modelling and integrates with `Pyroomacoustics` through the `Room.mirror_image` attribute for image source computation.

The **Algorithm Implementation layer** provides a standardised in-

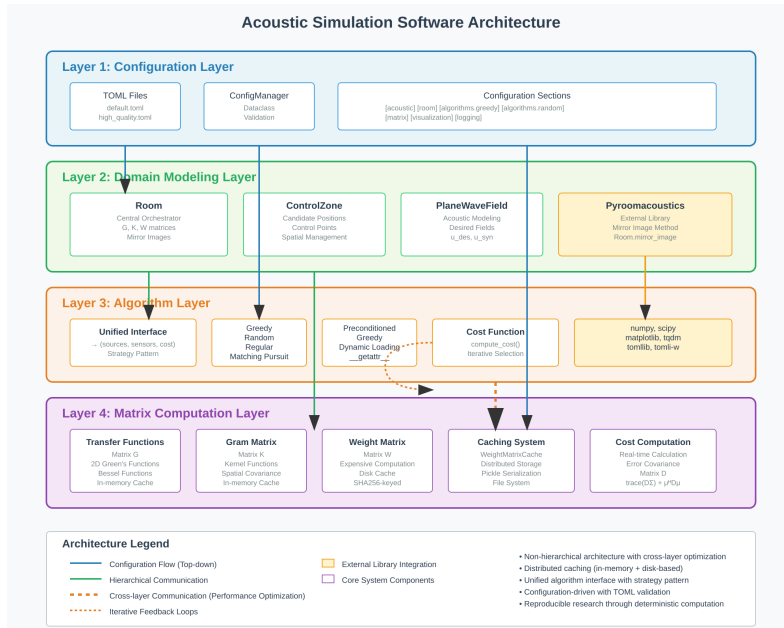


Figure A.1: Software architecture diagram

interface where all algorithms follow the pattern `algorithm(room, num_sources, num_sensors) → (sources, sensors, cost)`. Algorithms are dynamically loaded and include implementations for Greedy selection, Random selection, Regular placement, Matching Pursuit, and Preconditioned Greedy methods. This layer accesses matrix computations directly for performance during optimisation loops.

The **Matrix Computation layer** handles the computationally intensive operations, including transfer function computation using 2D Green's functions, Gram matrix computation with kernel functions, and weight matrix computation with persistent disk caching. The `WeightMatrixCache` uses SHA256-based storage for efficient retrieval of expensive computations.

A.1.2 Data Flow and Execution

Configuration flows from TOML files through the `ConfigManager` to all system layers simultaneously. Standard execution begins with configuration loading and validation, followed by room and control zone instantiation. The system then performs matrix precomputation through `Room.compute_matrices()` before algorithm initialisation and iterative optimisation with cost evaluation.

The cost function implementation follows the mathematical formulation

where the error metric $Q(L, M) = \text{trace}(D\Sigma) + \mu^H D \mu$ is computed using transfer function submatrices for selected sources and sensors. Adaptive regularisation maintains numerical stability based on condition numbers of the relevant matrices, with parameters $\lambda = 10^{-4} \times \text{cond}(K_{cc})$ for Gram matrix regularisation and $\eta = 10^{-4} \times \text{cond}(A)$ for synthesis regularisation.

A.1.3 Caching Strategy

The framework employs a two-tier caching system. Memory-based caching maintains transfer function matrices, Gram matrices, and spatial covariance matrices in the domain layer for rapid access during optimisation. Disk-based caching provides persistent storage for expensive weight matrix computations using binary serialisation and SHA256-based cache keys derived from configuration parameter hashing.

A.1.4 External Integration

The system integrates with Pyroomacoustics at the domain layer for mirror image computation and acoustic modelling. Scientific computing relies on NumPy for array operations, SciPy for specialised functions like Bessel functions in Green's function computation, and Matplotlib for visualisation capabilities. The framework also utilises TQDM for progress monitoring and TOML libraries for configuration management.

