



HAL
open science

Extending standard outdoor noise propagation models to complex geometries

Matthew Kamrath

► **To cite this version:**

Matthew Kamrath. Extending standard outdoor noise propagation models to complex geometries. Acoustics [physics.class-ph]. Université du Maine, 2017. English. NNT: 2017LEMA1038 . tel-01834558v1

HAL Id: tel-01834558

<https://theses.hal.science/tel-01834558v1>

Submitted on 10 Jul 2018 (v1), last revised 13 Jul 2018 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Thèse de Doctorat

Matthew KAMRATH

*Mémoire présenté en vue de l'obtention du
grade de Docteur de Le Mans Université
sous le sceau de l'Université Bretagne Loire*

École doctorale : Sciences pour l'Ingénieur

Discipline : 823 (Acoustique)

Spécialité : Acoustique

Unité de recherche : *Institut Français des Sciences et Technologies
des Transports de l'Aménagement et des Réseaux*

Soutenu : le 28 septembre 2017

Thèse N° : 2017LEMA1038

Extending standard outdoor noise propagation models to complex geometries

Jury

Rapporteurs :	Denis DUHAMEL , Professeur, Ecole Nationale des Ponts et Chaussées Jens FORSSÉN , Associate Professor, Chalmers University of Technology
Examineurs :	Fabienne ANFOSSO-LÉDÉE , Ingénieur Divisionnaire des Travaux Public de l'Etat, HDR, Ifsttar Régis MARCHIANO , Professeur d'université, UPMC Pierre-Olivier MATTEI , Chargé de recherche, HDR, CNRS, LMA Olivier RICHOUX , Professeur d'université, Université du Maine
Invités :	Fabrice JUNKER , Ingénieur, EDF Julien MAILLARD , Chercheur, CSTB
Directeur de Thèse :	Judicaël PICAUT , Directeur de recherche, Ifsttar
Co-directeur de Thèse :	Philippe JEAN , Chercheur, HDR, CSTB

Extending standard outdoor noise propagation models to complex geometries

Doctoral Dissertation in Acoustics
Defended September 28, 2017

Matthew Kamrath

Abstract

Noise engineering methods (e.g. ISO 9613-2 or CNOSSOS-EU) efficiently approximate sound levels from roads, railways, and industrial sources in cities. They model outdoor sound propagation by first finding the propagation paths between the source and receiver, then applying attenuations (e.g. geometrical divergence and atmospheric absorption) to each path, and finally summing all of the path contributions.

However, engineering methods are limited to only simple box-shaped geometries; for example, they cannot model a T-barrier. This dissertation develops and validates a hybrid method to extend the engineering methods to more complicated geometries by introducing an extra attenuation term that represents the influence of a real object compared to a simplified object that the engineering methods can model.

Calculating the extra attenuation term requires reference calculations to quantify the difference between the complex and simplified objects. Since performing a reference computation for each path is too computationally expensive, the extra attenuation term is linearly interpolated from a data table containing the corrections for many source and receiver positions and frequencies. The 2.5D boundary element method produces the levels for the real complex geometry and a simplified geometry, and subtracting these levels yields the corrections in the table.

Forming the table of corrections has a natural tradeoff between accuracy and computational cost; increasing the number of data points in the table decreases the interpolation error but increases the time and memory requirements. The computational restrictions of a standard desktop computer only allow the table to have about one million points, which is far below the number required to densely cover the large six-dimensional space. Nevertheless, minimizing the variance of a Gaussian process optimizes the source and receiver positions in the data table and reduces the interpolation error to an acceptable level.

This dissertation validates this hybrid method for a T-barrier with hard ground, soft ground, and buildings. All three cases demonstrate that the hybrid method is more accurate than standard engineering methods for complex cases; it reduces the mean error by approximately 2 dBA. In a final case, the hybrid method produces the results for a T-barrier in a much larger scene (i.e. 180 m x 80 m and frequencies up to 5 kHz) with buildings and mixed hard and soft ground. This case is too complex for ordinary noise engineering methods because of the T-shaped barrier and is cost prohibitive for reference methods because of the large domain size compared to a wavelength. Thus, the hybrid method provides a novel tradeoff between accuracy and cost.

Key words: urban outdoor sound propagation, noise engineering methods, boundary element method, hybrid method, Gaussian process regression, numerical acoustics, optimization

Résumé

Les méthodes d'ingénierie acoustique (e.g. ISO 9613-2 ou CNOSSOS-EU) approchent efficacement les niveaux de bruit générés par les routes, les voies ferrées et les sources industrielles en milieu urbain. Ces méthodes modélisent la propagation du son en milieu extérieur en recherchant les chemins de propagation entre source(s) et récepteur(s), puis en appliquant des atténuations (e.g. divergence géométrique et absorption atmosphérique) à chaque chemin, et enfin en sommant les contributions de tous les chemins.

Cependant, ces approches d'ingénierie sont limitées à des géométries de forme simple, le plus souvent de section rectangulaire ; elles ne peuvent pas, par exemple, modéliser des écrans acoustiques de forme en T. Ce mémoire développe donc, et valide, une approche hybride permettant l'extension des méthodes d'ingénierie à des formes plus complexes, en introduisant un terme d'atténuation supplémentaire qui représente l'effet d'un objet réel comparé à un objet simple de référence pouvant être appréhendé par les modèles d'Ingénierie.

Le calcul de cette atténuation supplémentaire nécessite des calculs de référence, permettant de quantifier la différence entre objets simple et complexe. Dans la mesure, où il est trop onéreux, numériquement, d'effectuer ce calcul pour tous les chemins de propagation, l'atténuation supplémentaire est obtenue par interpolation de données stockées dans un tableau et évaluées pour un large jeu de positions de sources, de récepteurs et de fréquences. Dans notre approche, les calculs de référence utilisent la méthode BEM en 2.5D, et permet ainsi de produire les niveaux de référence pour les géométries simple et complexe, tout en tabulant leur écart. Sur le principe, d'autres approches de référence pourraient être utilisées.

La construction de cette table corrective doit être un compromis entre précision et coût de calcul ; si l'on accroît le nombre de valeurs stockées, l'erreur d'interpolation sera réduite mais le prix à payer en termes de temps de calcul et de capacité mémoire peut devenir excessif. Les limites données pour un PC standard ne permettent guère de dépasser un million de valeurs, ce qui est bien insuffisant pour assurer une interpolation précise dans un espace à 6 dimensions. Néanmoins, en minimisant la variance d'un processus Gaussien on parvient à optimiser les positions des sources et des récepteurs, et par là-même, à réduire l'erreur d'interpolation à un niveau acceptable.

Ce travail valide cette approche hybride pour un écran en forme de T avec un sol rigide, un sol absorbant et un cas avec bâtiments. Ces trois cas démontrent que l'approche hybride est plus précise que l'approche d'ingénierie standard dans des cas complexes ; elle réduit l'erreur moyenne d'approximativement 2 dB(A). A l'aide d'un exemple plus complexe, nous montrons que l'approche hybride permet par exemple d'étudier un écran en forme de T placé dans une scène urbaine étendue (i.e. 180 m x 80 m) et pour des fréquences atteignant 5000 Hz. Ce cas serait trop complexe à mettre en œuvre, d'une part, avec les approches standards du fait du type d'écran considéré, et d'autre part, avec les approches de référence, de par ses dimensions relatives aux nombres d'ondes considérés. De cette façon, l'approche hybride propose un compromis entre précision et coût.

Mots-clés : Propagation sonore en milieu urbain, méthodes d'ingénierie acoustique, BEM, méthode hybride, processus Gaussien de régression, acoustique numérique, optimisation

Contents

Abstract	v
Résumé	vii
Contents.....	ix
Acknowledgements	xiii
Chapter 1 Introduction	1
§1.1 Objectives.....	1
§1.2 Context	1
§1.3 Literature review	3
§1.3.1 Geometrical approaches.....	3
§1.3.2 Engineering approaches	4
§1.3.3 Statistical approaches.....	5
§1.3.4 Frequency-domain approaches	6
§1.3.5 Time-domain approaches.....	7
§1.3.6 Hybrid approaches	8
§1.4 Model requirements.....	9
§1.5 Choosing a method.....	10
§1.5.1 Justifying using EMs.....	11
§1.5.2 Justifying using 2.5D BEM.....	12
§1.6 Thesis organization	13
§1.7 Summary.....	15

Chapter 2 Hybrid Method	17
§2.1 Objectives.....	17
§2.2 Method outline	18
§2.3 Engineering methods	20
§2.3.1 Scene setup.....	21
§2.3.2 Finding propagation paths.....	22
§2.3.3 Calculating the attenuations	23
§2.3.4 Summing the path contributions	26
§2.3.5 The hybrid method’s adjustment	26
§2.4 Precomputations	27
§2.4.1 Variables	27
§2.4.2 Deciding the source and receiver positions	30
§2.4.3 Choosing the frequencies and geometries.....	32
§2.4.4 Using flat, hard ground	33
§2.4.5 Post-processing	35
§2.5 Local source and receiver positions.....	38
§2.5.1 Single diffraction	38
§2.5.2 Multiple diffraction.....	40
§2.5.3 Reflection.....	42
§2.5.4 Lateral diffraction	43
§2.6 Interpolating the table of corrections.....	44
§2.6.1 Choosing a data structure.....	44
§2.6.2 Choosing an interpolation method	48
§2.6.3 Linear interpolation on a non-uniform grid	49
§2.7 Extrapolating the table of corrections.....	52
§2.8 Summary.....	57
Chapter 3 Validation	59
§3.1 Objectives.....	59
§3.2 Table of corrections	60
§3.3 T-barrier with hard ground.....	62
§3.3.1 Description.....	62
§3.3.2 Results.....	63
§3.3.3 Analysis.....	66
§3.4 T-barrier with soft ground.....	67
§3.4.1 Description.....	68
§3.4.2 Results.....	68

§3.4.3 Analysis.....	70
§3.5 T-barrier with buildings.....	71
§3.5.1 Description.....	71
§3.5.2 Results.....	73
§3.5.3 Analysis.....	74
§3.6 Real Scene.....	75
§3.6.1 Description.....	76
§3.6.2 Results.....	77
§3.7 Summary.....	79

Chapter 4 Optimization 81

§4.1 Objectives.....	81
§4.2 Introduction.....	81
§4.2.1 General optimization.....	82
§4.2.2 Objective functions.....	82
§4.2.3 Numerical optimization methods.....	83
§4.2.4 Convexity.....	84
§4.3 Point distribution options.....	85
§4.3.1 Linear distribution.....	85
§4.3.2 Exponential distribution.....	85
§4.3.3 Optimizing the distribution.....	86
§4.4 Gaussian processes.....	88
§4.4.1 Introduction.....	88
§4.4.2 Covariance function.....	89
§4.4.3 Calculations.....	91
§4.5 Determining the hyper-parameters.....	93
§4.5.1 Sample data.....	93
§4.5.2 The objective function.....	95
§4.5.3 Constraints.....	97
§4.5.4 Optimization method.....	97
§4.5.5 Improving the local results.....	97
§4.6 Minimizing the variance.....	97
§4.6.1 Objective function.....	97
§4.6.2 Constraints.....	98
§4.6.3 Optimization method.....	99
§4.7 Choosing N.....	99
§4.8 Comparison.....	100

§4.8.1 Description.....	101
§4.8.2 Results.....	104
§4.8.3 Analysis.....	106
§4.9 Summary.....	108
Chapter 5 Conclusion.....	109
§5.1 Objectives.....	109
§5.2 Principal results	109
§5.3 Novelty and importance	111
§5.4 Future work.....	112
§5.5 Summary.....	113
Extended summary	115
Résumé étendu	121
Bibliography.....	127
Publications.....	139

Acknowledgements

My high school physic teacher, Arlyn DeBruyckere, made me understand the power of symbolic manipulation. Before that class, I would write in all of the numbers that I knew in the equation, but he showed me a much more powerful way to do physics, so my interest grew. Thank you.

The summer after high school, I was an archery instructor at Many Point Scout Camp and the class asked me what I wanted to do in the future, so I said I wanted to earn a Ph.D. They thought referring to me as Dr. Kamrath was quite entertaining, so they proceeded to call me Dr. Kamrath all week. Thank you.

At St. Olaf College as I was playing saxophone in the band and studying math and physics, I discovered my passion for acoustics. My advisor, David Nitz, encouraged me to learn more about acoustics and even helped me do an independent study in acoustics using Kinsler and Frey's *Fundamentals of Acoustics*. Thank you.

When I arrived at Penn State, Daniel Russell taught my introduction to vibrations class with clarity and passion; he strived to do his best to teach us about acoustics. Once a week, we would have night classes, and after class he would often offer to drop me off at my home. He was genuinely interested in helping me succeed. Thank you.

During my first year at Penn State, I often studied with James Esplin for our introductory classes and the Ph.D. candidacy exam. I learned much more because we studied together and asked each other many difficult questions. Thank you.

While I was at Penn State, I also started attending the Acoustical Society of America (ASA) conferences. The ASA was very supportive of students attending the conferences by providing student travel subsidies and reduced registration. There were also many events specifically designed for students. The ASA conferences opened my mind to the

Acknowledgements

breadth of acoustics and allowed me to meet many new friends and legends in the field of acoustics. Thank you.

I also worked at KCF Technologies with my academic advisor Gary Koopmann and supervisor Mike Grissom. That internship gave me the opportunity to apply what I was learning in my classes and work on wide variety of topics including my master's research. Thank you.

During my second year at Penn State, I was also a graduate research assistant studying the effects of aircraft noise on children with my advisor Michelle Vigeant. She was very supportive and strived to create a welcoming and fun laboratory environment. This research taught me how to read scientific literature and developed my statistical skills. Thank you.

As I was finishing my MS in acoustics at Penn State, I began searching for a good Ph.D. project. Although I searched for opportunities at Penn State, I was not finding the right fit, so I started to look elsewhere. When the CSTB advertised a Ph.D. position modeling urban outdoor sound propagation, I was excited and started learning C++ and French before I was even accepted (both were brand new to me). I knew a little about the CSTB because I saw Alexandre Jolibois defend his Ph.D. at Penn State but not much. Victor Sparrow helped me decide what I should do. Thank you.

The CSTB funded this Ph.D. project and was very supportive. They helped me attend my first Inter-noise conference in Hamburg, Germany and sent me to the Boston ASA meeting. Through the CSTB, I had many advisors and stood on the shoulders of giants. Dirk Van Maercke and Jérôme Defrance provided excellent technical advice about the engineering methods. Julien Maillard was my co-advisor and helped me develop the C++ code. Philippe Jean was also my co-advisor and helped me run the BEM computations. They were both easily available and provided excellent advice. Thank you.

Judicaël Picaut was my academic advisor and really pushed me to improve the dissertation, clarify my thoughts, and do something novel. He knew what to do to make the dissertation better. Thank you.

Christophe Langrenne provided the FM-BEM reference calculations, which significantly added to the validation of the hybrid method. Without his contribution, we could only go up to 400 Hz, but with it we could go up to 1.6 kHz, which covers the most important frequencies. Thank you.

Fabienne Anfosso-Lédée, Pierre-Olivier Mattei, and Régis Marchiano were members of the jury and brought new perspectives to the project. Olivier Richoux and Fabrice Junker served on my thesis committee and jury and asked many excellent questions that helped shape the direction of the dissertation. Denis Duhamel and Jens Forssén were my

reviewers, served on the jury, and provided excellent feedback that significantly improved this dissertation. Thank you.

Finally, my friends and family have provided steadfast love and support. Without them, I would have never come this far. In high school, I had never used a programming language, but now I have used several and feel like I could learn any programming language. In high school, I was not allowed to drive in the twin cities, but now I live in France and can travel throughout Europe. My development is a testament to my friends and family. Thank you.

I have not listed everyone that has helped me along the way but know that your contributions, support, and encouragement are appreciated and never forgotten. There are so many people that have positively impacted my life and made this achievement possible. Thank you.

Chapter 1 Introduction

§1.1 Objectives

To aid the reader, each chapter begins with a set of chapter objectives and ends with a synthesis of the most important ideas in the chapter. In this chapter,

- §1.2 explains the problem context;
- §1.3 reviews current approaches;
- §1.4 defines what is required to solve the problem;
- §1.5 states and justifies the method that this dissertation develops;
- §1.6 lays out the organization of this dissertation.

§1.2 Context

Many people are exposed to high noise levels from cars, trains, planes, or industrial sources. For example, 120 million Europeans are exposed to L_{den} above 55 dB¹, which the World Health Organization (WHO) says produces “serious annoyance”². Moreover, these high noise levels contribute to increased disease and lower quality of life; specifically, high noise levels may negatively affect annoyance, cardiovascular disease, sleep, and cognitive performance³. WHO and the Joint Research Centre (JRC)⁴ estimate that environmental noise

- Annoys one in three Europeans during the day;
- Disturbs the sleep of one in five Europeans;
- Robs Europeans of over 1 million healthy years of life annually.

Accurately and efficiently modeling transportation noise in urban areas is crucial to evaluate and to reduce overall levels. However, this situation is difficult to model because

cities are both large and complex. Some methods (e.g. finite element or finite-difference methods, which this document calls reference methods) can model complicated shapes but are too expensive to use for several cities blocks while other methods (e.g. engineering or geometrical methods, which this document calls efficient methods) can model large scenes but cannot accurately model complicated geometries. This dissertation develops an approach to model outdoor sound propagation that is less expensive than the current reference methods and is more accurate than the current efficient methods.

To be more specific, this dissertation seeks to improve the accuracy of the engineering methods (EMs) at a reasonable cost. EMs are the standard methods that governments recommend to model urban outdoor noise propagation (e.g. CNOSSOS-EU⁵) because in many cases they efficiently estimate the long term noise levels. However, EMs assume that all of the objects in the scene have a very simple (box-like) shape, which can produce large errors for more complex shapes⁶. The method developed here augments the capabilities of EMs to include long objects with a constant, arbitrary cross-section (e.g. a noise barrier).

The literature provides many cases where this kind of extension would be useful. The principal case is a complex noise barrier^{6,7}, which could be rectangular (i.e. a thick straight barrier), T-shaped, Y-shaped, cylindrical, angled, curved, or absorbent. The barrier shape and surfaces could even come from an optimization algorithm^{8,9}. EMs can only model straight, thin barriers (hereafter called I-barriers), so anything more complicated would be an extension. In addition, even I-barriers should not be modeled next to a train or tram using EMs because this case would require too many reflections between the barrier and the train to be cost effective, so it is modeled with the boundary element method¹⁰.

Modeling roof shapes is also beyond the scope of EMs because they assume that all buildings are box-shaped. However, the shape of the roof can have a large impact on the sound attenuation of the building¹¹. In addition to shape, greening measures where surfaces are covered with plants can also lead to improved noise reduction¹².

These cases provide examples where the extra attenuation of the complex object compared to a reference object that is currently used in EMs is large compared to the uncertainty of EMs, which is 1-3 dB for a wide variety of cases^{13,14}. These applications are the kinds of cases where another approach has the potential to greatly improve the accuracy of EMs for a reasonable cost.

Therefore, the primary objective of this dissertation is to enable EMs to model more complicated geometries while minimizing the increased computational costs. This goal has been suggested at least twice before. First, Nord2000^{15,16}, a Nordic EM, discusses

“Special Screens” in §4.8 but does not give enough details to implement the method and does not verify the outlined approach. Second, the Hosanna project, in Work Package 6.2¹⁷ in §3.5 called “Coupling with Numerical Models”, describes their method in more detail but does not justify most of the decisions and does not discuss multiple diffraction points. The document validates the method for a lone barrier but does not validate it when the barrier is in a complex surrounding. In contrast, this dissertation both develops and validates a hybrid method to extend EMs to more complex geometries.

§1.3 Literature review

To decide what approach is most appropriate, the first step is to review the current literature on outdoor noise propagation. Researchers have developed many methods to help governments estimate current noise levels and to evaluate potential noise mitigation solutions. The major categories of methods are geometrical, engineering, statistical, frequency-domain, time-domain, and hybrid approaches. This section introduces a few methods for each approach and discusses the advantages and disadvantages of each model. The proposed approach is in §1.5.