A.1.5 Performance and Reproducibility

Performance optimisation centres on matrix precomputation during initialisation and vectorised operations throughout the computation pipeline. Cross-layer matrix access avoids unnecessary data copying while maintaining clear architectural boundaries. Reproducibility is ensured through deterministic operations with fixed random seeds, consistent numerical precision, and comprehensive parameter logging for result verification.

Appendix B

Simulation configurations

This appendix only shows one configuration file for brevity. All configuration files and all the code produced for this thesis can be found on <https://github.com/skoyamalab/SourceSensorPlacementSFC>

Listing B.1: Default configuration details

```
# Default Configuration for Acoustic Control Simulation
# This file contains all configurable parameters for the
  acoustic source and sensor placement optimization

[acoustic]
# Sound speed in m/s
sound_speed = 340.29
# Frequency for narrowband analysis in Hz
frequency = 500
# Noise source direction in radians  $\pi(-/4 = -45$  degrees)
noise_angle = -0.7853981633974483 #  $\pi-/4$ 
# Phase shift for signal generation
phase_shift = 0.0

[room]
# Room position and size (x_min, y_min, x_size, y_size)
position = [-2.5, -2.0]
size = [5.0, 4.0]
# Spatial resolution for discretization in meters
resolution = 0.05

[room.acoustics]
# Enable room acoustics (mirror image method) - false
```

```
    means free-field simulation
enable_room_acoustics = true
# Energy absorption coefficient (0.0 = no absorption,
  1.0 = full absorption)
e_absorption = 0.2
# Maximum reflection order for mirror image computation
max_order = 10

[control_zone]
# Control zone center coordinates (x, y)
center = [0.5, 0.3]
# Control zone width/diameter in meters
width = 1.0
# Sensor candidate selection strategy and parameters
sensor_strategy = "square"
sensor_layers = 4

[sources]
# Source candidate file path
candidate_file = "posCandidateSources.npy"
# Offset to apply to source positions (x, y)
offset = [0.5, 0.0]

[optimization]
# Number of sources to select (L)
num_sources = 5
# Number of sensors to select (M)
num_sensors = 5

[algorithms]
# Algorithm-specific parameters

[algorithms.greedy]
# Maximum iterations (0 = unlimited)
max_iterations = 0
# Show selection progress
show_selection = true

[algorithms.preconditioned_greedy]
# Preconditioning method: "correlation", "euclidean", "
  cosine", "magnitude_weighted"
method = "correlation"
```

74 | Appendix B: Simulation configurations

```
# Similarity threshold for grouping (0-1)
similarity_threshold = 0.95
# Maximum fraction of sources to remove (0-1)
max_source_reduction = 0.5
# Maximum fraction of sensors to remove (0-1)
max_sensor_reduction = 0.3
# Enable source preconditioning
enable_source_preconditioning = true
# Enable sensor preconditioning
enable_sensor_preconditioning = false
# Maximum iterations for greedy phase
max_iterations = 50

[algorithms.random]
# Number of runs to perform (best result selected)
num_runs = 100
# Maximum attempts per run to find valid selection
max_attempts = 10000

[cost_function]
# Regularization parameters
lambda_reg = 1e-4
eta_factor = 1e-11

[output]
# Output directory for results
output_dir = "results/default"
# Whether to save figures
save_figures = true
# Whether to show figures
show_figures = false
# Figure format: "png", "pdf", "svg"
figure_format = "png"
# Figure DPI
figure_dpi = 150
# Whether to save detailed reports
save_reports = true
# Report format: "txt", "md"
report_format = "md"
# Logging level: "DEBUG", "INFO", "WARNING", "ERROR"
log_level = "DEBUG"
```

```
[visualization]
# Colormap for field plots
colormap = "RdBu_r"
# Figure size (width, height) in inches
figure_size = [12, 8]
# Show grid on plots
show_grid = true
# Alpha transparency for plots
alpha = 0.7
```

Appendix C

Full Sized Figures

This appendix contains the figures from the report in a larger size for better legibility.