§1.3.1 Geometrical approaches

For efficiency, geometrical approaches assume that the objects are large compared to a wavelength. This assumption makes the computation time independent of frequency and instead depends on the number of reflections, diffractions, and polygons to define the geometry. When these variables are small in number, geometrical approaches can perform very well, but when they are large the computation time can be prohibitive. However, the main objection is that they do not allow for the geometric complexity that this application requires. Geometrical approaches may or may not include the phase information of each path. The image source, ray tracing, beam tracing, and geometrical theory of diffraction methods are all common geometrical approaches.

The image source method¹⁸ models specular reflections by inserting image sources and has been used to model outdoor street canyons¹⁹⁻²¹. Unfortunately, it does not usually model diffraction or atmospheric effects, and for complicated scenes the cost is unaffordable due to the number of image sources and the visibility/validity checks²².

The ray tracing method^{23,24} sends out many rays from the source and detects if any of them hit a small sphere centered at the receiver. Ray tracing has been used for urban scenes²⁵⁻²⁷, can model atmospheric effects^{28,29}, and can model diffraction over simple barriers including a Y-barrier³⁰. However, barriers could have an arbitrarily shaped cross-

section where the computational cost would be unreasonable because of the large number of diffractions required to accurately model such a complex barrier.

Similar to ray tracing, the beam tracing method³¹ sends out many beams from the source and detects if any of them hit a point receiver. Beams have a finite cross-sectional area whereas rays have an infinitesimal one. Beam tracing has been used to model outdoor sound propagation^{32,33}, but cannot model complex shapes just like ray tracing.

The geometrical theory of diffraction³⁴⁻³⁷ extends pure geometrical methods to include diffraction. If a ray from the source hits an edge or corner, then a source is added at that location. The geometrical theory of diffraction can model moderately complex barriers^{38,39} (e.g. T, Y, and inclined barriers), but the cost increases rapidly with the number of diffraction edges, so much more complicated barriers are cost prohibitive.

§1.3.2 Engineering approaches

EMs are standardized geometrical approaches, but they are given their own section here because of their importance for modeling outdoor noise propagation and because they make some additional simplifying assumptions. EMs find all of the propagation paths between the source and the receiver in a horizontal plane. The paths with reflections are then unfolded using image sources or receivers. The sound power is known for each source and attenuations are calculated for each path. The attenuations correspond to physical phenomena (e.g. geometrical divergence, atmospheric absorption, etc.), which are assumed to be mostly independent of each other. Then, the contribution of each path is summed incoherently to find the total sound pressure level. There are several EMs⁴⁰, including the international standard ISO 9613-2⁴¹, the French standard NMPB-Routes-2008^{42,43}, the Nordic standard Nord2000^{15,16}, and the European Union standard CNOSSOS-EU⁵. In addition, CSTB helped develop the Harmonoise EM^{44,45}.

The main advantage of EMs is their efficiency. The calculation times of EMs is independent of frequency and domain size and instead depend on the number of reflections and diffractions, which makes EMs feasible even when the domain size is much larger than a wavelength. The primary disadvantage of EMs is the limited geometric complexity. For example, EMs can only model I-barriers and cannot model more complicated shapes.

EMs are fundamentally different from reference methods because EMs incoherently sum the path contributions while reference methods coherently sum them. The difference is that EMs exclude the phase information of the complex acoustic pressure because they sum the squared pressures and the reference methods include the phase information because they sum the pressures. EMs coherently sum the rays within a propagation path (i.e. all the rays in the same vertical plane including the direct ray and ground

reflections), but the propagation paths are summed incoherently (i.e. rays that are not in the same vertical plane including the direct ray and a ray reflected from a vertical surface).

§1.3.3 Statistical approaches

Statistical approaches are methods that use statistics to make the computations more efficient or to gain physical insight. This section discusses the radiosity, diffusion, and statistical learning methods. The primary objection to the statistical approaches is that they cannot model complicated geometries. Moreover, most do not model atmospheric effects like refraction and turbulence.

The radiosity method²¹ assumes that boundaries reflect sound diffusively instead of specularly, which is in contrast to the geometrical approaches. The boundaries are broken up into patches where the accuracy and cost increase with the number of patches. The method calculates the energy exchange between all of the patches using ideas from heat transfer. The radiosity method has modeled urban outdoor sound propagation a few times^{21,46-49}, but unfortunately complicated geometries require many patches. In addition, specular reflection and diffraction must be treated separately⁵⁰.

The diffusion method⁵¹ uses the diffusion equation by imagining many sound particles emanating from the source. The particles can be absorbed by or reflected from the boundaries. Instead of tracing each particle, the diffusion method considers the probability density function of all of the particles. The diffusion method can model urban sound propagation analytically in simple cases⁵²⁻⁵⁴ and numerically using finite-differences in complex cases^{55,56}. However, in complicated cases the cost increases and estimating the diffusion coefficient is difficult.

Statistical learning methods^{57,58} are a very powerful and diverse set of tools to understand large and complex data sets. These methods are decades old, but they have recently seen renewed interest because computers have become more powerful and because big data has become increasingly available and important. Unfortunately, not many articles apply these methods to outdoor sound propagation. One exception⁵⁹ compares the parabolic equation method, three EMs, and four statistical learning methods for sound propagation over flat ground. In this case, they found that the statistical learning methods are more accurate than EMs; however, transitioning from this simple case to the much more complicated case in a city is unclear. The amount of data required might be prohibitive.

§1.3.4 Frequency-domain approaches

Frequency-domain methods assume a time-harmonic source, which simplifies the time derivatives of the differential equations. Frequency-domain methods perform single frequency computations and work well with boundary conditions that are defined in the frequency-domain. However, as the frequency and domain size increase the cost also increases, which quickly becomes prohibitive and is the main limitation of these methods. This section discusses the finite element, boundary element, parabolic equation, and equivalent sources methods. For completeness, many frequency-domain approaches also have time-domain implementations and vice versa, but the methods are categorized here based on the most common approach taken in the literature on outdoor noise propagation.

The finite element method⁶⁰⁻⁶⁴ solves differential equations (e.g. the Helmholtz equation in acoustics) by dividing the problem's domain into many pieces called elements. Within each element, the field variable (e.g. pressure in acoustics) is usually assumed to be constant or vary linearly (or quadratically). Minimizing the error between this piece-wise approximation and the true solution yields a system of linear equations, which is solved for the pressure throughout the domain. The finite element method uses perfectly matched layers⁶⁵ or infinite elements⁶⁶ to model outdoor (i.e. infinite domain) problems. The finite element method is a very powerful method because it is very flexible and can model arbitrary geometries, but as the number of elements increases, the cost also increases. The finite element method is much too expensive for this application.

The boundary element method (BEM)⁶⁷⁻⁶⁹ is an approximation of the Kirchhoff-Helmholtz integral equation, which relates the pressure in a domain to the pressure and its normal derivative on the domain's boundary. The boundary is discretized and the boundary conditions are specified everywhere on the boundary. Then, the approximate integral and the boundary conditions together produce a system of linear equations that yields the unknown boundary values. Lastly, the pressure anywhere in the domain is calculated from the boundary values. BEM is often used to calculate the attenuation of noise barriers^{6,70,71} and outdoor sound propagation more generally⁷². The fast multipole BEM⁷³⁻⁷⁷ provides a substantial improvement to BEM in terms of computation time because it 1) makes the matrix sparse by putting the boundary contributions into a hierarchy, 2) uses an iterative solver, and 3) does not store the entire matrix equation. However, like the finite element method, BEM is too expensive for this case.

The parabolic equation method⁷⁸⁻⁸⁰ is based on the parabolic approximation of the Helmholtz equation that models one-way propagation from a monopole within a limited vertical angle. This method has at least two important implementations that approximate the spatial derivative in different ways. The Crank-Nicholson parabolic equation method approximates it by a truncated Taylor series and the Green's function parabolic equation method approximates it with a Fourier transform. The parabolic equation method has

been used for many long range outdoor propagation problems where refraction^{81–83} and turbulence^{84,85} are important. However, this method cannot model complicated shapes and usually does not model reflections from vertical objects⁸⁶.

The equivalent source method^{87,88} divides the problem's domain into subdomains that have known Green's functions. The Green's function gives the pressure everywhere inside a subdomain due to a source inside a subdomain. The interfaces between the subdomains are meshed, and each element is represented by an acoustic source with unknown source strength. Enforcing continuity on pressure and velocity produces a system of equations, which give the source strengths. Inserting these source strengths into a sum of corresponding Green's functions yields the pressure in the subdomain. The equivalent source method has been used for predicting noise propagation around and over buildings^{89–91}. However, finding the appropriate Green's function may be difficult for complicated geometries and surfaces, and the matrices can rapidly become ill-conditioned.

§1.3.5 Time-domain approaches

Time-domain methods take the pressure information from the current and previous time steps and the boundary conditions to predict future pressure fields. Time-domain methods are well-suited to model transient phenomena (e.g. pulses, turbulence, and wind/temperature changes) and moving sources. This section discusses the finite-difference time-domain method and the transmission line matrix method.

The finite-difference time-domain method⁹² marches forward in time using the pressure information at the current and previous time steps to determine future pressure values. Approximating the derivatives of the wave equation by truncated Taylor series yields the future pressure values. The finite-difference time-domain method discretizes both time and space. This powerful method has been used to model urban acoustics with real geometries and atmospheric effects^{11,12,93–96}, but because of the computational cost the simulations have been restricted to 2D models or to below 500 Hz.

The transmission line matrix method⁹⁷ models a problem's domain as a Cartesian grid of pipes called transmission lines. Pressure pulses travel through the transmission lines, and conserving the pressure and energy at each intersection determines the behavior of the pressure pulse. Since these interactions are determined *a priori*, this method does not require matrix inversions. The transmission line matrix method has been used to model outdoor sound propagation^{98–101} but is too expensive for this application because the domain size is so large compared to a wavelength.

§1.3.6 Hybrid approaches

Hybrid methods combine more than one method to achieve a goal (e.g. to reduce cost while maintaining similar accuracy). Many hybrid methods try to optimize cost and accuracy by applying a reference method to the complex geometries and an efficient method to the simple geometries, which is a spatial decomposition. One approach uses the finite-difference time-domain method as the reference method and the parabolic equation method as the efficient method^{102–104}. Others use BEM as the reference method with the parabolic equation method^{105,106} or ray tracing^{107–110} as the efficient method. Another approach uses both spatial and frequency decomposition using the equivalent source method for locations close to objects and at low frequencies and ray tracing otherwise¹¹¹. However, these methods are still very expensive because the vast majority of the computation time is still dedicated to the reference method.

In order to reduce the computation time, others have suggested that the attenuation of a complex geometry should be computed beforehand and often in a smaller number of dimensions. There are two main approaches in the literature. The simplest method^{15,17} tabulates the attenuation based on a large number of reference calculations, and interpolating the data yields a specific attenuation. The other method^{7,112} goes one step further and fits a curve to the tabulated results so that the function and its coefficients are stored instead of the raw data, and evaluating the function gives the desired attenuation.

The main limitation of these methods is that they have only been used in the simple case of a barrier with hard, flat ground. None of the methods explain how to use these corrections in a complex scene with other objects or ground types. Specifically, the methods need to specify what to do when a path includes reflections and diffractions. The method developed in this dissertation seeks to determine how to incorporate simple corrections into complicated urban scenes.

§1.4 Model requirements

After reviewing the main approaches for modeling outdoor noise propagation, the next step is to establish criteria to evaluate if a method is appropriate for the proposed application, namely making road noise maps in large, complicated urban areas. Since EMs are the current standard, some of these requirements use EMs as a reference point. The requirements include physical phenomena, frequency range, geometric extent, geometric complexity, and computational restrictions. Table 1.1 summarizes the requirements.

First, the approach must reasonably approximate all of the most important physical phenomena for outdoor sound propagation including geometrical divergence, atmospheric absorption, ground effects, reflections, diffraction, refraction, and turbulence.

Second, the approach must be able to calculate noise levels over a broad range of frequencies. Specifically, the minimum requirement is the 125 Hz to 4 kHz octave bands, which corresponds to the range in CNOSSOS-EU. This range contains the majority of A-weighted sound power from cars, trucks, and trains and is where human hearing is most sensitive. Aircraft noise would require a lower minimum frequency. Preferably, the method could also calculate third-octaves and would have an even larger range.

Third, the approach must be able to handle 3D problems with propagation distances of at least 100 m and potentially up to 1 km. Urban problems commonly have propagation distances over 100 m but over 1 km is rare except for aircraft noise propagation. Methods capable of longer propagation distances are preferable to ones with shorter maximum distances.

Forth, the approach must be able to model objects with an arbitrary 2D cross-section. This requirement balances flexibility and cost. For this application, objects are usually considered long in the dimension that is perpendicular to the propagation, so allowing object to change shape along that dimension is not nearly as important as the vertical cross-section of the object. Currently, the cross-section has to be shaped like a wedge or a rectangle. An arbitrary cross-section gives much greater flexibility.

Fifth, the computation time must be less than 24 hours on a standard desktop, which currently has a quad-core 3.4 GHz processor with 16 GB RAM and 512 GB SSD¹¹³ if it is new and slightly lower performance if it is older. Computations that require more than one day to complete are impractical for many users who may have clients waiting for their conclusions. Preferably, the computation time should be as short as possible.

Table 1.1: Method requirements

This table summarizes the requirements and their justifications that are presented in this section.

Name	Requirement	Justification
Physical phenomena	All of them from EMs	The new approach should not lose capabilities
Frequency range	125 Hz to 4 kHz octave bands	This range is from CNOSSOS-EU
Geometric extent	Propagation distances of at least 100 m and potentially up to 1 km	Most applications are between 10 m and 1 km
Geometric complexity	Objects with an arbitrary cross-section	This requirement is the main improvement
Computational restrictions	Computation time < 24 hours on a standard desktop	Longer computation times are unreasonable; many users cannot access more powerful computers

§1.5 Choosing a method

After discussing the main methods to model urban outdoor sound propagation and establishing a way to evaluate them, the next step is to determine the best method for this application. Most of the methods do not meet the requirements (Table 1.2); all of the non-hybrid approaches fall short because they cannot model complex geometries or are too computationally expensive. For example, combining the frequency range and geometric extent requirements means that the domain is much larger than a wavelength. Also including the computational restrictions requirement eliminates all methods that model the entire domain and increase in cost with frequency. Specifically, these requirements exclude the frequency-domain and time-domain approaches. In addition, the geometric complexity requirement excludes the geometrical, engineering, and statistical methods. Thus, the hybrid approaches are the only ones that remain.

However, currently there are no developed, validated, and published methods that fit these requirements, which is the motivation for this dissertation. Most approaches either do not model complex geometries or are computationally prohibitive for this application. The HOSANNA project¹⁷ documents the most promising approach, but it is not fully developed or verified. The closely related ray tracing/BEM hybrid methods^{108,110} first model the source near complex geometries to find the equivalent sources on a fictitious interface and then use ray tracing to determine the level at the receiver. However, for 3D problems and high frequencies the number of sources would be very large. Therefore, this dissertation seeks to develop and validate a hybrid method to model urban outdoor sound propagation.

Table 1.2: Comparing potential approaches

This table considers if any of the main categories of approaches satisfy the requirements from §1.4.

Approach	Physical phenomena	Geometric complexity	Geometric extent	Geometric complexity	Computational restrictions
Geometrical	Yes	Yes	Yes	No	Yes
Engineering	Yes	Yes	Yes	No	Yes
Statistical	No	Yes	Yes	No	Yes
Frequency-domain	Yes	Yes	Yes	Yes	No
Time-domain	Yes	Yes	Yes	Yes	No
Hybrid	Yes	Yes	Yes	Yes	Yes

The first major decision is to determine which methods should be combined. Generally, there is an efficient method that can model a simplified version of the entire scene, and a reference method that can model complicated objects but is too expensive to apply to the entire domain. The engineering, geometrical, and statistical approaches provide the options for the efficient method, and the frequency-domain and time-domain approaches are possibilities for the reference method. Following the example in the HOSANNA project, the hybrid method developed in this dissertation combines EMs and 2.5D BEM. The rest of this section justifies these choices.

§1.5.1 Justifying using EMs

The hybrid method uses EMs because they are

1. Able to model important physical phenomena;
2. Very efficient;
3. Easily extendable;
4. Widely utilized.

First, EMs estimate the effects of many critical physical phenomena⁵ including geometrical divergence, atmospheric absorption, ground effects, reflections, diffraction, refraction, and turbulence. While many of these effects are only approximated, EMs have been validated with both reference methods and measurements^{13,17,59,114}.

Second, EMs estimate noise levels very efficiently. In particular, EMs make a high frequency approximation so that the calculation time is independent of frequency. While the geometrical methods also make this assumption, EMs are still theoretically less expensive because they make additional simplifications (e.g. combining paths within a vertical slice and summing the propagation paths incoherently) that should make EMs less expensive.

Third, EMs are easily extendable. Their assumption that sources of attenuation can be calculated independently makes including another attenuation term trivial, although calculating its value may be very complicated. Augmenting a statistical method would be much more complicated.

Fourth, EMs are widely used in governmental and legal contexts where their results have monetary implications and can affect the health and well-being of millions of people. Governments use them to estimate and mitigate the noise exposure of their constituents. For example, the directive 2002/49/EC¹¹⁵ requires European Union member countries to create noise maps using the EM called the Common Noise Assessment Methods for Europe⁵ (CNOSSOS-EU). These noise maps enable European citizens to know their noise exposure and help appropriate authorities create action plans to mitigate the negative effects of noise. Moreover, urban planners use them to predict how changing a city's infrastructure would change the noise exposure of the residents; consultants use them to evaluate if someone can sue their neighbor for breaking a noise ordinance; and researchers use them to estimate noise levels.

§1.5.2 Justifying using 2.5D BEM

The hybrid method uses the 2.5D BEM because it models:

- Arbitrary cross-sectional geometries;
- Infinite domains and infinite, hard ground efficiently;
- Point sources.

The most important requirement of the reference method is that it can model geometries with an arbitrary cross-section because this aspect is the improvement that the reference method provides for EMs. This requirement reduces all of the options discussed in §1.3 to three main options: BEM, the finite element method, and the time-domain finite-difference method. This requirement does not include arbitrary 3D geometries because the cost would be even greater and it would not separate the paths that diffract over and around the complex object. Thus, since they are more efficient, this requirement also prefers 2D or 2.5D approaches, which assume a constant cross-section in the third dimension, to truly 3D approaches.

The next most important requirement is that the reference method models infinite domains and infinite hard ground efficiently because many outdoor propagation problems include these features. While all three methods can model infinite domains, only BEM does so without substantial extra expense. BEM simply changes the Green's function and does not require a larger mesh.

The last requirement is that the reference method models point sources because EMs use point sources. This requirement excludes 2D BEM because it only models coherent line

sources and does not model point sources, which does affect the predicted attenuation¹¹⁶. Thus, 2.5D BEM fits the requirements the best.

Specifically, this dissertation uses MICADO¹¹⁷ as the reference method. MICADO is a variational 2.5D BEM approach. Variational BEM reformulates the integral equation to reduce the influence of the irregular frequencies that correspond to internal resonances such that fewer points are required per wavelength compared to direct approaches¹¹⁷. 2.5D methods assume a constant cross-sectional area in one dimension so that a 3D problem can be transformed into a set of 2D problems that can be much less expensive to compute¹¹⁸. Since the 2.5D approach reduces the problem dimension by one and BEM reduces the dimension by one (i.e. only the boundary of the domain must be meshed), MICADO reformulates the 3D problem as a set of 1D problems. All of these improvements in efficiency combine to produce a reference method that can be used up to at least 5 kHz.