C.1 Algorithm plots

C.1.1 Greedy Algorithm Plots

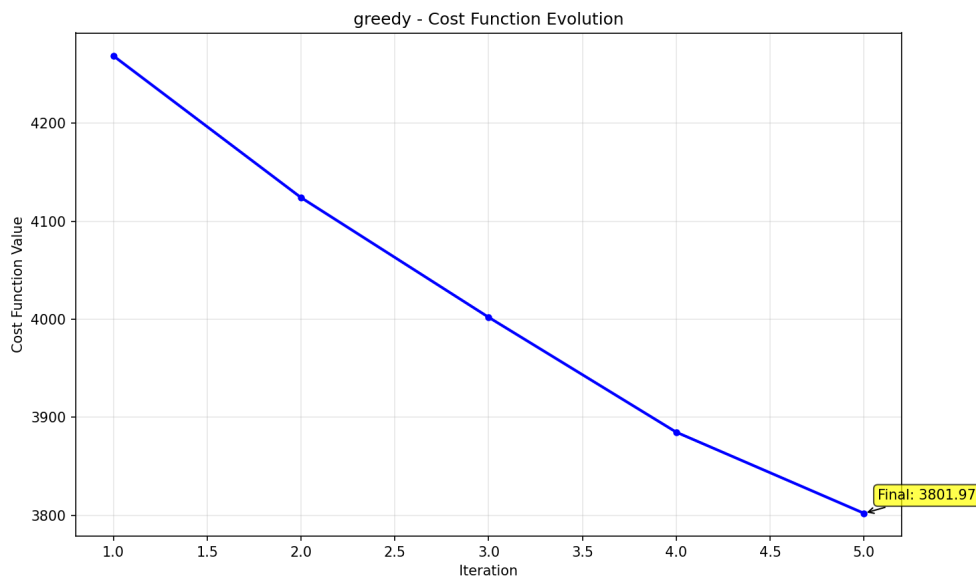


Figure C.1: Greedy algorithm convergence in default configuration. 5.2a

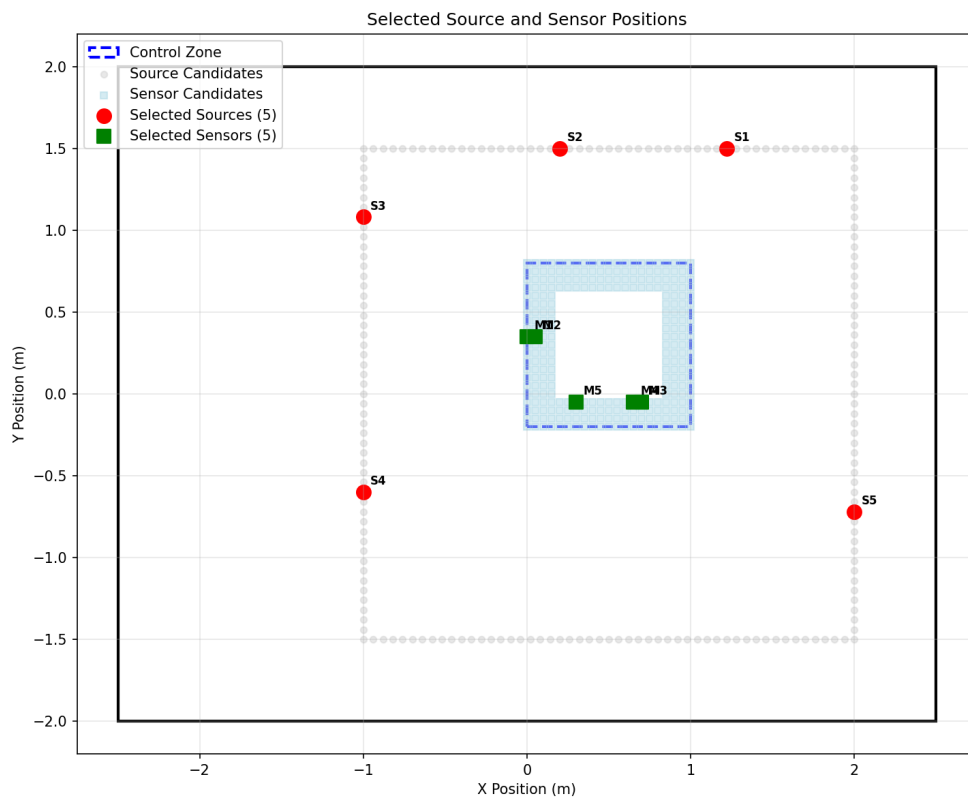


Figure C.2: optimised source and sensor positions (Greedy, default). 5.2b




Figure C.3: Acoustic field visualisation (Greedy, default). 5.2c

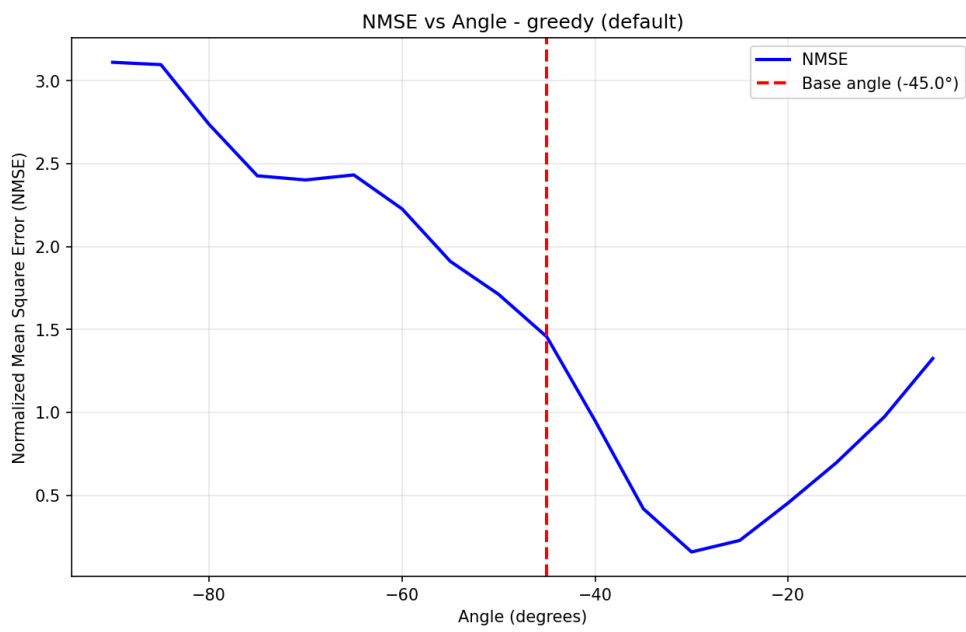


Figure C.4: NMSE evaluation (Greedy, default). 5.2d

C.1.2 Regular Algorithm Plots




Figure C.5: optimised source and sensor positions (Regular, default). 5.3a




Figure C.6: Acoustic field visualisation (Regular, default). 5.3b

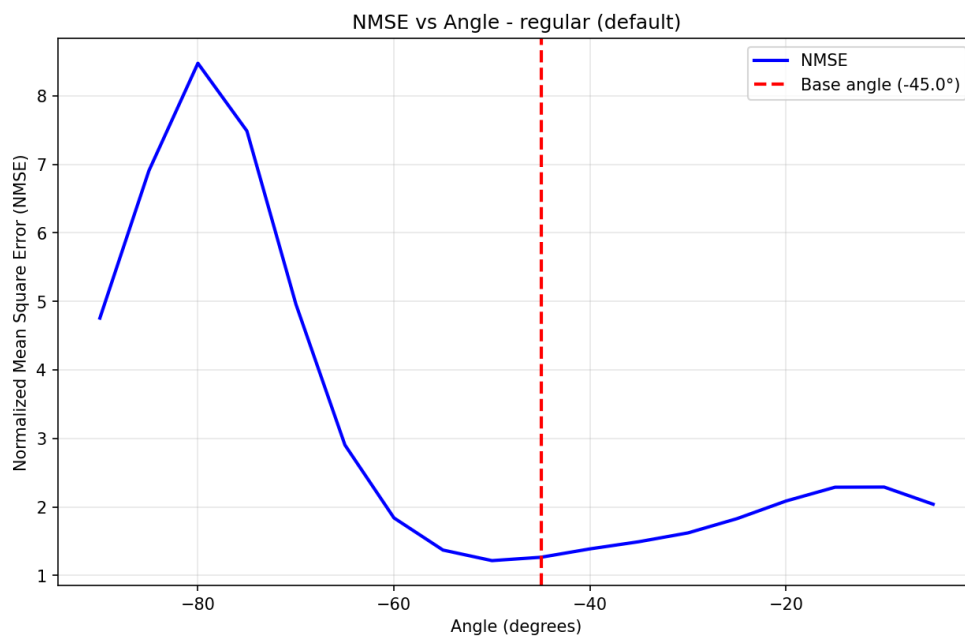