§1.6 Thesis organization

Again, the main purpose of this work is to find a method that can model urban noise propagation in large, complicated scenes for a reasonable cost. To that end, the dissertation is divided into five chapters: 1) introduction, 2) explanation, 3) validation, 4) optimization, and 5) summary of the hybrid method.

Chapter 1 (i.e. the current chapter) introduces the problem of urban outdoor noise propagation and proposes an approach to address the problem. First, the chapter explains why modeling long-term, urban noise levels is important but difficult, which provides the motivation for the dissertation. Then, the chapter presents the state of the art for outdoor noise propagation to hopefully find a method that is better than standard EMs and defines what is required to be better (i.e. more accuracy for a reasonably extra cost). The chapter concludes that none of the current methods meet the criteria and proceeds to suggest a hybrid method that combines EMs and 2.5D BEM. The chapter closes by outlining the entire dissertation.

Chapter 2 explains and justifies the hybrid method. It starts with an outline of the entire method and then goes into detail on each of the parts. The chapter describes the relevant aspects of the two main parts of the hybrid method, namely EMs and the precomputations using 2.5D BEM. After understanding these two big parts, then the chapter details how to connect them by first converting the source and receiver positions from global coordinates (where the origin could be anywhere) to local coordinates (where the origin is at the base of the complex object). Finally, the chapter explains how to interpolate (or

extrapolate) the precomputed BEM data at the local coordinates of the source and receiver to obtain the influence of the complex object for the given path.

Chapter 3 validates the hybrid method for a wide variety of cases. It starts with the simplest case and works its way up to a real situation. The first case is a T-barrier with hard flat ground, which evaluates the error that is introduced by interpolating a finite set of points that are far apart compared to a wavelength. The next case is the same T-barrier on soft ground, which demonstrates how invariant the T-top correction is to the ground type. Afterward, the chapter considers a T-barrier with buildings. This case evaluates the conversion from global to local coordinates for the source and receiver positions. Finally, the chapter considers a real case to demonstrate that the hybrid method is capable of handling larger cases.

Chapter 4 proposes a method to optimize the locations of the data points that are calculated with 2.5D BEM. First, the chapter explains why it is more appropriate to minimize the uncertainty of a predicted value rather than directly minimizing the error of a predicted value, which is far more common. Then, it explains why a Gaussian process is an appropriate approach to model the uncertainty. Next, it specifies how to implement a Gaussian process and determine its hyperparameters using the Newton method. Then, the chapter states how to minimize the uncertainty and choose the number of points for each dimension. The chapter closes with a comparison of the interpolation error using three different data distributions.

Chapter 5 concludes the dissertation by looking back at what has been accomplished and by looking forward to the work that is still to be done. It synthesizes the principle results while highlighting the novelty and importance of the work and identifies ways to further improve and extend this work.

§1.7 Summary

Millions of people are being exposed to high noise levels that are associated with negative health outcomes like annoyance and sleep disturbance. In response, governments have taken action to estimate and mitigate the noise levels using noise EMs. However, these methods cannot model complicated geometries or surfaces, which may offer significant noise reduction.

Other outdoor noise propagation models do not meet the requirements of an improved EM. The geometrical, engineering, and statistical methods cannot model the desired geometric complexity, and frequency-domain and time-domain methods are prohibitively expensive. Many hybrid methods are too expensive as well.

However, combining EMs with 2.5D BEM to create a hybrid method satisfies the criteria because this approach reduces the computation time compared to reference methods and improves the accuracy compared to efficient methods. This dissertation develops (Chapter 2), validates (Chapter 3), and optimizes (Chapter 4) this hybrid method.

Chapter 2 Hybrid Method

§2.1 Objectives

The primary objective of this chapter is to articulate and justify the hybrid method. To that end,

§2.2 outlines the entire hybrid method;

§2.3 describes EMs;

§2.4 specifies how to create the data table of corrections;

§2.5 explicates how to find the local source and receiver positions;

§2.6 elucidates how to interpolate the data to obtain the extra attenuation;

§2.7 details how to extrapolate the data for points that are outside the data.

In detail, §2.2 overviews the entire hybrid method. The primary purpose of the outline is to become acquainted with all of the major pieces of the hybrid method and to see how they fit together. Justification of each part occurs in the sections that follow. Since the core of the hybrid method is EM, the first detailed section (§2.3) describes the most relevant parts of EMs. Then, §2.4 describes how to make the table of correction factors using 2.5D BEM. The correction factors quantify the influence of the complex object. After describing EMs and BEM separately, the remaining sections illustrate how they interact. Specifically, §2.5 details how to convert the source and receiver positions that EM uses (i.e. global coordinates) to the source and receiver positions that are necessary to use with the BEM data (i.e. local coordinates). Finally, §2.6 states how to interpolate the data table at the local source and receiver positions to yield the extra attenuation, and §2.7 describes how to extrapolate the data if the local coordinates lie outside of the data table.

§2.2 Method outline

This section gives a big picture view of the hybrid method, but it does not give all of the details and certainly does not justify the method. The rest of this chapter is dedicated to clearly explaining the method and justifying the details of its implementation.

Conceptually, the hybrid method trusts that EMs estimate long-term noise levels with sufficient accuracy unless the scene includes something beyond their capabilities. For example, EMs only model very simple geometries. When a scene has a complex object, EMs need some way to quantify the attenuation of the complex object for each propagation path on the surrounding sound levels.

This situation is where 2.5D BEM can step in. It can model geometries that are far more complex than EMs can, but it is also much more expensive. To minimize this cost, BEM models the complex object and separately a reference object in very simple surroundings. The difference of these calculations is the influence of the complex object compared the reference object that can be included in EMs. These BEM computations yield the extra attenuation of the complex object for each propagation path, which can then be included into EMs.

However, even in very simple surroundings, BEM is too expensive to run for every propagation path. Thus, a large but finite set of data points is precomputed and then interpolated whenever EMs need it. Since the data points are far apart compared to a wavelength, this approach does introduce some interpolation error, but optimizing the locations of the data points reduces the error. Since the origin of the coordinate system is likely different between the real problem and the simplified problem that BEM uses, interpolating the BEM data requires that the source and receiver positions are converted to the corresponding local coordinate system. In addition, when the source or receiver coordinates fall outside of the precomputed data, the extra attenuation must be extrapolated instead of interpolated.

Concretely, the hybrid method has two major steps:

1. Calculate and tabulate the attenuation of each complex object compared to a reference object at a large number of source/receiver positions and frequencies;
2. Run EM with an extra attenuation term for the complex objects.

The first step also has two parts: the BEM computations and the post processing of the data. In the first part, 2.5D BEM produces the noise levels for a complex case and a simplified case. Both cases usually use a homogeneous atmosphere; flat, hard ground; and the same sources, receivers, and frequencies. However, the complex case has a special object that EMs cannot model accurately, and the reference case has an object that is similar to the special object but can be modeled with EMs. The source/receiver

positions and frequency are on a non-uniform grid over the entire range of interest. In the post processing part, numerically integrating the BEM results over frequency yields third-octave bands instead of individual frequency calculations. Finally, subtracting the reference case from the complex case produces a 6D table of corrections (over the source/receiver positions and frequencies) that is stored in a .txt file so that EM can use it later. This step only has to be performed once per complex object, and the results are stored in a database. This procedure saves significant computation time when the corrections are used multiple times.

In the second step, EM runs normally except it also calculates an extra attenuation term. EMs first create the scene based on the user inputs, and each complex object is added to the scene as its reference object with an associated table of corrections, which is loaded into memory. Next, a path finding algorithm searches for the important paths between the source and receiver, which is unchanged by the hybrid method. For each path, the hybrid method calculates all the attenuations including the extra attenuation term to find the path contribution.

For paths that do not diffract over a complex object or that only diffract over complex objects that do not have the largest path length difference of all of the objects in the path, the extra attenuation is zero. For paths that diffract over a complex object that has the largest path length difference, the extra attenuation is obtained by:

1. Calculating the local source and receiver positions based on the geometry of the path
2. Linearly interpolating the table of corrections at the local source and receiver position

Finally, summing the path contributions yields the level at the receiver. Figure 2.1 summarizes these steps in the hybrid method and Figure 2.2 describes how to calculate the extra attenuation.




Figure 2.1: Hybrid method outline

The hybrid method has the following steps: 1) calculate the levels for the complex and reference cases for a large number of sources, receivers, and frequencies; 2) integrate over frequency to produce frequency bands and subtract the cases; 3) setup the scene including complex objects; 4) find the relevant propagation paths between the source and the receiver; 5) calculate the attenuations including an extra attenuation term for the complex objects; and 6) sum all of the path contributions to find the level and the receiver. The first two steps are unnecessary when the corrections are already accessible in a database.




Figure 2.2: Calculating the extra attenuation

To find the extra attenuation, 1) determine if a path has a complex object that has the largest path length difference of any object in the path, 2) compute the local source and receiver positions for the large complex objects, 3) interpolate the table of corrections that comes from BEM at the local coordinates to produce the extra attenuation.

§2.3 Engineering methods

The core of the hybrid method is an EM, so this chapter first explains the most relevant parts of EMs before detailing how to improve them. For simplicity, this section only focuses on CNOSSOS-EU⁵ because it is the most recent standard, but the hybrid method is flexible enough to work with any EM.

CNOSSOS-EU⁵ accomplishes five main tasks:

1. States its objectives and requirements;
2. Describes the frequency spectrum and directivity of common urban noise sources;
3. Formulates how to calculate sound propagation;
4. Specifies how to position receivers on façades and to assign people to those receivers;
5. Provides guidelines for its use.

This dissertation focuses almost entirely on improving the sound propagation calculation but sometimes looks to the other chapters for guidance. For example, CNOSSOS-EU states:

CNOSSOS-EU should be designed to produce plausible noise maps showing plausible results... A parameter is considered essential if the range of values the parameter can take yields variations in L_{den} or L_{night} of more than ± 2.0 dB(A) 95% C.I. (all other parameters remaining unchanged).⁵

These requirements guide the decisions of the hybrid method and are indeed part of the impetus for developing the hybrid method.

Modeling outdoor sound propagation in CNOSSOS-EU⁵ has the following steps:

1. Input the scene (i.e. the source/receiver locations and the geometry);
2. Find 2D propagation paths between the source and receiver in a horizontal plane;
3. Calculate the attenuations for each propagation path;
4. Sum the path contributions to find the total sound level at the receiver.

The rest of this section details how to implement each of these steps.

§2.3.1 Scene setup

The first step is to input all of the data that the method needs. In CNOSSOS-EU⁵, the user specifies

- The ground's shape and type;
- The (box-shaped) objects' locations, dimensions, and surfaces types;
- The sources' locations and directional sound power for each frequency band;
- The receivers' locations;
- The probability of favorable wind speed and temperature gradients.

For example, consider a fictional case with a road, a barrier next to the road, and a building further away from the road (Figure 2.3). The scene setup describes the geometry including surface types, the source, the receiver, and the atmosphere.

CNOSSOS-EU also specifies where the sources and receivers should and should not be put, which is useful for designing the hybrid method. For example, in CNOSSOS-EU the source and receiver should be within 800 m of each other; although, the total propagation distance with reflections may be up to 2 km. In addition, sources are 5 cm high for roads and should be defined on a curve, which is then approximated by incoherent point sources. Receivers are typically spaced up to 5 m apart in the horizontal plane but the height should be known to the nearest 10 cm. The receivers should be at least 2 m from the ground or any building and are 4 m high for noise maps. Receivers must be at least 2 m from a building façade because EMs incoherently sum the direct path and the path reflected from the building, which causes a 3 dB error for a receiver on the façade. This difference must be added to points close to vertical surfaces. All of these CNOSSOS-EU recommendations are useful when deciding where to put the sources and receivers.




Figure 2.3: Example scene

This simple, fictitious scene illustrates the kinds of input data that EMs require. Specifically, they need the source and receiver positions, the ground type, and the locations and geometries of the barrier and the building.

§2.3.2 Finding propagation paths

After the scene has been established, the next step is to find the propagation paths between the source and the receiver. For simplicity and efficiency, EMs only search for paths in a 2D horizontal plane. From this view (Figure 2.4), paths appear to go through objects when they diffract over objects. EMs are 2.5D approaches because they first look for paths in a 2D horizontal plane and then calculate attenuations based on the geometry of the 2D vertical plane. This approach is not truly 3D because it would not find paths through a tunnel or under a bridge.

There are three main types of propagation paths: direct, reflected, and diffracted (Figure 2.4). Direct paths go straight between the source and the receiver as seen from above but may diffract over objects like a barrier. Reflected paths are reflected from vertical surfaces like a building. Diffracted paths contain lateral diffractions where the sound is diffracted around a vertical edge like the end of a barrier. Paths can also contain both reflections and lateral diffractions. EMs incoherently sum the contributions of multiple propagation paths.

Within each propagation path, there is a coherent set of ray paths to account for the ground affect. A propagation path and a ray path are distinct because a propagation path is the coherent set of ray paths in a vertical plane through the source and the receiver. Thus, the direct ray and the rays that are reflected from the ground are summed coherently and do not need to be found separately. For example, the red direct propagation path in Figure 2.4 actually contains four paths when viewed in the vertical plane. All four paths diffract over the barrier, but the paths have different numbers of ground reflections: 1) zero, 2) one on the source side of the barrier, 3) one on the receiver

side, and 4) two (i.e. one on each side). The fourth ray with two ground reflections is typically ignored in EMs. Thus, EMs coherently sum the first three ray paths but incoherently sum the direct, reflected, and diffracted propagation paths.

The algorithm stops searching for paths when it has found all of the propagation paths up to a given order of reflection and lateral diffraction or has traveled more than a preset distance (e.g. 2 km). Usually, the direct path is the most important followed by the reflected and then the diffracted paths. Thus, only a small number of reflections and lateral diffractions are required in many cases. Commonly, only two or three reflections are included and lateral diffraction can be ignored for extended sources (e.g. an incoherent line source).




Figure 2.4: Types of propagation paths

The three main types of propagation paths are direct, reflected, and diffracted. A propagation path may also be a combination of these basic path types.

§2.3.3 Calculating the attenuations

After finding the propagation paths, EMs can calculate the path contributions. EMs differ on how to calculate some of the attenuations, so for simplicity the method detailed here is taken from CNOSSOS-EU, the current EU standard. For more details or clarification, see the CNOSSOS-EU standard⁵.

CNOSSOS-EU calculates the path contributions for multiple atmospheric conditions and frequencies, but to explain only the essentials and to not recreate the standard here, this section assumes homogeneous conditions and only one frequency. Thus, the path contribution is

$$L = L_{W,dir} - A \quad (2.1)$$

where $L_{W,dir}$ is the directional sound power for the given frequency band, which is given by the user. In this equation and the following equations, L , $L_{W,dir}$, and all of the attenuation terms are in dB. A is the attenuation along the path:

$$A = A_{div} + A_{atm} + A_{boundary} \quad (2.2)$$

where

$A_{div} = 20 \log_{10}(d) + 11$ is the attenuation from geometrical divergence;

$A_{atm} = \alpha_{atm} d/1000$ is the attenuation from atmospheric absorption;

$A_{boundary} = A_{ground} + A_{dif}$ is the attenuation from the ground and diffracting objects.

In the above equations, d is the direct distance between the source and the receiver and α_{atm} is found in a table in ISO 9613-1¹¹⁹. Since this dissertation is concerned with including complicated diffracting objects, let there be a diffraction object. In this case, $A_{ground} = 0$ because the ground effects are included in A_{dif} . Then, the attenuation from diffraction is

$$A_{dif} = \Delta_{dif(S,R)} + \Delta_{ground(S,O)} + \Delta_{ground(O,R)} \quad (2.3)$$

where

$\Delta_{dif(S,R)}$ is the attenuation from diffraction between the source (S) and receiver (R);

$\Delta_{ground(S,O)}$ is the attenuation from the ground effect on the source side;

$\Delta_{ground(O,R)}$ is the attenuation from the ground effect on the receiver side.

Equation (2.3) is only an approximation and does not include a term for the ray path that reflects off of the ground on both sides of the barrier because it is considered negligible. For simplicity, this dissertation does not describe the ground effect terms in more detail, but the diffraction term is important. The attenuation from pure diffraction (i.e. without ground effects) must be in the range [0,25] dB and is

$$\Delta_{dif} = \begin{cases} 10C_h \log_{10}(3 + 40C'' \delta/\lambda) & \text{if } 40C'' \delta/\lambda \geq -2 \\ 0 & \text{otherwise} \end{cases} \quad (2.4)$$

where

$$C_h = \min(1, fh/250), \tag{2.5}$$

$$C'' = \begin{cases} \frac{3+r}{1+r} & \text{where } r = 75 \left(\frac{\lambda}{e}\right)^2 & e > 0.3 \text{ m} \\ 1 & e \leq 0.3 \text{ m} \end{cases}, \tag{2.6}$$

and

δ is the path length difference between the direct path and the diffracted path in meters;

f is the frequency at the center of the frequency band in hertz;

λ is the wavelength at the frequency f in meters;

h is the height of the highest diffraction edge in meters;

e is the distance between the diffraction point that is closest to the source and the diffraction point that is closest to the receiver in meters.

To better understand Equation (2.4), Figure 2.5 plots the pure diffraction (Δ_{dif}) versus the path length difference (δ), the primary variable for estimating the attenuation of objects, for multiple distances between the diffraction points (e) and for $C_h = 1$. The plot illustrates that the attenuation increases as the path length difference increases or as the distance between the diffraction points increases.




Figure 2.5: Pure diffraction

The pure diffraction term in CNOSSOS-EU increases with path length difference and distance between diffraction points.

In CNOSSOS-EU, C'' (from Equation (2.6)) is described as “a coefficient used to take into account multiple diffractions”⁵. When $e \rightarrow 0$ (i.e. there is only one thin barrier) $C'' \rightarrow 1$, and when $e \rightarrow \infty$ (i.e. the diffraction points are far apart) $C'' \rightarrow 3$. This correction is very simple for such a complex phenomenon, which suggests that paths with multiple diffractions are not very important to the overall level in most cases. Getting the attenuation right in these cases might not be very important because more direct paths are the dominate contribution. Thus, making large improvements for paths with multiple diffractions is not a priority for the hybrid method.

§2.3.4 Summing the path contributions

After the path contributions have been calculated, the last step is to incoherently sum all of the contributions. When calculating a noise metric like L_{den} , the contributions are weighted according to their frequency (e.g. A-weighting), time of day (e.g. day, evening, or night), and atmospheric condition (e.g. homogenous or downward refracting).

§2.3.5 The hybrid method's adjustment

With some background into EMs, the hybrid method can now be more clearly defined. For the hybrid method, Equation (2.2) becomes

$$A = A_{\text{div}} + A_{\text{atm}} + A_{\text{bounday}} - A_{\text{extra}} \quad (2.7)$$

where A_{extra} represents the attenuation of the complex objects or surfaces compared to a simplified object and comes from interpolating the BEM results, and computing these results is the subject of §2.4.

The minus sign is necessary because of how A_{extra} is defined and the minus in Equation (2.1). Conceptually, A_{extra} is the difference of two terms:

$$A_{\text{extra}} = L_{\text{complex}} - L_{\text{simplified}} \quad (2.8)$$

where L_{complex} is the level for the complex object and $L_{\text{simplified}}$ is the level for the simplified object. The definition of A_{extra} was chosen considering $L_{\text{simplified}}$ as a reference just like 20 μPa is the reference pressure in air (i.e. the reference level is subtracted from the real level). However, this definition is arbitrary and setting A_{extra} to its negation would remove the minus sign from Equation (2.8) and make A_{extra} act like an insertion loss. Physically, the different sign indicates that when A_{div} (or A_{atm} or A_{bounday}) increases the total level [i.e. L in Equation (2.1)] decreases whereas when A_{extra} increases the total level increases.

§2.4 Precomputations

Understanding EMs make stating the objective for the rest of this chapter clear: to efficiently and accurately calculate the extra attenuation, A_{extra} . This section specifies and justifies how to create the data table using 2.5D BEM. However, before calculating the data points, the number of variables and the coordinate system must be determined and substantiated (§2.4.1). Then, calculating the table of corrections has three parts:

1. Deciding the source/receiver positions and the frequencies;
2. Calculating the levels for the complex and reference cases;
3. Post-processing the data to obtain the table of corrections.

§2.4.1 Variables

The choice of variables is important because choosing an inappropriate coordinate system could make the problem more difficult and choosing an inappropriate number of variables could lead to unnecessary inaccuracy. This section justifies the choice of the coordinate system (§2.4.1.1) and of the number of variables (§2.4.1.2). Finally, §2.4.1.3 defines the variables used in the hybrid method.

§2.4.1.1 Coordinate system

Coordinate systems are chosen mostly out of convenience, which depends on the application. Spherical, cylindrical, and Cartesian coordinates are all common coordinate systems. On the one hand, researchers studying the attenuation of wedges using diffraction theory commonly use spherical coordinates because the attenuation of the wedges is a simple function of the angles to the source and to the receiver with the origin at the point where the ray path intersects the top of the wedge. However, in the current application the attenuation is not a simple function of the angles because the noise barriers are much more complicated than simple wedges. Thus, for more complex shapes defining the source and receiver positions in terms of angles does not have a large benefit.

On the other hand, Cartesian coordinates are convenient for a few reasons. First, noise maps are usually calculated at a constant height, so having data points in horizontal sheets is helpful. Second, since most cases use a flat, horizontal ground, extra care would need to be taken with cylindrical or spherical coordinates to ensure that points are not in the ground for certain angles to the source and the receiver because the origin is above the ground. Third, the steps that use the input variables are easier and less costly to perform with Cartesian coordinates. Specifically, the Cartesian coordinates simplifies the interpolation process and the 2.5D BEM computations. Thus, the hybrid method uses Cartesian coordinates because they are the most convenient.

§2.4.1.2 Number of parameters

The number of parameters is important because too few or too many would result in an inaccurate model because the total amount of data is limited by the computer memory. For each additional variable, the number of data points per variable decreases rapidly. For example, assuming the data points are on a grid, the total size of the data file in bytes (s) is approximately

$$s \approx pn^v \quad (2.9)$$

where p is the precision in bytes, n is the average number of points per dimension, and v is the number of dimensions or variables. The equation is only approximate because 1) it does not include the data required to specify the variables corresponding to the data points and 2) it is not exactly correct when the variables do not all have the same number of values. Solving this equation for n yields

$$n \approx e^{\frac{\log(\frac{s}{p})}{v}}. \quad (2.10)$$

The maximum designed file size is 0.5 GB so that all of the data can easily be loaded into memory, which commonly has an upper limit of 16 GB in new computers. Even if there are several complex objects, which potentially would each have a 0.5 GB data, they could all fit in memory. Assuming double precision numbers (i.e. 8 bytes per number), Table 2.1 gives the number of points per variable rounded to the nearest integer for one to ten variables. According to the table, minimizing the number of variables is important so that the number of points per variable can be as high as possible, which leads to greater accuracy if the number of variables is constant. However, if the number of variables is too small, then the model will not accurately predict the levels either.

Table 2.1: Number of variables vs number of points per variable

This table demonstrates the trade-off between the number of variables (v) and the number of points per variable (n) in Equation (2.10) with $s = 0.5$ GB and $p = 8$ B.

v	1	2	3	4	5	6	7	8	9	10
n	67M	8192	406	91	37	20	13	10	7	6

At the very least, the attenuation of the complex object compared to a reference object depends on the source frequency and the source and receiver positions. Since both the source and receiver are in a 3D space, their positions alone seem to require six variables, and including frequency makes seven variables. However, since the object has a constant cross-section and is infinitely long in the y -direction, the model only needs the difference of the y -positions (i.e. Δy) instead of each y -position (Figure 2.6). In contrast, the Nord 2000 EM¹⁵ does not include Δy , but this choice is unjustifiable because the changes in the y -direction are just as large as the changes in the x -direction or the z -direction. Excluding

any of these variables would likely lead to large errors in the predicted attenuation. Thus, the hybrid method requires six variables (i.e. the height and perpendicular distance of the source and receiver to the complex object, the distance parallel to the complex object between the source and receiver, and the frequency).

There are several possible additional variables including the ground impedance before or after the complex object, the vertical sound speed gradient, and the location and geometries of additional objects. There are multiple reasons to not include any more variables. First, while these variable affect the attenuation of a barrier compared to no barrier, it is less clear that they significantly affect a smaller change between a complex barrier and a reference barrier. Second, since there are many options, choosing which one to include seems arbitrary. Third, EMs already have a method to incorporate these interactions, so doing so here would be like modeling something twice. Finally, and most importantly, including more variables reduces the resolution of all of the other variables because the total number of data points is limited. Thus, increasing the number of variables would likely decrease the accuracy instead of increase it.

§2.4.1.3 Variable definitions

The hybrid method uses the following variables, which are also illustrated in Figure 2.6:

- x_s is the displacement of the *source* from the object that is perpendicular to the object. This variable is always negative;
- x_r is the displacement of the *receiver* from the object that is perpendicular to the object. This variable is always positive;
- z_s is the height of the *source* above the bottom of the object;
- z_r is the height of the *receiver* above the bottom of the object;
- Δy is the distance between the source and the receiver that is parallel to the object
- f is the frequency of the source.




Figure 2.6: Input variables

These six input variables are calculated from the propagation path to interpolate the 2.5D BEM data.

§2.4.2 Deciding the source and receiver positions

Deciding the source and receiver positions is important because they directly affect the accuracy of the hybrid method. This task requires choosing three aspects for each variable: 1) the range, 2) the number of points, and 3) the distribution of those points. The source and receiver positions also depend on the data structure, but this discussion is postponed until §2.6.1 because the data structure directly impacts the choice of interpolation method.

First, the range of each variable depends on the application and balances two extremes. On the one hand, if the variable range is too small, then many paths will have to be extrapolated instead of interpolated, which reduces the accuracy. On the other hand, if the variable range is too large, then the limited number of available points will be unnecessarily spread out, which also reduces the accuracy. The exact range for each variable does not matter too much, but certainly it should be within a factor of two of the optimal value (i.e. 100 m versus 110 m is not very important but 100 m versus 200 m is important).

The lower limits for all of the variables (upper limit for x_s) are fairly set. Specifically, all of the variables are bounded by zero, and x_s and x_r have a slightly tighter boundary to make room for the complex object (e.g. $x_s \leq -1$ m and $x_r \geq 1$ m).

However, the other boundary is much more case dependent. One must judge what distance is best for the specific application from practical experience, but here are some guidelines for road and rail noise:

- x_s : The sources are usually less than eight lanes away for road noise and less than four tracks away for rail noise.
- z_s : The sources are usually less than 1 m high for road noise and less than 4 m high for rail noise. CNOSSOS-EU uses a height of 5 cm for road noise.
- x_r : This variable is the hardest one to choose. CNOSSOS-EU states that the source and receiver should not be more than 800 m apart, so that is the largest possible range, and the minimum is about 100 m. A range of 100-200 m is still likely to contain most receivers while not being unnecessarily large. One must decide considering where most receivers will be located and the area where the correction factor is significant.
- z_r : Multiple receiver heights are important. When people are outside, usually their ears are less than 2 m high, but they can also be inside buildings, which are usually less than ten stories tall. CNOSSOS-EU compromises by modeling noise maps at 4 m high.
- Δy : It does not have to be as big as x_r because complex objects are usually parallel to the noise source, so the noise from the source with the smallest Δy is likely to be the largest contribution. 50 m is likely a good starting point, but this variable could also be decided by taking one-half to one-fourth of x_r .

Second, choosing the number of points depends on the range of each variable. All else being equal, a variable with a larger range would require more points than a smaller range. In addition, the number of points depends on how the value changes with respect to each variable. Assuming linear interpolation for concreteness, a variable may need more points if the absolute value of the second derivative of the correction with respect to that variable is large compare to the same quantity for the other variables. For example, if the correction does not change with respect to a variable, then the table only needs one value for that variable. If the correction changes linearly (i.e. has a constant slope), then only two values are necessary. However, if the slope is not constant (i.e. the second derivative is non-zero), then more points are required to accurately represent the underlining function. Unfortunately, the second derivative is usually unavailable. This analysis assumes linear interpolation for concreteness, but using a different interpolation method would produce a similar conclusion. The difference is that using a lower order interpolation method like nearest neighbor would require lower order derivatives and using a higher order interpolation method like cubic would require higher order derivatives. However, all of these derivatives are generally unknown. However, practical experience suggests that the sound level is more sensitive to the source and receiver heights than their distance to or along the barrier, especially in the region where the path length difference is near zero (recall Figure 2.5). Furthermore, CNOSSOS-EU also suggests putting up to 5 m between receivers horizontally, but the height must be known

to the nearest 10 cm. This recommendation suggests a higher concentration of points for the heights of the sources and receivers compared to the other variables.

Third, the distribution of the points also depends on the derivatives of the correction, but again this information is generally unknown. Assuming the points are on a grid so that each dimension can be treated separately, the following are potential spacing options:

- Linear;
- Exponential;
- Optimized.

The linear spacing is characterized by the function $x_n = c_1 n + c_0$ where x is the location of the point, n is the index of the point, and the coefficients are constants. The linear spacing works well when the second derivative is about the same size for the whole range and the number of points is fairly large. However, neither of those conditions is met for this application.

The exponential spacing is characterized by the function $x_n = e^{c_1 n} + c_0$. This spacing allows the points to be concentrated close to the complex object, but they are likely too close together near the complex object and too far apart far from the complex object (e.g. ten points between 0.1 m and 1 m, between 1 m and 10 m, and between 10 m and 100 m).

The optimized spacing is the most rigorous way to find the best spacing. Ideally, the point locations would be optimized by minimizing the overall interpolation error. However, this option is too computationally expensive. Calculating the approximately one million points at twenty frequencies for the table of corrections is already quite expensive but manageable. Reference data would require many more points because otherwise the optimization would just move the data points close to the reference points, which might give a low error compared to the reference data but potentially high error for a different reference dataset. Chapter 4 goes into great detail on how to optimize the number of points and their locations by minimizing the uncertainty of a predicted value using a Gaussian process. For the rest of this chapter, just assume that source and receiver positions have been optimized or chosen by hand as well as the user can.

§2.4.3 Choosing the frequencies and geometries

After the source and receiver positions have been determined, a few more variables must be established. Specifically, the exact frequencies and the geometries. EMs define the frequency bands (e.g. 50 Hz-5 kHz in third-octave bands). However, since BEM makes single frequency computations and not third-octave band calculations, the frequency bands must be approximated by numerically integrating each frequency band. The number of frequencies per frequency band should be sufficient to converge and the spacing of the frequencies should be taken into account when performing the integral. For

example, one could use ten or more uniformly spaced frequencies within each third-octave band.

The geometry is usually hard, flat ground with either the complex object or the simplified object. The complex object is well-defined and is provided by the user. Its main restriction is that the object must have a constant cross-section. However, the simplified object needs a little explanation. The reference object should be an object that EMs can model and that is as similar to the complex object as possible. For example, if the complex object is a T-barrier, then the reference object is an I-barrier with the same height as the T-barrier. “Similar” is intuitively meaningful but is slightly ambiguous. Making a technical definition such as “having the same path length difference” is not useful because the path length difference depends on the source and receiver positions and because it is unlikely to produce more accurate results. The main point is that the smaller the difference between the complex and reference objects, the greater the accuracy of the hybrid method.

A further complication is that the reference object in EM might be slightly different than the reference object in 2.5D BEM. For example, thin barriers have a width in 2.5D BEM but not in EM. Having two definitions of the reference object does not cause any problems as long as they are similar. Furthermore, the two definitions provide a small opportunity to mitigate the error between EM and 2.5D BEM for the simplified object. To reduce the error, minimize the error between the levels predicted by EM and 2.5D BEM with respect to the shape of the simplified object in 2.5D BEM. To reduce the number of BEM computations, the shape variables could be reduced to just the width of the object. However, this optimization would likely produce a negligible improvement for a substantial cost, so it is not implemented in this dissertation.

Now, given the source and receiver positions, the frequencies, and the geometries, one could produce the results for the two main cases: the complex case and the reference case. This dissertation does not explain the BEM computations any further because there are entire books dedicated to BEM^{67–69}. Next, §2.4.4 details why the complex and reference geometries are modeled on hard, flat ground and illustrates an example where this approach does not work.

§2.4.4 Using flat, hard ground

For the BEM computations, the complex and reference objects are on an infinite, hard ground, which must be justified. First and most importantly, this surrounding is efficient. Using BEM with this type of ground means that none of the ground has to be meshed, which dramatically reduces the cost. In addition, if each ground type and the locations of those ground types required a different table of corrections, then the number of data

tables would grow very rapidly. Second, a flat, hard ground accurately represents many urban grounds (e.g. concrete and asphalt). Moreover, Figure 2.7 illustrates that the attenuation due to a complex object compared to a reference object is mostly invariant for extremely different ground impedances. In particular, the changes between the different source positions are much larger than the changes between the different ground conditions.




Figure 2.7: Comparison of T-top correction for hard and soft ground

There are four different source positions, which are shown by the black circles, and two ground types (hard and grass with $\sigma = 80 \text{ kPa}\cdot\text{s}\cdot\text{m}^{-2}$), which is shown by a black or green ground, respectively. The results are 2D results for 50 Hz – 5 kHz using the spectrum from a car (engine and tire noise) at 50 km/h from the Hosanna Project¹²⁰ (Task 2.3, p 43-44). The individual third-octaves have larger differences, but these plots demonstrate that the differences are not systematic.

However, flat, hard ground does not work in all cases. Specifically, it does not work when the ground is not approximately flat (Figure 2.8). These cases must be treated as special cases and will be more costly to compute, but there is no theoretical reason why the hybrid method could not also be applied to these cases given an appropriate reference geometry.




Figure 2.8: Problems when the ground is not flat

Using a flat ground is problematic for this situation where the ground is not approximately flat. The pink area would not have a valid correction term unless this situation is treated as a special case. The green area would also be incorrect because the ground effect would be very different.

§2.4.5 Post-processing

The post-processing of the BEM data has four main steps:

1. Convert single frequency calculations into frequency band calculations;
2. Subtract the levels for the reference case from the complex case;
3. Order the points;
4. Calculate the decay function, which §2.7 defines.

First, numerically integrating the single frequency calculations yields the frequency band levels. Using equally spaced frequencies within each frequency band, the integral is approximately

$$L_f = 10 \log_{10} \left(\frac{f_{\max} - f_{\min}}{N} \sum_{i=1}^N \frac{|p_i|^2}{p_{\text{ref}}^2} \right) \quad (2.11)$$

where f_{\max} and f_{\min} are the maximum and minimum frequencies of the frequency band centered at f , N is the number of single frequency calculations in the frequency band, p_i is the complex pressure calculated with 2.5D BEM, and $p_{\text{ref}} = 20 \mu\text{Pa}$.

Second, after the numerical integration is complete, subtracting the levels for the reference case from the levels for the complex case produces the corrections, ΔL . This definition is the negation of insertion loss. For notational consistency, the attenuations in the data table are called the table of corrections (ΔL) versus the extra attenuation (A_{extra}), which is interpolated from the table of corrections.

Third, ordering the table of corrections allows them to be searched very efficiently for a certain value. Specifically, there are five position columns – i.e. one for each position variables $\mathbf{x} = (x_s, z_s, x_r, z_r, \Delta y)$ – followed by one column for each frequency band from

least to greatest where each row contains a unique position and all of the frequencies (Table 2.2). The positions are ordered from least to greatest where $\mathbf{x} > \mathbf{x}'$ if $x_s > x_s'$. If $x_s = x_s'$, then $\mathbf{x} > \mathbf{x}'$ if $z_s > z_s'$. If $x_s = x_s'$ and $z_s = z_s'$, then $\mathbf{x} > \mathbf{x}'$ if $x_r > x_r'$. This pattern continues for z_r and finally Δy .

In general, the row position (n) starting from zero is given in terms of the number of values for each variable (\mathbf{M}), the index starting from zero of the value of each variable (\mathbf{m} where $m_i \in [0, M_i - 1]$), and the number of dimensions (k) as

$$n = \sum_{i=1}^k \left(m_i \prod_{j=i+1}^k M_j \right) \quad (2.12)$$

where any product over an invalid range is one (e.g. from $k + 1$ to k because the values are decreasing and M_{k+1} does not exist). For example, for $k = 5$

$$n = m_1 M_2 M_3 M_4 M_5 + m_2 M_3 M_4 M_5 + m_3 M_4 M_5 + m_4 M_5 + m_5. \quad (2.13)$$

Figure 2.9 gives a 2D example where $k = 2$ and $\mathbf{M} = [5 \ 4]$. The plot gives \mathbf{m} and n as $(m_1, m_2) \rightarrow n$. In practice, BEM can output the values in this order so that the values do not need to be sorted later.




Figure 2.9: Order grid example for 2D data

This plot shows how to order the points in Table 2.2 with only two position variables. Each point has the following information: $(m_1, m_2) \rightarrow n$. In this case, Equation (2.12) simplifies to $n = m_1 M_2 + m_2$.

Fourth, the decay function needs to be calculated and put in the header information of the data file, but the details of computing the decay function are postponed until §2.7 because the process requires multiple steps.

Table 2.2: Fictional example of table of corrections

This table demonstrates how to order all of the rows of the table based on the five position variables. First, iterate through all of Δy , then z_r through x_s . All of the frequencies are stored in the same row. The table shows that each row (i.e. source/receiver pair) requires 26 numbers (i.e. 5 position variables and 21 frequencies).

x_s (m)	z_s (m)	x_r (m)	z_r (m)	Δy (m)	ΔL (dB)		
					50 Hz	...	5 kHz
-2	0	1	0	0	-1	...	-5
-2	0	1	0	1	-1	...	-5
-2	0	1	1	0	-1	...	-5
-2	0	1	1	1	-1	...	-5
-2	0	2	0	0	-1	...	-5
-2	0	2	0	1	-1	...	-5
-2	0	2	1	0	-1	...	-5
-2	0	2	1	1	-1	...	-5
-2	1	1	0	0	-1	...	-5
-2	1	1	0	1	-1	...	-5
-2	1	1	1	0	-1	...	-5
-2	1	1	1	1	-1	...	-5
-2	1	2	0	0	-1	...	-5
-2	1	2	0	1	-1	...	-5
-2	1	2	1	0	-1	...	-5
-2	1	2	1	1	-1	...	-5
-1	0	1	0	0	-1	...	-5
-1	0	1	0	1	-1	...	-5
-1	0	1	1	0	-1	...	-5
-1	0	1	1	1	-1	...	-5
-1	0	2	0	0	-1	...	-5
-1	0	2	0	1	-1	...	-5
-1	0	2	1	0	-1	...	-5
-1	0	2	1	1	-1	...	-5
-1	1	1	0	0	-1	...	-5
-1	1	1	0	1	-1	...	-5
-1	1	1	1	0	-1	...	-5
-1	1	1	1	1	-1	...	-5
-1	1	2	0	0	-1	...	-5
-1	1	2	0	1	-1	...	-5
-1	1	2	1	0	-1	...	-5
-1	1	2	1	1	-1	...	-5

§2.5 Local source and receiver positions

The previous section details how to precompute the table of corrections using 2.5D BEM, and the next section explains how to interpolate the data. However, to interpolate the data, first the source and receiver locations must be converted from their values in the full problem (i.e. the global coordinates) to their values in the BEM problem (i.e. the local coordinates). These coordinates are likely not the same because the location of the origin and the orientation of the coordinates is likely different. Thus, the conversion requires a translational and rotational transformation. Moreover, the BEM scene does not contain any other objects (e.g. buildings), so before applying the transformations, the source or receiver positions might change to incorporate reflections or diffractions. This section details how to transform the source and receiver positions for four cases:

1. Single diffraction;
2. Multiple diffraction;
3. Reflection;
4. Lateral diffraction.

In practice, paths that are combinations of these cases can be broken down into these four cases.

§2.5.1 Single diffraction

Consider a path that diffracts over one object. If the object is not complex, then $A_{\text{extra}} = 0$ dB and the local variables are not calculated. If the object is complex, then the global coordinates of the source and receiver must be converted to the local coordinates.

Since EMs search for paths in a 2D horizontal plane and then calculate attenuations in the 2D vertical plane, EMs have two sets of coordinates for each point. The global coordinates $[\mathbf{r}_{\text{global}} = (x_{\text{global}}, y_{\text{global}})]$ are the coordinates for the whole scene and therefore do not necessarily have a special relationship to any of the objects in the scene. The path coordinates $[\mathbf{r}_{\text{path}} = (d_{\text{path}}, z_{\text{path}})]$ are different for each path where d_{path} is the distance along the path and z_{path} is the height of the path above the ground. The source, diffraction point, receiver, and the complex object's ends are respectively denoted by S, D, R, B_1 , and B_2 in Figure 2.10.




Figure 2.10: Single diffraction

Single diffraction includes cases that diffract over only one horizontal edge without reflections or lateral diffractions. This case is the simplest case and all of the more complicated cases are converted to this setup.

Before finding the local coordinates, the angle θ must be found first (Figure 2.10). The angle $\theta < 90^\circ$ is the smallest angle between the propagation path and the complex object. A slight modification of the dot product yields

$$\theta = \cos^{-1} \left| \hat{\mathbf{r}}_{\text{global},D \rightarrow S} \cdot \hat{\mathbf{r}}_{\text{global},B_1 \rightarrow B_2} \right| \quad (2.14)$$

where $\hat{\mathbf{r}}_{\text{global},D \rightarrow S}$ is the 2D unit vector (the z -dimension is not included) parallel to the vector from D to S , and $\hat{\mathbf{r}}_{\text{global},B_1 \rightarrow B_2}$ is analogous from B_1 to B_2 . The absolute value guarantees that θ is less than 90° . To avoid dividing by zero, the vectors must have finite length, which physically means that the source, diffraction point, and receiver cannot all be collocated and the barrier must have a finite length. Then, from geometry the local coordinates are

$$x_s = [d_{\text{path}}(S) - d_{\text{path}}(D)] \sin \theta; \quad (2.15)$$

$$z_s = z_{\text{path}}(S); \quad (2.16)$$

$$x_r = [d_{\text{path}}(R) - d_{\text{path}}(D)] \sin \theta; \quad (2.17)$$

$$z_r = z_{\text{path}}(R); \quad (2.18)$$

$$\Delta y = [d_{\text{path}}(R) - d_{\text{path}}(S)] \cos \theta. \quad (2.19)$$

These equations assume that d_{path} is increasing from the source to the receiver and $0^\circ < \theta < 90^\circ$ so that $x_s \leq 0$ m and $z_s, x_r, z_r, \Delta y \geq 0$ m.

§2.5.2 Multiple diffraction

When there are multiple diffracting objects in a propagation path, then the objects are ranked based on their path length differences (δ). The path length difference is the distance from the source to the diffraction point plus the distance from the diffraction point to the receiver minus the direct distance between the source and the receiver (Figure 2.11).

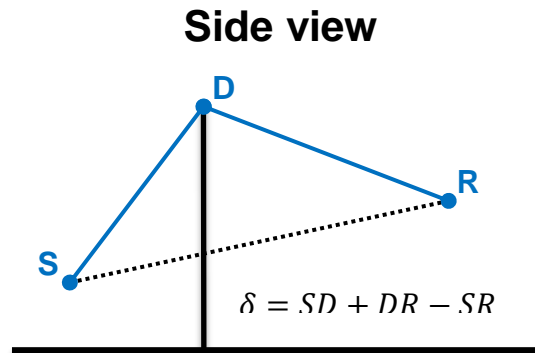


Figure 2.11: Path length difference

The path length difference is the distance SD plus DR minus SR and is the standard way to rank the importance of objects in outdoor sound propagation.




Figure 2.12: Multiple diffraction

Multiple diffraction includes paths that diffract over more than one horizontal edge. This figure illustrates the local variables when the left-hand barrier has the largest path length difference and is complex. Essentially, the multiple diffraction case is converted to the single diffraction case by ignoring barriers that do not have the largest path length difference when calculating the extra attenuation.

If a complex diffracting object has the largest path length difference, then the local variables are calculated the same way as in the single diffraction case (§2.5.1) ignoring the other diffracting objects (Figure 2.12). Otherwise, $A_{\text{extra}} = 0$ dB, so the local variables are unnecessary.

This approximation is very simple in that it is either on or off and only affects complex objects with the largest path length difference. The simplicity matches the simplicity of the multiple diffraction approximation in CNOSSOS-EU (§2.3.3). Other methods like Harmonoise⁴⁴ make more complicated approximations for the barriers, but most EMs do not. The primary reason for not using the Harmonoise diffraction approximation is that it would increase the required ranges of the sources and receivers in the data table, which would increase the distance between data points and would likely lead to a less precise value for A_{extra} .

The on/off nature of this approximation means that the pressure level is discontinuous at where the path length differences are equal (Figure 2.13). Specifically, the correction is applied where the complex object has the largest path length difference but is not applied where it does not have the largest path length difference, and these points could be adjacent near where two objects have equal path length differences. While undesirable, this discontinuity is unlikely to be significant in practice when there are many distributed sources that will wash out the discontinuities. Furthermore, the discontinuity would only be noticeable when the multiple diffraction path is the dominant contribution at the receiver. If one still desired to smooth this discontinuity, A_{extra} could be scaled by $\delta_{\text{complex}}/\delta_{\text{largest}}$ where δ_{complex} is the path length difference of the complex object and δ_{largest} is the largest path length difference of the entire path, which could be the path length difference of the complex object. However, this more complicated approximation is usually unnecessary.




Figure 2.13: Correction discontinuity

There is a discontinuity where the path length differences of the two barriers are equal (i.e. where the extra attenuation is being turned on or off).

§2.5.3 Reflection

EMs model propagation paths with reflections from vertical surfaces using image sources and receivers and an additional attenuation for the loss due to the reflection. The hybrid method also uses image sources and receivers. The calculations are identical to those in §2.5.1 except that the actual source or receiver may be replaced by the image source or receiver. The image of a 2D point (x and y) is that point reflected across a line (e.g. the intersection of the vertical plane that reflects the path and the xy -plane).

Specifically, if there is a reflection in the propagation path between the source and the complex object then the source becomes the corresponding image source, and if there is a reflection in the propagation path between the complex object and the receiver then the receiver becomes the corresponding image receiver. Figure 2.14 gives an example with an image receiver.




Figure 2.14: Reflection

Reflected paths use image sources and receivers to unfold the path so that it is straight as seen from above. The figure shows the local variables when the left-hand barrier is large and complex. Letting $S \rightarrow S'$ and $R \rightarrow R'$ converts this case to the single diffraction case.

Since r_{path} is the same for a point and its image, Equations (2.15)-(2.19) are unchanged. The only difference is in Equation (2.14) for θ where $\hat{r}_{\text{global},D \rightarrow S}$ potentially becomes $\hat{r}_{\text{global},D \rightarrow S'}$. However, an easier solution, which works for modeling lateral diffractions as well, is to let S become the closest propagation point on the source side. The propagation point could be the source, a reflection point, or a lateral diffraction point. For example, if the propagation path is $S \rightarrow D_1 \rightarrow D_2 \rightarrow R$ and the complex object is at D_2 , then $\hat{r}_{\text{global},D \rightarrow S}$ in Equation (2.14) would be $\hat{r}_{\text{global},D_2 \rightarrow D_1}$. This approach simplifies the

calculations because the location of the image source does not need to be calculated. The vector is a unit vector, so only the direction is important.

Paths can reflect off of the complex object, but unless the path diffracts over the complex object $A_{\text{extra}} = 0$ dB.

§2.5.4 Lateral diffraction

For calculating the local source and receiver positions, lateral diffractions (i.e. diffractions around vertical edges) are treated exactly the same as reflections (§2.5.3). Paths can diffract around the complex object, but unless the path diffracts over the complex object $A_{\text{extra}} = 0$ dB. However, the cases look slightly different because reflections use an image source or receiver whereas lateral diffractions extend the path from the diffraction point on the complex object in a straight line through the lateral diffraction point the same distance as the rest of the path beyond the diffraction point (Figure 2.15).




Figure 2.15: Lateral diffraction

Laterally diffracted paths diffract around one or more vertical edges. This case is converted to the single diffraction case by straightening the propagation path as seen from above while keeping the length of the propagation path constant.

The justification for this approach to lateral diffractions is relatively tenuous compared to the approach to reflections. For reflections, using image sources is well established across many different models but there is more than one common way to model diffraction. There are approaches that model diffraction by putting additional sources on the diffraction edge so that in Figure 2.15 R would become D_2 . However, the hybrid method extends the path past the lateral diffraction point to be consistent with EMs, which use the

same idea when calculating the attenuation from diffraction around a thin noise screen. Thus, the primary reason for this approach is model consistency.

§2.6 Interpolating the table of corrections

Interpolation methods estimate the value of a function based on data at any given point inside the dataset. The hybrid method uses an interpolation method to estimate A_{extra} based on the data provided by BEM. This section discusses two major decisions: the most appropriate data structure (§2.6.1) and interpolation method (§2.6.2). Then, (§2.6.3) details how to implement the chosen interpolation method using data with the specified structure.

§2.6.1 Choosing a data structure

The data structure is important because it affects the efficiency of searching and interpolating the data. There are several possible data structures that all have their positives and negatives. Four widely used examples (Figure 2.16) are

- Scattered or unstructured;
- 2^k -trees (e.g. quadtrees in 2D);
- Uniform rectangular grid;
- Non-uniform rectangular grid.




Figure 2.16: Comparison of data structures

Scattered data has no predefined structure. Quadtree is formed by bisecting rectangles (i.e. four points become nine points) as necessary. Uniform grid data must be on a rectangular grid and have the same distance between each data point along each axis. Non-uniform grid data must also be on a rectangular grid but can vary the distance between the data points along each axis. All of the plots have about the same number of points.

A data structure is distinct from but closely related to the distribution of the points from §2.4.2. The data structure refers to how the data is organized and the distribution refers to how the points are (or are not) concentrated. Using randomly generated points and a scattered data structure, the point distribution could be, for example, uniform or Gaussian

about the origin. In addition, points on a grid could be linearly, exponentially, or randomly distributed. Figure 2.17 illustrates scattered and gridded data that use random points from uniform and Gaussian distributions. Changing the data structure affects the efficiency of searching and interpolating the data, but changing the data distribution does not affect it.

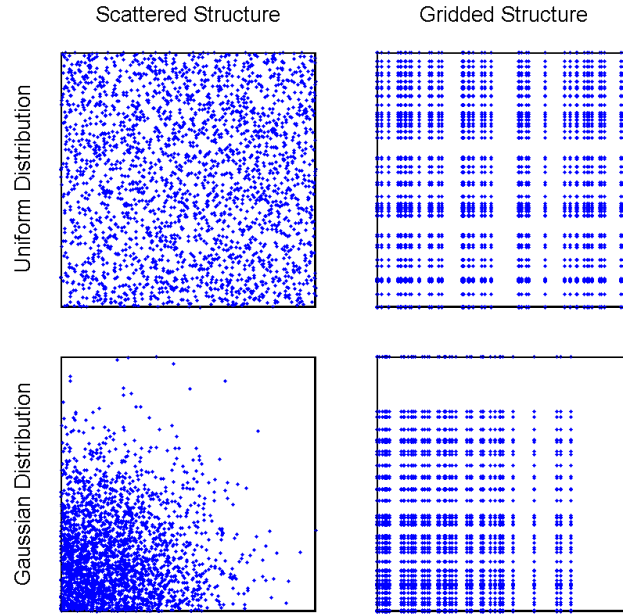


Figure 2.17: Data distribution vs structures

Data structures and distributions are distinct but related. To highlight the difference, this plot gives two data structures and two data distributions. The left-hand plots use a scattered structure, and the right-hand plots use a gridded structure. The top plots use random data from a uniform distribution, and the bottom plots use random data from a Gaussian distribution.

For this application, the efficiency of searching and interpolation are especially important because there are so many points. Evaluating each data structure for feasibility requires an approximate number of data points, which is limited by the maximum file size. Ignoring any header information and considering the file as a large 2D matrix, the file size (s) is

$$s = pn_{\text{rows}}n_{\text{cols}} \tag{2.20}$$

where p is the precision, which is 8 bytes (B) for standard double precision numbers; n_{rows} is the number of rows in the data file, which is the number of unique source/receiver pairs; and n_{cols} is the number of columns in the data file, which is 26 (i.e. 5 position variables and 21 frequencies) according to Table 2.2. Solving Equation (2.20) for n_{rows} yields

$$n_{\text{rows}} = \frac{S}{pn_{\text{cols}}}. \quad (2.21)$$

Assuming that the maximum file size is 0.5 GB as in §2.4.1.2 and noting that 1 GB = 1024^3 B instead of 10^9 B produces 2.6×10^6 source/receiver pairs. Assuming points on a grid and an equal number of points in each of the five dimensions gives approximately $\sqrt[5]{2.6 \times 10^6} \approx 19$ points per dimension.

Data without a set structure (i.e. scattered data) certainly gives complete freedom but requires the computational overhead to organize the data. For example, a Delaunay triangulation is a common way to organize scattered data. In 2D, a Delaunay triangulation groups points into non-overlapping triangles where the smallest angle in each triangle has been maximized. A geometrical way to find these triangles is to draw a circle through all of the points of a proposed triangle (Figure 2.18). If and only if no other point is inside the circle, then proposed triangle is a Delaunay triangle. Applying this procedure to the entire data set creates the Delaunay triangulation.




Figure 2.18: Delaunay triangulation

There are four points that can be connected to create two Delaunay triangles. The two triangles on the left are Delaunay because the last data point is not inside the circle through the points of the triangle. The two triangles on the right are not Delaunay because the last data point is inside the circle.

This process can be applied in any number of dimensions but the cost is larger for higher dimensions. Figure 2.19 demonstrates that calculating the Delaunay triangulation with Qhull^{121,122}, a free C library for calculating convex hulls, for 5D data points on a uniform grid requires $O(N^{4/3})$ computation time and $O(N)$ memory. Recall that this application uses about 20^5 source/receiver pairs, which would require about 109 minutes and 97 GB to calculate. Searching for points and interpolating would be an additional cost. Therefore, scattered data is too computational expensive.




Figure 2.19: Qhull performance versus number of data points

The computation time increases super linearly $O(N^{4/3})$, and the memory increases approximately linearly $O(N)$ for 5D data points on a uniform grid. Randomly spaced points may require even more time.

2^k -trees are better known as quad-trees because this approach is often applied to 2D images. In 2D, the entire rectangular domain is defined by just its four corners. If the interpolation error is too large at a few sample points, then the rectangle is broken up into four pieces, which is why they are called quad-trees. The same process is repeatedly applied to each sub-rectangle until the interpolation error is small enough everywhere or a certain tree height is reached. The structure is a tree because each rectangle is attached to a parent rectangle and four children rectangles such that the number of children grows with each generation just like a family tree.

2^k -trees introduce some structure into the data while allowing points to be concentrated where necessary. However, since this application is 5D, it requires 32-trees instead of quad-trees. While there are many C/C++ libraries for 2D and even 3D, unfortunately the author was unable to find a C/C++ library for 5D. In addition, developing such a library is well beyond the scope of this dissertation, so this option could not be pursued further.

Although, assuming a C/C++ library did exist, there would still be several issues to consider. First, there would be some overhead cost of setting up the 2^k -tree, which would need to be evaluated against its benefits. Since this application is 5D, the height of the tree would be fairly short. Memory restrictions limit the total number of points to about 1 million, which means for a full tree the height could only be four ($32^4 \approx 1$ million). Moreover, the memory restrictions also make deciding where to concentrate the points more difficult. With only 20 points per dimension that are far apart compared to a wavelength, it is also likely that there are not enough points to effectively concentrate the points, which means that the points would be mostly uniformly spaced. Lastly, interpolation in a 2^k -tree would require special attention to guarantee that there are not any discontinuities traveling from one subregion to the next.

Next, a uniform rectangular grid can search and interpolate in constant time, so it is very fast, and has very little overhead costs when the data is sorted. The main disadvantage of

a uniform rectangular grid is that the structure is very limiting. This data structure does not allow any concentration of points.

Finally, consider a non-uniform rectangular grid. This option requires almost no overhead cost; instead, the data just need to be in a specified order, and searching for a point is very fast, $O(\log_2 N)$. The biggest downside is that this structure drastically limits where one can put the points. However, it provides more flexibility than a uniform grid, and even more flexible structures are too expensive (i.e. scattered data) or unavailable (i.e. 2^k -trees), so a non-uniform rectangular grid is the best option.

§2.6.2 Choosing an interpolation method

No interpolation method is best for all applications, so the details of this application are important. Specifically, this application has a large number of data points ($N \approx 10^6$). Thus, the interpolation method must be able to be efficient even with a large number of data points. In addition, this interpolation problem is 5D. Spatial interpolation problems usually are only 2D or 3D, and the methods developed for a 2D problem may not be applicable to a 5D problem. Lastly, for this application the data points have a well-defined structure instead of scattered data, so the interpolation method should be able to take advantage of this structure.

There are many types of interpolations methods¹²³; a few examples are:

- Inverse distance weighted;
- Natural neighbor;
- Regression;
- Spline.

Inverse distance weighting¹²⁴ calculates a weighted sum of the data points. The weights are based on the inverse distance between the interpolation point and the data points where the distance function is

$$d(\mathbf{x}, \mathbf{x}', p) = \left[\sum_i (\mathbf{x}_i - \mathbf{x}')^p \right]^{1/p} \quad (2.22)$$

where the normal Euclidian distance function uses $p = 2$. The model can be adjusted by altering the weights (e.g. setting $p = 4$ instead of $p = 2$) or including more or fewer surrounding points. The main reason against using inverse distance weighting is that it does not take advantage of the data being on a rectangular grid like the spline methods do and is therefore slightly more expensive for this case.

Natural neighbor¹²⁵ also calculates a weighted sum of the data points, but the weights are based on an area instead of a distance. A Voronoi diagram is a way to break up a region into cells that each contain only one data point and the subregion that is closest to that data point. The weights in natural neighbor are related to the cell areas of the Voronoi diagrams with and without the interpolation point. The main objection to using natural neighbor is that the Voronoi diagrams are expensive to compute in 5D with so many data points.

Regression is fitting a function to the entire (continuous) dataset. Regression often uses a simple polynomial as the function but also includes non-parametric functions like Gaussian process regression¹²⁶. Regression is different from all of the other methods because the function usually does not go through all of the points. The main problem with regression is that it requires a matrix inversion that is very costly for a large number of parameters or large datasets.

Splines¹²⁷ are piecewise approximations. Nearest-neighbor, linear, and cubic interpolation are all examples of splines. Splines are efficient even for a large dataset because they only use a very small subset of the data, can be used for any number of dimensions when an appropriate function is selected, and take advantage of the structure of the data. However, splines are really a family of interpolation methods, so the next choice is which one to use.

Nearest-neighbor is very fast but is not continuous. Linear interpolation is almost as fast (i.e. requires a 32x32 matrix to be multiplied by a vector) and is continuous, but its derivatives are discontinuous. Cubic interpolation allows the function and its first derivative to be continuous, but the author could not find or easily derive a 5D version. Even if such a formulation exists, cubic interpolation would be more expensive because it would require inverting a 1024x1024 matrix. Thus, linear interpolation provides the best compromise between efficiency and accuracy for this application.

§2.6.3 Linear interpolation on a non-uniform grid

Linear interpolation has two parts: 1) finding the closest data points that contain the interpolation point and 2) linearly interpolating based on those data points.

§2.6.3.1 Finding the local data points

Given a large k -dimensional dataset on a non-uniform rectangular grid, the goal is to find the data points that define the smallest hyperrectangle that contains a given interpolation point. The most efficient algorithm for this task is a kd -tree, which can find the points in $O(\log_2 N)$ time. This section explains how this algorithm is applied to this case. The whole process has the following steps:

1. For each axis, obtain an ordered (i.e. least to greatest) list of the distinct values;
2. For each axis, use a binary search to find the index of the largest value that is less than or equal to the interpolation point for that axis;
3. Permute the indices and the indices plus one to obtain all 2^k points.

The first step is to obtain all of the distinct values in order from least to greatest for each axis. This step is straightforward because of how the data is stored. Since the data is on a rectangular grid, let \mathbf{X} be a k -dimensional matrix where each entry is a vector \mathbf{x} of the independent variables of the corresponding data point, e.g. $\mathbf{x} = (x_s, z_s, x_r, z_r, \Delta y)$. Moreover, \mathbf{X} is defined such that x_1 (i.e. the first value in \mathbf{x}) increases along the first dimension of \mathbf{X} , x_2 increases along the second dimension, and so forth. Then, to obtain an ordered list of the distinct values for each independent variable, simply traverse the matrix along each dimension. Since the matrix was already ordered, the time for this step is proportional to the number of distinct values.

Second, for each of the k dimensions use a binary search to find the index of the largest value that is less than or equal to the interpolation point and not the last index along that dimension. Essentially, the binary search algorithm evaluates a point in the middle of the dataset and determines if the value is too small, just right, or too big. If the value is not right, then since the data are ordered half of the data points are eliminated. The algorithm keeps cutting the number of points in half until it finds the right point.

For example, let a be the value of the interpolation point for a dimension and let \mathbf{b} be a vector of the n distinct values in that dimension ordered from least to greatest with indices 0 to $n - 1$. In addition, let there be three integers $i_{\min} = 0$, $i_{\text{mid}} = \text{floor}(\frac{n-1}{2})$, and $i_{\max} = n - 1$, which are minimum, middle, and maximum possible indices for the desired index.

For this application, the binary search algorithm has the following steps:

1. If $b_{i_{\min}} \leq a \leq b_{i_{\max}}$, continue but otherwise return an error;
2. If $b_{i_{\text{mid}}} \leq a$, then calculate $b_{i_{\text{mid}+1}}$;
 - a. If $b_{i_{\text{mid}+1}} \geq a$, then terminate and return i_{mid} ;
 - b. If $b_{i_{\text{mid}+1}} < a$, then set $i_{\min} = i_{\text{mid}}$ and $i_{\text{mid}} = \text{floor}(\frac{i_{\min}+i_{\max}}{2})$ and go to step 2;
3. If $b_{i_{\text{mid}}} > a$, then set $i_{\min} = i_{\text{mid}}$ and $i_{\text{mid}} = \text{floor}(\frac{i_{\min}+i_{\max}}{2})$ and go to step 2.

This algorithm does not allow i_{mid} to be equal to $n - 1$ so that there is always at least one index greater than i_{mid} , which is important for the next step. This step runs in $O(\log_2 N)$ time

Now, the algorithm has identified the largest index that is less than or equal to the interpolation point, which means that incrementing any combination of the indices by one would provide a new point of the box that contains the interpolation point. There are 2^k combinations including the original point. The time of this step is independent of the total number of data points.

§2.6.3.2 Interpolating the data

The difficult part of interpolating the data is that it is 5D. Previous literature¹²⁸ used k -linear interpolation without giving the equations explicitly. A k -linear interpolation equation is

$$y(\mathbf{x}) = \sum_{m_1, m_2, \dots, m_k=0}^1 c_{m_1, m_2, \dots, m_k} x_1^{m_1} x_2^{m_2} \dots x_k^{m_k}. \quad (2.23)$$

The sum is over all combinations of the m 's being zero and one, and c_{m_1, m_2, \dots, m_k} represents 2^k coefficients. For example, the bilinear interpolation equation is $y(\mathbf{x}) = c_{0,0} + c_{1,0}x_1 + c_{0,1}x_2 + c_{1,1}x_1x_2$. The coefficients come from evaluating the equation at the corners of the hyperrectangle that contains the interpolation point. By converting to local coordinates between zero and one, the matrix equation for the coefficients can be solved analytically for any number of dimensions, so finding the coefficients does not require a matrix inversion. Specifically,

$$\mathbf{c} = \mathbf{T}_k \mathbf{y} \text{ where } \mathbf{c} = \begin{bmatrix} c_{0,0,\dots,0} \\ c_{0,0,\dots,1} \\ \vdots \\ c_{1,1,\dots,1} \end{bmatrix} \text{ and } \mathbf{y} = \begin{bmatrix} y(0,0, \dots 0) \\ y(0,0, \dots 1) \\ \vdots \\ y(1,1, \dots 1) \end{bmatrix} \quad (2.24)$$

and \mathbf{T}_k is given by the following recursive relationship:

$$\mathbf{T}_0 = 1 \text{ and } \mathbf{T}_{n+1} = \begin{bmatrix} \mathbf{T}_n & 0 \\ -\mathbf{T}_n & \mathbf{T}_n \end{bmatrix}. \quad (2.25)$$

The order of \mathbf{c} and \mathbf{y} is important, so to clarify the subscripts make a binary number that gives \mathbf{c} and \mathbf{y} a unique order. After converting the interpolation point to local coordinates, it can be plugged into Equation (2.23) to yield the interpolated value.

Except for in this section, this dissertation omits the prefixes (e.g. using linear instead of bilinear in 2D or pentilinear in 5D) because they usually do more to obscure the meaning than to clarify it. Although, linear interpolation in 2D technically uses $y(\mathbf{x}) = c_{0,0} + c_{1,0}x_1 + c_{0,1}x_2$ and bilinear uses $y(\mathbf{x}) = c_{0,0} + c_{1,0}x_1 + c_{0,1}x_2 + c_{1,1}x_1x_2$. From the context (i.e. interpolation on a rectangular grid), the meaning should be clear.

§2.7 Extrapolating the table of corrections

Since the data table that is calculated with BEM has a finite range, an interpolation point might be outside of this range. There are two cases when this problem may occur. First, x_s or x_r cannot have a range that includes zero because the object must have a finite width in the BEM computations, so x_s or x_r could have values that are too small to be within the data. Generally, the level where the source or the receiver is within 1 m of the barrier is not of interest. Second, the interpolation point could be far from the barrier in any direction. In this case, this long path is unlikely to be the main contribution to the sound level at that location unless it is the direct path, so again the extrapolated value is likely unimportant. In addition, extrapolation is notoriously difficult because the data may not accurately represent the unexplored region. For this reason, the ranges of the variables should be sufficiently large to handle most cases. However, since an interpolation point could be outside of the data table, which would make it an extrapolation point, this section details an extrapolation method.

The simplest approach would be to set the extra attenuation to zero anywhere outside of the data. This approach is very fast but is discontinuous. For most applications, this discontinuity is insignificant, but there is a more rigorous approach that does not cost much extra.

The suggested approach is to find the closest data point and multiply the value at that data point by a decay function that ranges from one at the data point to zero far from the data point. This approach removes the largest discontinuity and is still fairly efficient, $O(\log_2 N)$. There could still be a discontinuity at the points that are equidistant from two data points that have different values, but this discontinuity is even less important than the previous discontinuity, and removing it would be more expensive.

Finding the decay function, which would occur during the post-processing of the BEM data, requires the following steps:

1. Take a sample of the BEM data;
2. Calculate the distance between each of the points and the absolute value of the difference of the corrections for each frequency;
3. Sort the points by their distance;
4. Smooth the data using a moving average for each frequency;
5. Calculate one minus the value of the smoothed data normalized by the maximum value for each frequency;
6. Fit all the data to a curve.

The first step is to take a sample of the data because using all of the data is too memory intensive for the steps that follow (i.e. the process is $O(N^2)$). Choosing what points to include in the sample is important because the end result varies depending of what points are selected. Without any other information, the best option might appear to select the points at random, but this approach usually produces points that are far from each other, which is not useful for modeling a relationship that varies rapidly with distance. All of the points are dissimilar, so the similarity between the points as function of distance cannot be accurately represented. The next logical option is to choose the points that are closest together, but this choice solves one problem by creating another. The points that are most densely packed are also next to the complex object and the ground. Since this relationship between points varies throughout space, the points next to the complex object or even the ground to some extent are not very representative of the majority of dataset. Thus, the sample should balance between points that are close together and points that are far from objects. For example, the source and receiver could be the closest points that are at least 5 m from the complex object and at least 1 m high with six points in each direction for a total sample size of $6^5 = 7776$. These numbers are a general guideline and may change depending on the complex object. Finite computer memory limits the sample size.

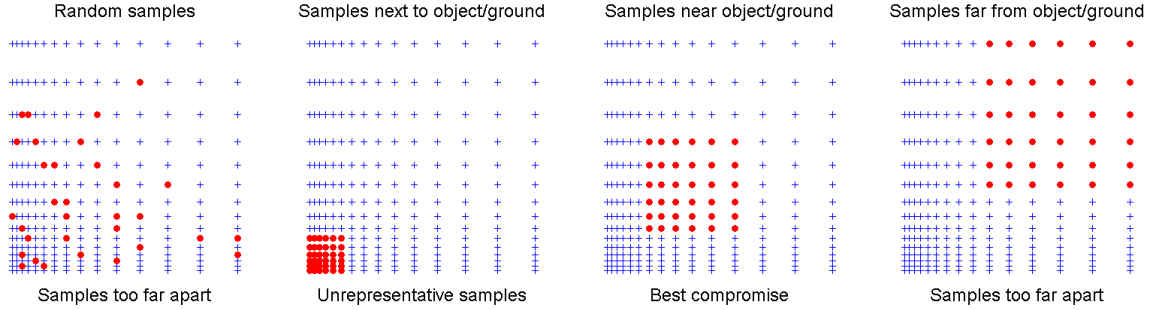


Figure 2.20: Choosing samples to model extrapolation

Here are four different sampling options to estimate how similar points are as a function of distance. The random samples and samples far from the object/ground are so far apart that most of the points are dissimilar. The samples next to the object/ground are unrepresentative of the extrapolation regions far from the complex object/ground because the correction factors change more rapidly in this region. The samples near but not right next to the object/ground provide the best compromise between samples being close to each other and samples being representative of the exterior region.

The second step calculates two relationships between each of the points and joins them as an $(d, \Delta A)$ pair. The first quantity is distance between each of the points. This distance is the standard Euclidian distance extended to 5D:

$$d(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{\sum_{i=1}^5 (\mathbf{x}_{1,i} - \mathbf{x}_{2,i})^2} \text{ where } \mathbf{x} = (x_s, z_s, x_r, z_r, \Delta y). \quad (2.26)$$

The second quantity is the absolute value of the difference in the correction factors:

$$\Delta A(A_1, A_2) = |A_1 - A_2| \quad (2.27)$$

where this calculation is performed for each frequency to create a vector of attenuation differences (ΔA).

The third step sorts all of the points according to their distance. Sorting is a standard function in most computer languages, so no further details are given here. The sorting is required to group the points into similar distances in order to track how closely related the points are.

The fourth step smooths the data using moving average. For example, the width could be 10^5 points without overlap. The best width depends on the number of data points in the sample, and overlap is unnecessary with so many points. This step takes the data from looking like a cloud to a wiggly curve, which makes the normalization factor in the next step (i.e. ΔA_{\max}) much less sensitive to extreme data values. If ΔA_{\max} is much larger than most values, then the function will never decay close to zero, which would indicate that points are similar when they really are dissimilar.

The fifth step converts the attenuation differences into the decay factor. The decay factor must be one when the distance is zero and decay to zero as the distance is increased. The range of the attenuation differences could be $[0, \infty)$. To convert this range to the range from zero to one, the value is divided by the maximum value for that frequency and then subtracted from one. Specifically,

$$f(d) = 1 - \frac{\Delta A(d)}{\Delta A_{\max}} \quad (2.28)$$

where again this calculation is performed for each frequency.

The sixth and final step is to fit all of the data to a curve. Originally, there was a curve for each frequency. However, since the variability within each frequency was greater than the variability between frequencies, the many curves were simplified to one curve for all frequencies. The curve has the following form:

$$f(d) = \frac{1}{c_1 d + 1} \quad (2.29)$$

In general, there may be higher order terms in the denominator, but all of the coefficients must be greater than or equal to zero so that the function never increases with distance. In some cases, an exponential function could work as well.

Figure 2.21 gives an example with real data from a T-barrier (3 m high and 1 m wide), which is described in detail in Chapter 3. The variability is quite high for different samples, frequencies, and distances, so the optimized value is not very precise. This process is just to obtain an approximate shape for the decay function and might not even have to be repeated for every data set as a rule of thumb could be used instead (e.g. the values of c_1 tend to be between 0.5 and 1.5). As stated at the beginning of this section, the extrapolation function should not affect the level at most relevant receivers. If it does, then the range of the variables should be increased.




Figure 2.21: Decay function

The decay function is a rough approximation of how similar two points are as a function of distance and is given by fitting a curve to the data. Each of the curves represents a different third-octave frequency band and the thick black curve is the optimized curve.

After determining the decay function, then hybrid method can use it to extrapolate the data by multiplying the value of the closest data point by the decay function evaluated at the appropriate distance. The closest data point is found using a *kd*-tree that is similar to the one described in §2.6.3.1. The main difference is that the description of the point that needs to be found is slightly different. The previous algorithm searched for the point with the largest dimensions that were all less than or equal to the interpolation point whereas the current algorithm searches for the point that is closest to an extrapolation point. Thus, instead of checking if a data point is less than the interpolation point, it checks if the data point is closer than its neighbors one dimension at a time. In addition, since the point is known to be outside of the dataset, the searches can be optimized by checking the extreme data points first.

§2.8 Summary

This chapter describes the whole hybrid method in detail and justifies the most important decisions. The hybrid method provides a general framework for extending EMs to complex geometries and surfaces that are currently not allowed in EMs by interpolating the 2.5D BEM results.

Broadly speaking, the hybrid method has two parts: a reference method and an efficient method. The reference method is 2.5D BEM, which tabulates the effects of complex geometries in a very simplified surrounding (i.e. hard, flat ground) for a large number of source/receiver locations and frequencies. Compared to full 3D BEM, the hybrid method drastically reduces the computation time from prohibitive to manageable for slightly less accuracy.

The efficient method is EM, which approximates all the attenuations in the full scene except for the extra attenuations associated with the complex objects. Compared to EMs, the hybrid method trades slightly longer computation time for greater accuracy in the complicated cases. The link between these parts is a table of corrections, which comes from BEM and is interpolated based on variables from EM.

More specifically, the hybrid method has the following major steps:

1. Determine the source and receiver locations on a non-uniform grid;
2. Compute the levels for the complex and reference cases;
3. Integrate, subtract, and sort the results to obtain the table of corrections;
4. Calculate the decay function;
5. Input the scene into EM including the tables of corrections associated with any complex objects;
6. Find the most important propagation paths between the source and receiver;
7. For each path with a large complex object:
 - a. Determine the local source and receiver locations;
 - b. Linearly interpolate (or extrapolate as necessary) the table of corrections;
 - c. Add A_{extra} to the path contribution;
8. Finish EM by incoherently summing all of the path contributions.

Steps 1-4 only need to be performed when a table of corrections does not exist for a required complex object.

EMs strive to very efficiently provide a reasonable approximation of the overall level. Thus, efficiency is also very important in the development of the hybrid method. Efficiency is an important justification for the following:

- Using EMs;
- Utilizing 2.5D BEM on a hard, flat, and infinite ground;
- Having only six input variables (x_s , z_s , x_r , z_r , Δy , and f) in Cartesian coordinates;
- Modeling multiple diffraction very simply;
- Linearly interpolating on a non-uniform grid;
- Precomputing, sorting, and tabulating the BEM results.

While this chapter explains and justifies the hybrid method, it does not present the results of the hybrid method. This chapter argues that the hybrid method works in theory, and Chapter 3 demonstrates that the hybrid method works in practice by considering four cases.

Chapter 3 Validation

§3.1 Objectives

After detailing the hybrid method in the last chapter, this chapter provides examples using the hybrid method compared to EMs and reference methods. This chapter's primary objectives are to demonstrate that the hybrid method

- Is significantly more accurate than EMs for complex cases;
- Has a reasonable computational cost for even large scenes.

To these ends, this chapter presents the following four cases:

1. T-barrier with hard ground (§3.3);
2. T-barrier with soft ground (§3.4);
3. T-barrier with buildings (§3.5);
4. T-barrier in a real situation (§3.6).

The T-barrier with hard ground case demonstrates that the interpolation process is sufficiently accurate despite the data points being very far apart compared to a wavelength. The T-barrier with soft ground case illustrates that the correction for the T-top is fairly invariant with respect to the ground type. The T-barrier with buildings case verifies the process of finding local coordinates with multiple diffractions, reflections, and lateral diffractions. Finally, the T-barrier in a real situation demonstrates that the hybrid method can be used in larger and more complicated situations. This chapter uses the same T-barrier and associated table of correction factors for all of the cases, which is a strong advantage of the hybrid method. Thus, before detailing the cases, the next section describes the table of corrections including the I-barrier and T-barrier, the source and receiver locations, and the frequencies.

§3.2 Table of corrections

This section describes the table of correction factors for the cases developed in this chapter. All of these cases use the same simplified and complex barriers. The simplified barrier is an I-barrier that is infinitely long, 3 m tall, and 20 cm wide. The complex barrier is a T-barrier that has the same dimensions as the I-barrier except it also has a 1 m wide and 20 cm thick T-top. Figure 3.1 gives the cross-section of each barrier. The barrier surfaces are all hard and the ground is infinite, flat, and hard.




Figure 3.1: Barrier cross-sections

The left plot shows the cross-section of the I-barrier, which is the simplified object and is 3 m tall and 20 cm wide. The right plot illustrates the T-barrier, which is the complex object and is 3 m tall, 20 cm thick, and 1 m wide at the top.

The sources are 1-35 m from the barrier and 0-4 m high. The receivers are 1-200 m from the barrier and 0-35 m high. The distance parallel to the barrier between the source and receiver is 0-60 m. The frequencies are the 50-5000 Hz third octave bands. Table 3.1 gives all the positions of the sources and receivers and the frequencies. An optimization procedure that is similar to the one described in Chapter 4 produces the positions by minimizing the variance of a Gaussian process.

Since the points are on a grid, there are $18 \times 14 \times 22 \times 21 \times 26 \approx 3$ million source/receiver points and a total of $3 \text{ million} \times 21 \approx 60$ million data points. Each data point requires about 20 BEM computations (a factor of about 10 for the number of frequencies/third octave band and a factor of 2 for the simplified and complex shapes). Thus, calculating the dataset requires about 1 billion boundary element computations. Nevertheless, since these computations are independent, they can be parallelized and completed in less than 24 hours using Micado¹¹⁷. The computation time may be longer using other BEM approaches because Micado is highly optimized and uses a variational approach that requires fewer points per wavelength than direct BEM approaches. While

this computation time is substantial, the table of corrections can be used for a large number of applications that all use the same complex object. Indeed, all four cases presented in this chapter use the same table of corrections to demonstrate this point.

Table 3.1: Source/receiver locations and frequencies

This table gives all of the source/receiver positions and frequencies for the table of corrections.

n	x_s (m)	z_s (m)	x_r (m)	z_r (m)	Δy (m)	f (Hz)
1	-35.00	0.00	1.00	0.00	0.00	50
2	-31.79	0.13	1.30	1.00	1.00	63
3	-28.69	0.28	1.60	1.24	1.24	80
4	-23.69	0.58	2.30	1.49	1.50	100
5	-19.19	0.91	3.19	2.05	2.03	125
6	-15.62	1.26	4.28	2.68	2.68	160
7	-12.66	1.44	5.68	3.43	3.04	200
8	-10.23	1.64	7.47	4.32	3.45	250
9	-8.19	2.04	9.76	5.35	4.33	315
10	-6.56	2.48	12.54	5.95	5.37	400
11	-5.44	2.72	16.03	6.58	5.96	500
12	-4.50	2.96	20.50	8.00	6.61	630
13	-3.69	3.47	26.17	9.67	8.02	800
14	-2.97	4.00	33.24	11.63	9.70	1000
15	-2.36		42.09	13.94	11.68	1250
16	-1.83		53.34	16.67	13.98	1600
17	-1.37		67.37	19.80	16.69	2000
18	-1.00		85.08	22.83	19.88	2500
19			105.28	26.40	23.60	3150
20			130.55	30.41	27.99	4000
21			161.59	35	33.10	5000
22			200.00		39.14	
23					43.30	
24					48.50	
25					53.98	
26					60.00	

§3.3 T-barrier with hard ground

The first test case is a T-barrier with hard, flat ground. This case is geometrically identical to the case for calculating the table of corrections. Furthermore, 2.5D BEM produces both the table of corrections and the reference results for this situation. Thus, this case helps to quantify two sources of error. First, EMs and BEM have different modeling assumptions that inherently result in slightly different values for even the simplified object (i.e. the I-barrier). This error between the methods establishes a baseline that the hybrid method does not seek to improve but rather to maintain for more complicated geometries. Without fundamentally changing either method, the only way to reduce this error is to optimize the width of the simplified object with respect to this error, which §2.4.3 outlines. Second, interpolating the table of corrections where the points are far apart compared to a wavelength produces interpolation error. This error can be reduced by optimizing the locations of the points in the table of corrections, which Chapter 4 describes in detail. This simple case quantifies these two sources of error.

§3.3.1 Description

The first test case has the following parameters (Figure 3.2):

- The **ground** is hard and is the plane $z = 0$ m;
- The **barrier** is an infinitely long T-barrier along the y -axis with the same cross-section as shown in the right-hand plot of Figure 3.1;
- The **source** is a monopole at $(x, y, z) = (-3, 17.5, 0.3)$ m;
- The **receivers** are on a uniform grid at $x = [-5, -4.75, \dots, 25]$ m, $y = [0, 0.25, \dots, 35]$ m, and $z = 1.5$ m;
- The **frequencies** are 50-5000 Hz. During the analysis, the spectrum is A-weighted and weighted to look like a car (i.e. engine noise plus rolling noise) traveling at 80 km/h according to the spectra provided in the Hosanna Project¹²⁰ (Task 2.3, p 43-44);
- The **medium** is homogeneous air where the speed of sound is $c = 340$ m/s, the density is $\rho = 1.3$ kg/m³, and the reference pressure is $p_{\text{ref}} = 20$ μ Pa;
- The **engineering method** is Harmonoise with up to 6 reflections, 2 lateral diffractions, and 1 km path length. However, lateral diffraction around the barrier is turned off, so in this case the lateral diffraction order is essentially zero. In §3.5 and §3.6, when the scene has other objects, the lateral diffractions are important.




Figure 3.2: Barrier with hard ground

There is a T-barrier with hard ground. The source is 3 m from the barrier and 30 cm high. The receivers are 1.5 m high and cover the entire area with a uniform grid with 25 cm spacing in both the x and y directions.

This case is analyzed with two methods: the hybrid method and 2.5D BEM. The hybrid method uses the Harmonoise EM^{44,45} with up to six reflections, two lateral diffractions, and a maximum path length of 1 km. To model an infinite barrier using the EM, lateral diffractions around the barrier have been turned off. CNOSSOS-EU was not used simply because it was not yet included in the development code that was used to implement the hybrid method. The sound propagation models in Harmonoise and CNOSSOS-EU are very similar, so they produce very similar results. Moreover, the hybrid method is designed to work with any EM, and the conclusion that EMs benefit from the hybrid method should be true regardless of the particular EM that is used.

In this case, 2.5D BEM is being used in two distinct ways. 2.5D BEM provides 1) the table of corrections and 2) reference results to evaluate the hybrid method. The first case has about 3 million source/receiver pairs over a wide range of all of the variables (i.e. x_s , z_s , x_r , z_r , Δy). The second case has about 17 thousand source/receiver pairs with only one source position (x_s and z_s) and one receiver height (z) but a small range for x and y that are densely sampled.

§3.3.2 Results

To evaluate the hybrid method, Figure 3.3 illustrates the total sound pressure levels with the I-barrier and T-barrier using 2.5D BEM and the hybrid method. The top plots give the total levels calculated with 2.5D BEM, and the bottom plots give the same plots calculated with the hybrid method. In other words, the top plots show what the hybrid

method should ideally predict, and the bottom plots show what the hybrid method actually predicts. The left-hand plots illustrate the total levels for the I-barrier, and the right-hand plots do the same for the T-barrier. The hybrid method for the I-barrier gives the results of the underlining EM, which is the Harmonoise method.

Comparing the top-left plot with the bottom-left plot demonstrates how well the Harmonoise method approximates the 2.5D BEM results for the I-barrier. For this case, these two left-hand plots quantify the error between the two methods, which is usually less than 1 dB, and illustrate the best agreement that the hybrid method could reasonably be expected to provide in more complicated cases.

There are two notable differences between the I-barrier plots. First, the decay in level going away from the source is more circular for EM and more oval for 2.5D BEM. Thus, the source appears to be a line segment along the y -direction in the 2.5D BEM calculations and a point in the EM calculations. This difference results from a simplifying assumption of EM that the barrier is always perpendicular to the propagation path. Second, the area around the line $x = 20$ m has higher levels in the 2.5D BEM computation than the EM computation. This dissimilar area results from a large difference between the methods in the 1.6 kHz third-octave band in that area. This area demonstrates that the approximation of the ground effect for single diffraction in the EMs is sometimes insufficient to model the ground effect accurately.

The T-barrier (right-hand) plots show a similar level of agreement as the I-barrier (left-hand) plots. Like for the I-barrier, the hybrid method and 2.5D BEM predict slightly different shapes for the greater than 55 dBA area in the T-barrier plots. These differences have two causes. First, the differences between 2.5D BEM and EM for the I-barrier, which the previous paragraph describes, still apply to the T-barrier case because the hybrid method does not make any attempt to reduce the differences between the methods for an I-barrier. Second, the hybrid method also has some interpolation error because the points in the table of corrections are far apart compared to a wavelength.




Figure 3.3: I-barrier and T-barrier with hard ground

All of the plots show the total pressure level for 50 Hz – 5 kHz where the spectrum is A-weighted and weighted to look like a car spectrum. The top plots are calculated with 2.5D BEM and give what the hybrid method would ideally predict, and the bottom plots are calculated with the hybrid method and give what the hybrid method actually predicts. The left-hand plots give the I-barrier results, and the right-hand plots give the T-barrier results. Overall, the hybrid method produces results that are very similar to the 2.5D BEM results.

To quantify the interpolation error without the error between 2.5D BEM and EM for an I-barrier, Figure 3.4 shows the influence of the T-top (i.e. the total pressure of the T-barrier minus the total pressure of the I-barrier, which is the negative of the insertion loss of the top). Thus, the left-hand plot in Figure 3.4 is the top-right plot in Figure 3.3 minus the top-left plot in Figure 3.3, and the center plot in Figure 3.4 is the bottom-right plot in Figure 3.3 minus the bottom-left plot in Figure 3.3. These two plots overall look fairly similar with the differences mostly coming from interpolation error in the hybrid method. Some differences exist on the source side of the barrier where the extra attenuation is assumed to be zero, but these differences are ignored in the hybrid method because they are usually small.

The right-hand plot in Figure 3.4 is the error of the hybrid method compared to 2.5D BEM, which is the center plot minus the left-hand plot. The error of the hybrid method is fairly small and usually less than 1 dB. The main point of Figure 3.4 is that the error of

the hybrid method (right-hand plot) is small compared to the largest correction (left-hand plot).




Figure 3.4: T-top effect with hard ground

The T-top effect (i.e. the total field with the T-barrier minus the total field with the I-barrier) is calculated with 2.5D BEM in the left plot and with the hybrid method in the center plot. The right plot gives the hybrid plot minus the 2.5D BEM plot. The main differences are due to interpolation error.

§3.3.3 Analysis

After visually inspecting the results in the last section, this section statistically analyzes the total error between the engineering or hybrid method and 2.5D BEM for the I-barrier and T-barrier. In the analysis sections, the error is

$$E = L_{\text{HYD}} - L_{\text{BEM}} \quad (3.1)$$

where L is a vector of the total level (A-weighted and weighted to look like a car traveling at 80 km/h) for receivers where $x \geq 2$ m, and the subscript designates the hybrid or BEM method.

To better understand the error, Figure 3.5 provides several statistical metrics for three different cases. For each case, the red line indicates the median (i.e. the second quartile); the blue box shows the center 50% of the data (i.e. the first and third quartiles); the whiskers represent the full range of the data; μ is the mean; σ is the standard deviation; and RMS is the root mean squared error.

The left box calculates the error of EM for the I-barrier compared to 2.5D BEM (i.e. in Figure 3.3 the bottom-left plot minus the top-left plot). This case establishes the baseline performance of EM for the I-barrier, which the hybrid method ideally would preserve for more complicated geometries. The center box calculates the error of EM for the T-barrier compared to 2.5D BEM (i.e. in Figure 3.3 the bottom-left plot minus the top-right plot). This case results from naively using the I-barrier in EM in place of the T-barrier; in other words, this case assumes that the effect of the T-top is negligible. The large increase in

error indicates that this assumption is a poor one. The right box calculates the error of the hybrid method for the T-barrier compared to 2.5D BEM (i.e. in Figure 3.3 the bottom-right plot minus the top-right plot). This case demonstrates how well the hybrid method does for the T-barrier.

Figure 3.5 demonstrates that on the one hand EM approximates the 2.5D BEM results well for the I-barrier, especially considering the computational savings. In this case, the RMS error is approximately 0.6 dBA. On the other hand, EMs model complex geometries poorly; for the T-barrier, the RMS error is about 2 dBA on average and up to 5 dBA. However, the hybrid method substantially improves the results for complex cases; for the T-barrier, the RMS error is brought back to the level of EM for the I-barrier.




Figure 3.5: Error comparison with hard ground

The horizontal lines for each box and whisker plot from bottom to top are the smallest value, first quartile, second quartile, third quartile, and largest value, and the values below each box are the mean (μ), standard deviation (σ), and root mean squared error (RMS). The left box represents the error between EM and 2.5D BEM for the I-barrier. The center box represents the error between EM and 2.5D BEM for the T-barrier where EM is modeling the I-barrier. The right box represents the error between the hybrid method and 2.5D BEM for the T-barrier.

§3.4 T-barrier with soft ground

The next case is the same T-barrier as the last case with soft ground instead of hard ground. The primary purpose of this case is to determine if the same correction factors, which are calculated with hard ground, can be used regardless of the ground type. The motivation for wanting the correction factors to be independent of the ground type is twofold: model simplicity and reducing the number of 2.5D BEM computations.

§3.4.1 Description

All of the parameters are the same as the last case except the entire ground has a flow resistance of $200 \text{ kPa} \cdot \text{s}/\text{m}^2$, which is similar to grass-covered ground. 2.5D BEM approximated the impedance of the ground using the Delany-Bazley model¹²⁹. Figure 3.6 shows the geometry where the green area represents soft (grassy) ground.




Figure 3.6: Barrier with soft ground

There is a T-barrier with soft ground. The source is 3 m from the barrier and 30 cm high. The receivers are 1.5 m high and cover the entire area with a uniform grid with 25 cm spacing in both the x and y directions.

§3.4.2 Results

The results for the soft ground case are very similar to those in the hard ground case. Figure 3.7 shows many of the same patterns as the previous case. For example, the spatial spreading of the source in the y -direction for the 2.5D BEM plot remains. The largest difference between the hard ground case and the soft ground case is a 2-3 dB drop in overall level associated with mostly losing the ground reflection. In addition, the ground affect at $x = 20 \text{ m}$ in Figure 3.3 is mostly absent from Figure 3.7 because the ground affect is much weaker for soft ground than for hard ground.




Figure 3.7: I-barrier and T-barrier with soft ground

All of the plots show the total pressure level for 50 Hz – 5 kHz where the spectrum is A-weighted and weighted to look like a car spectrum. The top plots are calculated with 2.5D BEM and give what the hybrid method would ideally predict, and the bottom plots are calculated with the hybrid method and give what the hybrid method actually predicts. The left-hand plots give the I-barrier results, and the right-hand plots give the T-barrier results. Overall, the hybrid method produces results that are very similar to the 2.5D BEM results.

Again, to evaluate the interpolation error and the error from using a table of correction calculated with hard ground in a soft ground case, Figure 3.8 plots the T-top correction (i.e. the total field with the T-barrier minus the total field with the I-barrier) calculated with 2.5D BEM (left plot) and the hybrid method (center plot) and the difference (i.e. hybrid results minus 2.5D BEM results) in the right plot. Similar to the hard ground case, for soft ground the average error in the T-top correction is small compared to the average T-top correction. This result suggests that the table of corrections, which is only calculated with hard ground results, can be applied even when the ground is soft.

There are two main reasons why the T-top correction appears mostly independent of the ground type. First, the ground effect is already disrupted by the I-barrier, and the path length differences for the T-barrier are not very different from the I-barrier. If instead of using an I-barrier the reference case was flat ground, then the correction would be quite different for hard and soft grounds. Second, since the source is fairly broadband, the

errors in individual third-octaves, which can be significant, are averaged away when summed over the whole frequency range because the direction of the error is random. If a narrowband source is modeled instead of the typical broadband transportation noise, then there might be larger errors depending on the case.

Finally, the hybrid method predicts slightly different values for the hard ground case (center plot of Figure 3.4) and the soft ground case (center plot of Figure 3.8) because these results are summed across frequencies. For each third-octave, the results would be identical, but these results are affected by the relative level of each third-octave. The soft ground does not affect all frequencies the same amount.




Figure 3.8: T-top effect with soft ground

The T-top effect (i.e. the total field with the T-barrier minus the total field with the I-barrier) is calculated with 2.5D BEM in the left plot and with the hybrid method in the center plot. The right plot gives the hybrid plot minus the 2.5D BEM plot. The main differences are due to interpolation error and using a table of corrections calculated with hard ground instead of soft ground.

§3.4.3 Analysis

As with the hard ground case, the next step is to statistically analyze the results. The process is the same as before with different data. Figure 3.9 shows that error between EM and 2.5D BEM for the I-barrier is small on average with a small spread (RMS = 0.4 dBA). Then, the error greatly increases when the EM is applied to the T-barrier (RMS = 1.7 dBA), but is substantially reduced by the hybrid method (RMS = 0.6 dBA). Thus, much improvement is still achieved in soft ground cases when using the table of corrections that was calculated using hard ground.

Comparing the hard ground (Figure 3.5) and soft ground (Figure 3.9) results reveals that the hybrid method has similar RMS error for both the hard and soft ground cases. One would expect that the soft ground case has slightly larger error because the correction factors are calculated using a hard ground. This expectation should hold true in general when EM performs equally well for both hard and soft ground with the I-barrier.

However, in this case EM has lower RMS error for soft ground (0.4 dBA) than for hard ground (0.63 dBA) because of the error near $x = 20$ m between the left-hand plots in Figure 3.3 that is mitigated in the left-hand plots in Figure 3.7.




Figure 3.9: Error comparison with soft ground

The horizontal lines for each box and whisker plot from bottom to top are the smallest value, first quartile, second quartile, third quartile, and largest value, and the values below each box are the mean (μ), standard deviation (σ), and root mean squared error (RMS). The left box represents the error between EM and 2.5D BEM for the I-barrier. The center box represents the error between EM and 2.5D BEM for the T-barrier where EM is modeling the I-barrier. The right box represents the error between the hybrid method and 2.5D BEM for the T-barrier.

§3.5 T-barrier with buildings

The next case is a T-barrier with buildings. This case is a big step up in the level of complexity of the test case. Both of the previous cases were 2.5D problems. This case is the first truly 3D problem. This aspect makes the case much more realistic and interesting. In addition, in this case the propagation paths could interact with something besides just the complex object. The primary purpose of this case is to test how well the reflection, lateral diffraction, and multiple diffraction approximations work.

§3.5.1 Description

This problem is an augmentation of the first test case, so many of the parameters are the same. Specifically, the ground, barrier, source, receivers, medium, and EM are all the same. The main difference is the addition of the buildings, which are shown in Figure 3.10. Table 3.2 also gives the coordinates of the building corners. All of the building surfaces are hard, and lateral diffractions are allowed around the buildings.




Figure 3.10: Barrier with buildings

There is a 35 m long T-barrier with two buildings 10 m from the barrier. The source is 3 m from the barrier and 30 cm high. The receivers are 1.5 m high and cover the entire area with a uniform grid with 25 cm spacing in both the x and y directions.

Table 3.2: Building geometries for test case

The two buildings are box-shaped and have the coordinates given below.

	Building 1		Building 2	
Corner	x (m)	y (m)	x (m)	y (m)
1	10	5	15	20
2	20	5	20	25
3	20	15	15	30
4	10	15	20	25
z (m)	6		8	

Another major change is that the reference method is switching from 2.5D BEM to fast-multiple BEM (FM-BEM), but the correction factors are still the ones from §3.2 and are still calculated with 2.5D BEM. Since the buildings have finite length, the 2.5D BEM, which provides the reference calculations for the last two examples, is not appropriate for this case. Instead, FM-BEM is used, and Christophe Langrenne from LMSSC/CNAM¹³⁰ performed the computations. FM-BEM efficiently approximates 3D BEM by reducing the number of interactions between the elements by grouping nearby elements and solving the matrix equation iteratively so that only part of the matrix must be stored. This implementation of FM-BEM was verified using 3D Micado up to 400 Hz for the same geometry. FM-BEM provides the results for this case because it 1) is a reference method, 2) can model this case at most of the important frequencies, and 3) is readily available.

Lastly, since this case is much more computationally intensive than the last two cases, the frequency range is only 50-1600 Hz. The upper frequency limit is bounded by the computational ability of the FM-BEM method. For comparison, regular 3D BEM is limited to less than about 400 Hz, so using FM-BEM instead of ordinary BEM increased the maximum frequency by a factor of four.

§3.5.2 Results

Again, this case is evaluated using the same plots as the last two cases, but the data is different. Comparing the I-barrier plots in Figure 3.11 shows good overall agreement between Harmonoise and FM-BEM, but Harmonoise tends to predict slightly lower levels. The differences are largest in the shadow regions of the buildings where EMs are known to be less accurate. The T-barrier plots illustrate a similar level of agreement as the I-barriers.




Figure 3.11: I-barrier and T-barrier with buildings

All of the plots show the total pressure level for 50 Hz – 1.6 kHz where the spectrum is A-weighted and weighted to look like a car spectrum. The top plots are calculated with FM-BEM and give what the hybrid method would ideally predict, and the bottom plots are calculated with the hybrid method and give what the hybrid method actually predicts. The left-hand plots give the I-barrier results, and the right-hand plots give the T-barrier results. Overall, the hybrid method produces results that are very similar to the FM-BEM results.

As in the previous cases, Figure 3.12 plots the T-top correction to attempt to remove the inherent disagreement between Harmonoise and FM-BEM. In the previous two cases, the error was mostly due to interpolation error, but in this case the error is a combination of the error from 1) interpolation, 2) the reflection approximation, 3) the lateral diffraction approximation, and 4) multiple diffraction approximation. Since the agreement is fairly good in between the barrier and the buildings, the interpolation error and reflection error appear small. The lateral and multiple diffraction errors affect the same regions, so evaluating them independently in this case is difficult. The largest error being behind the shorter building tentatively suggests that the multiple diffraction approximation is the largest source of error of the four options. However, this largest error is in the same location as where EMs are the least accurate, so decreasing this error is not a priority.




Figure 3.12: T-top effect with buildings

The T-top effect (i.e. the total field with the T-barrier minus the total field with the I-barrier) is calculated with FM-BEM in the left plot and with the hybrid method in the center plot. The right plot gives the hybrid plot minus the FM-BEM plot. The main differences are due to interpolation error and the approximations for reflections and diffractions.

§3.5.3 Analysis

Once more, Figure 3.13 gives the error for modeling the I-barrier with EM, the T-barrier with EM, and the T-barrier hybrid method compared to FM-BEM. However, the receivers are restricted to only in between the barrier and the buildings (i.e. $2 \text{ m} \leq x \leq 8 \text{ m}$ and $2 \text{ m} \leq y \leq 33 \text{ m}$) where EM should be the most accurate. Behind the buildings, the mean error of EM for the I-barrier is -3.2 dB , so correction in this area is not as important. In front of the buildings, EM does a reasonable job approximating the FM-BEM results, but has noticeably more error this time; for the first time the zero error line is not within the center 50% of the data (i.e. the left blue box in Figure 3.13 does not overlap with zero error line).

As expected, modeling the T-barrier as an I-barrier in EM produces fairly poor results (center box of Figure 3.13); the RMS error is 2.8 dBA , the mean error is 2.5 dBA , and the range is also very large at almost 7 dBA . However, the hybrid method reduces the RMS

error to 0.7 dBA and the mean error to 0.0 dBA; the range is also reduced to 5 dBA with over 50% of the data within 0.5 dBA.

While having a mean error of 0.0 dBA appears good, since EM has a mean error of -0.7 dBA for the I-barrier, it actually indicates that the correction does not quite restore the baseline established by EM. In other words, if EM had been more accurate for the simplified case, the mean error would be larger for the hybrid method (i.e. if the mean error of EM was 0.0 dBA for the I-barrier, then the mean error of the hybrid method would be 0.7 dBA for the T-barrier). Still, the hybrid method performs much better than the plain EM for the T-barrier.




Figure 3.13: Error comparison with buildings

The horizontal lines for each box and whisker plot from bottom to top are the smallest value, first quartile, second quartile, third quartile, and largest value, and the values below each box are the mean (μ), standard deviation (σ), and root mean squared error (RMS). The left box represents the error between EM and FM-BEM for the I-barrier. The center box represents the error between EM and FM-BEM for the T-barrier where EM is modeling the I-barrier. The right box represents the error between the hybrid method and FM-BEM for the T-barrier.

§3.6 Real Scene

The final case is a real scene near Grenoble, France. It has two line sources, a T-barrier, a few buildings, and a partially soft ground. The primary purpose of this example is to demonstrate that the hybrid method can be applied to cases that are larger and more complex than the simple test cases. However, because of the size and complexity, a reference method is unavailable for this case to check the accuracy. This kind of problem is one where the hybrid method provides a new choice in the trade-off between accuracy and computation time. The hybrid method is more accurate than EMs alone and computationally possible in contrast to the reference methods.

§3.6.1 Description

The scene is a real tramway next to a few buildings. Figure 3.14 is a screenshot from Google Maps of the area, and the GPS coordinates (i.e. the latitude and longitude in decimal notation) are (45.194798, 5.755008). The geometry was simplified and approximated to the nearest meter using the measure distance feature in Google Maps. The simplified geometry is the top-left plot of Figure 3.15.



Figure 3.14: Barrier with buildings and soft ground

This picture is from Google Maps and represents the real scene that this section models. There are four buildings with grassy areas in between, a parking lot, and a tramway. The GPS coordinates are (45.194798, 5.755008).

Here are the important parameters:

- The **ground** is soft with a flow resistance of $200 \text{ kPa} \cdot \text{s/m}^2$ in the grassy areas around the buildings [i.e. the rectangle from (32,13) m to (180,80) m] and is hard otherwise for the parking lot and tramway areas;
- There are four **buildings**, which have hard surfaces and are three stories (~ 12 m) tall. Table 3.3 gives the corner positions;
- The **source** is a tram that runs in both directions, so there are two line sources going from (0,7) m to (180,7) m and from (0,10) m to (180,10) m. Both line sources are 30 cm high. While the real scene has a tramway, the source is actually modeled as a car traveling at 80 km/h with the spectrum (50-5000 Hz) from Hosanna that was used previously. Using a car instead of a tram allows the same table of corrections to be used in this case as all of the previous cases because modeling a tram would require the tram's geometry to be included in the 2.5D BEM model to properly account for multiple reflections between barrier and the tram body;
- The **receivers** are on a uniform grid where $x = [0,2, \dots 180]$ m, $y = [0,2, \dots 80]$ m, and $z = 1.5$ m;

- The same **T-barrier** as in the previous cases (Figure 3.1) is added to the real scene and is parallel to the tramway going from (0,11) m to (180,11). The T-barrier is 3 m tall and 1 m from the closest tram;
- The **medium** is homogeneous air where the sound speed is 340 m/s, the density is 1.3 kg/m^3 , and the reference pressure is $p_{\text{ref}} = 20 \text{ } \mu\text{Pa}$;
- The **engineering method** is Harmonoise with up to 3 reflections, 2 lateral diffractions, and 1 km path length. Lateral diffraction around the barrier is turned off.

Table 3.3: Building geometries for real case

Each of the building corners are given below. A diagram of the building layout is also in Figure 3.15.

Corner	Building 1		Building 2		Building 3		Building 4	
	x (m)	y (m)	x (m)	y (m)	x (m)	y (m)	x (m)	y (m)
1	17	21	87	21	135	21	6	47
2	69	21	97	21	145	21	32	47
3	63	37	104	29	152	29	32	70
4	17	37	86	65	134	65	15	70
5			75	63	123	63	15	72
6			69	56	117	56	6	72
z (m)	12		12		12		12	

§3.6.2 Results

The results for this case are slightly different because there is no reference solution. Figure 3.15 gives the geometry (top-left), the total level for the I-barrier (top-right), the total level for the T-barrier (bottom-left), and the T-top correction (bottom-right). The hybrid method produced all of these levels. The main point here is that the hybrid method produces plausible results for the I-barrier and the T-barrier and that the absolute value of the T-top effect is significant (i.e. sometimes greater than 2 dBA) and should be included.

To give some idea of the computation times, there are 3731 receivers, which took 79 min 9 s for the I-barrier and 79 min 35 s for the T-barrier on a standard desktop computer (i.e. 3.2 GHz quad-core processor with 16 GB of RAM) or 1.3 s per receiver on average. Running a shorter computation multiple times reveals that there is about 0.2% uncertainty in the computation time or about 10 seconds for this computation. The additional computation time is almost entirely dedicated to loading the table of corrections, which took 29 s. Thus, the additional time for the hybrid method not including the time to load the file is -3 ± 20 s, which is negligible compared to the total computation time of EM (i.e. 4749 s) and even negligible compared to the uncertainty in the computation time.




Figure 3.15: I-barrier and T-barrier with buildings and soft ground

All of the plots show the total pressure level for 50 Hz – 5 kHz where the spectrum is A-weighted and weighted to look like a car spectrum. The top-left plot gives the simplified geometry that was input into the hybrid method. The top-right plot shows the total level with the I-barrier, and the bottom-left plot displays the total level with the T-barrier. The bottom-right plot illustrates the T-top correction, which is the T-barrier plot minus the I-barrier plot. In this case, the T-top effect demonstrates a substantial difference between the barriers that would be important to include in a model.

§3.7 Summary

This chapter considers four cases that each try to quantify different sources of potential error. First, the T-barrier with hard ground case demonstrates that EM adequately models the simple I-barrier but not the more complicated T-barrier and that the interpolation error is small compared to the typical values of the extra attenuation. Second, the T-barrier with soft ground case indicates that for broadband sources the table of corrections does not need to be recalculated for each ground type but rather that a table of corrections calculated with hard ground reasonably approximates even soft grounds when summed over all the frequencies. Third, the T-barrier with buildings case shows that the reflection, lateral diffraction, and multiple diffraction approximations all produce plausible results, especially considering relatively large error between EM and FM-BEM for this case. Fourth, the real case demonstrates that the hybrid method is computationally feasible even for larger and more complicated scenes and produces plausible result that should be included.

Finally, all of the cases confirm that the hybrid method performs about 2 dBA better on average than the plain EM for the T-barrier, and at select source receiver pairs the improvement is as much as 5 dBA. The actual improvement depends on the case (i.e. the complex object, the source/receiver positions, and the surrounding geometry), but these results demonstrate that the hybrid approach is substantially more accurate than the regular EM by the criterion set forth in CNOSSOS-EU. The CNOSSOS-EU standard states, “A parameter is considered essential if the range of values the parameter can take yields variations in L_{den} or L_{night} of more than ± 2.0 dB(A) 95% C.I. (all other parameters remaining unchanged).”⁵ Thus, using this criterion the hybrid method is an essential addition to EMs.

The next chapter discusses how to maximize the improvement that the hybrid method offers EMs by minimizing the interpolation error. The optimization turns out to be more complicated than it at first appears because the locations of the values, not the values themselves, must be optimized.

Chapter 4 Optimization

§4.1 Objectives

The previous chapter validates the hybrid method with four examples, but the results depend on the data points in the table of corrections. If the data point locations are selected poorly, then the interpolation error will be larger than necessary. This chapter's primary objective is to determine optimized source and receiver locations for the table of corrections to reduce the interpolation error. To accomplish this task,

- §4.2 introduces data location optimization;
- §4.3 discusses potential point distribution options;
- §4.4 presents Gaussian processes;
- §4.5 optimizes the hyper-parameters of a Gaussian process;
- §4.6 minimizes the maximum variance of the interpolated values;
- §4.7 selects the number of data points for each dimension;
- §4.8 evaluates the benefit of the optimization procedure.

§4.2 Introduction

Optimization¹³¹ is a diverse and developed field with applications in many disciplines including acoustics⁵⁹. Before discussing the specific optimization problem of minimizing interpolation error, this section

- introduces optimization more generally;
- explores objective functions
- discusses numerical optimization methods;
- explains convexity.

§4.2.1 General optimization

Optimization problems generally have

- Variables, \mathbf{x} ;
- Objective function, $f(\mathbf{x})$;
- Equality constraints, $\mathbf{g}(\mathbf{x})$;
- Inequality constraints, $\mathbf{h}(\mathbf{x})$.

The variables are values that can be changed to improve the objective function. The objective function yields a value that must be minimized or maximized and depends on the variables. Often, the variables cannot take on any value and are constrained by certain relationships. These relationships are called constraints, and there are two kinds of constraints: Equality and inequality. Equality constraints are functions that equal a constant and inequality constraints are functions that are bounded by a constant. An optimization problem is usually written as follows:

$$\text{minimize } f(\mathbf{x}) \quad \text{such that} \quad \begin{array}{l} \mathbf{g}(\mathbf{x}) = 0 \\ \mathbf{h}(\mathbf{x}) \leq 0 \end{array} \quad (4.1)$$

An optimization problem may not come in this form, but all of them can be converted to this form. For example, if the problem was originally a maximization problem, then multiply the objective function by negative one; if the constraints are not less than or equal to zero, then subtract a constant and/or multiply by negative one to reach the desired form.

§4.2.2 Objective functions

In real world applications, defining the objective function is often difficult. Ideally, the objective function should be computationally fast and accurately represent the quantity that should be optimized. Consider a simple example of fitting a curve to a dataset. The true optimization problem is to minimize the error between the curve and unknown future data point. Since the actual objective function cannot be evaluated, then the objective function is approximated by the total squared error between the curve and the data. This definition reduces computation time but decreases the accuracy.

The difference between these two problems emerges when the number of parameters is increased until it equals the number of data points. Now, the curve goes through every data point so that the objective function is as small as possible, but this curve probably does not predict a new data point very well, which is called overfitting. The objective function needs another term that quantifies how complex the model is. All else being equal, the model complexity should be as small as possible.

Another issue with fitting a curve to a dataset involves the definition of error. Commonly, the error is the total squared error, but there are many other options that may be better in certain situations. Instead, the objective function could be the sum of the absolute value of the error, the maximum error, or the median error. The total squared error is commonly used because it has an analytical solution, but all of the others are valid options that may produce different results.

An example that is closely related to the real application of this dissertation is weather forecasters deciding where to put new weather stations. The objective is to predict the weather as accurately as possible for the largest number of people for the smallest number of weather stations. Since the forecasters cannot build many weather stations before deciding where to put them, the objective function again cannot be measured directly. Instead, the forecasters choose locations where many people live and the uncertainty in their predictions is high. The uncertainty depends on the distance to the nearest weather station (i.e. a larger distance creates larger uncertainty) and the terrain because the weather in the mountains may change more rapidly than in the plains. Thus, the objective function could be the maximum of some combination of the uncertainty in an area and its population.

§4.2.3 Numerical optimization methods

Once the optimization problem is setup, then one of many established numerical optimization algorithms can be applied to find a solution. These algorithms often use the first and/or second derivatives of the objective function to improve the convergence rate. When there is more than one variable, then the first derivatives are stored in a vector called the gradient and the second derivatives are stored in a matrix called the Hessian. They have the following forms:

$$\nabla_i(f) = \frac{\partial f}{\partial x_i} ; \quad \mathbf{H}_{i,j}(f) = \frac{\partial^2 f}{\partial x_i \partial x_j} \quad (4.2)$$

Entire books have been written about numerical optimization algorithms¹³¹, but some common optimization approaches include:

- Steepest decent
- Newton
- Quasi-Newton

The steepest (or gradient) decent method starts at an initial guess and takes small steps in the direction of the negative gradient until it converges to a solution. The benefit of this method is that it only requires the first derivatives, which can be calculated analytically

or numerically using finite differences, but it converges slowly compared to the other methods.

The Newton method is similar to the steepest decent method except it uses both the gradient and Hessian (i.e. the second order derivatives) to obtain a much faster convergence rate. The main drawback is that the Hessian grows quadratically with the number of variables, so it can be expensive to compute when there are many variables.

The quasi-Newton methods use the gradient and an approximation of the Hessian to choose a direction. Since the actual Hessian is not computed, the quasi-Newton methods take slightly more and slightly quicker steps than the Newton method, especially as the number of variables increases.

Since all of these methods are well-established in the literature, the details of their implementations are not discussed in detail. Indeed, these methods do not usually need to be implemented because many computer programming languages with an advanced mathematics library already have implementations of these methods. For this project, the optimization is performed using the Matlab function `fminunc()`¹³², which is in the optimization toolbox.

§4.2.4 Convexity

Convexity is an important concept in optimization. Informally, a function is convex if any two points in the function's domain can be connected by a straight line. Formally, a function $f(\mathbf{x})$ is convex if for any two points in the function's domain (\mathbf{x} and \mathbf{x}') and for all $p \in [0,1]$

$$f[p\mathbf{x} + (1 - p)\mathbf{x}'] \leq pf(\mathbf{x}) + (1 - p)f(\mathbf{x}') \quad (4.3)$$

For example, $f(x) = x^2$ is convex, but $f(x) = \sin x$ is not convex. Convexity is important because if an optimization problem is convex, then a local minimum is also a global minimum. If an optimization problem is not convex, then a local minimum does not have to be a global minimum. In other words, for non-convex problems, an optimization algorithm may find a local minimum that is much larger than the global minimum (i.e. a poor local minimum). For non-convex problems, optimization algorithms are often run multiple times from random initial positions to improve the probability of finding the global minimum.

§4.3 Point distribution options

There are many different ways to distribute points, and this section discusses only a small fraction of the ways. Specifically, it highlights linear and exponential distributions, and then discusses how to optimize more complicated models.

§4.3.1 Linear distribution

A very common option is to linearly distribute the points. For example, in a 1D problem the x -positions could be

$$x_n = c_1 n + c_0 \quad (4.4)$$

where c_1 and c_0 are constants, $n \in \{0, 1, 2, \dots, N-1\}$, and N is the total number of points. Given x_0 and x_{N-1} (i.e. the position of the first and last points), then all the points are given by the following equation:

$$x_n = \frac{x_{N-1} - x_0}{N-1} n + x_0. \quad (4.5)$$

The linear model is very easy to compute and works well when the function of x is equally important or complicated in the entire range. However, for the current application the complexity of the function is expected to be highest near the complex object or ground and decrease moving away from the complex object. These points are also expected to be used more frequently, so are more important than the points far away.

§4.3.2 Exponential distribution

To concentrate the points while keeping the simplicity of the linear model, one can use the following exponential model:

$$x_n = e^{c_1 n} + c_0. \quad (4.6)$$

Since there are only two variables, the constants are also determined by the range and number of points to be

$$x_n = e^{\ln(x_{N-1} - x_0 + 1) \frac{n}{N-1}} + x_0 - 1. \quad (4.7)$$

These two distributions give a sample of very common distributions with only two constants. Increasing the number of constants makes the distribution much more customizable and potentially difficult to use. There are again many options, but one

option that allows the concentration of points near the barrier to be customized is to add a coefficient to the exponential term:

$$\mathbf{x}_n = c_2^2 e^{c_1 n} + c_0. \tag{4.8}$$

The constant c_2 is squared to show that it must be a non-negative real number. Similar to the last two models, two of the constants can be eliminated using the range and number of points to give:

$$\mathbf{x}_n = c_2^2 e^{\ln\left(1 + \frac{x_{N-1} - x_0}{c_2^2}\right) \frac{n}{N-1}} + x_0 - c_2^2. \tag{4.9}$$

However, since the equation is underdetermined, c_2 could be any real number. If $c_2^2 = 1$, then this case reduces to the exponential model. If c_2^2 is sufficiently large, then this case reduces to the linear model using the Taylor series approximations $\ln(1 + z) \approx z$ and $e^z \approx 1 + z$, which are valid when $|z| \ll 1$. Finally, if $0 < c_2^2 < 1$, then the points are even more concentrated towards x_0 than for the exponential model with $c_2 = 1$. Figure 4.1 illustrates all these options. Thus, c_2^2 allows the points to go from super concentrated near x_0 to being uniformly spaced between x_0 and x_{N-1} .




Figure 4.1: Point distribution options

This figure illustrates several different options that can all be created using the optimized exponential model in Equation (4.9).

§4.3.3 Optimizing the distribution

The next step is to optimize c_2 , but doing this task requires a definition of the objective function and probably some data to help evaluate the objective function. This particular example using Equation (4.9) only has one variable, but in general a different function could have been selected or even the position of each point could be a variable. However,

the next step is an optimization problem regardless of the specific form; only the number of variables would change.

Since the end goal is to interpolate these data points, the most natural definition of the objective function is the sum of the squared interpolation error. However, this definition makes optimizing the set of variables very expensive. To understand why, consider what would have to be done in order to do the optimization. First, a sample dataset would have to be calculated using a reference method (e.g. 2.5D BEM) or measured. Depending on the size of the sample dataset, the computation time would be significant, and more flexible models with more variables often require even more data. Furthermore, each evaluation of the objective function would require a full table of corrections to be calculated, which requires on the order of one day to complete. An optimization problem usually requires many objective function evaluations, so this optimization could easily require months of computation time. There must be a better way.

The primary problem is the large number of BEM computations that must be performed for each evaluation of the objective function. To avoid this computational expense, the objective function is the maximum variance of an interpolated value. Variance is the standard deviation squared and represents the level of uncertainty in an interpolated value. This objective function is reasonable because as the variance approaches zero, the interpolation error should also approach zero. However, the current interpolation method (i.e. linear splines) does not estimate the uncertainty, so another interpolation method that does this must be found.

To search for an interpolation method that estimates the uncertainty of the values, here is the list given in §2.6.2 of common interpolation methods:

- Inverse distance weighted;
- Natural neighbor;
- Regression;
- Spline.

Inverse distance weighted, natural neighbor, and spline methods do not model the uncertainty of a predicted value, so consider the regression methods (e.g. linear, polynomial, and