Figure C.7: NMSE evaluation (Regular, default). 5.3c

C.1.3 Matching Pursuit Plots




Figure C.8: optimised source and sensor positions (Matching Pursuit, default). 5.4a




Figure C.9: Acoustic field visualisation (Matching Pursuit, default). 5.4b

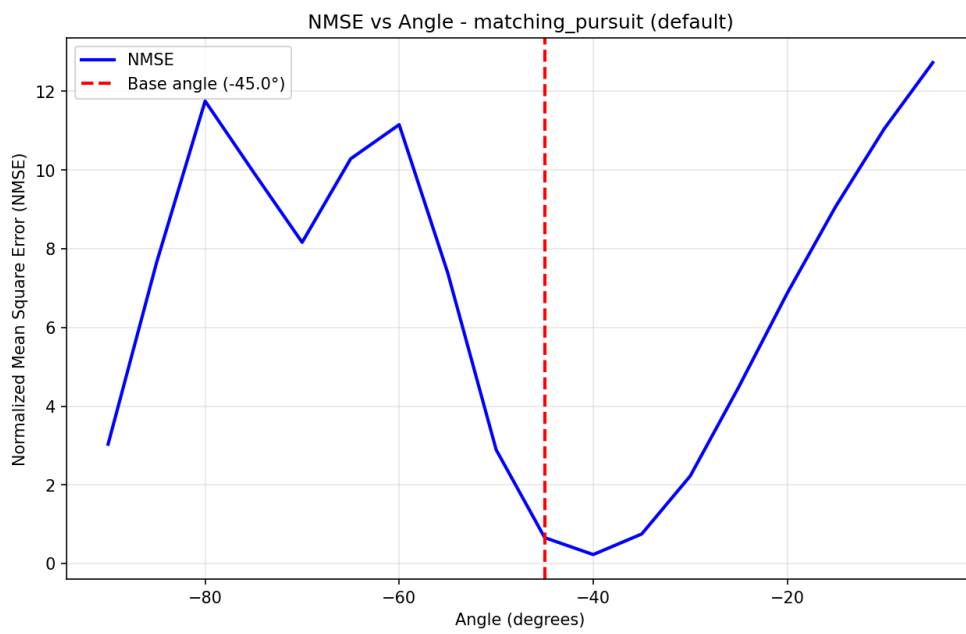


Figure C.10: NMSE evaluation (Matching Pursuit, default). 5.4c

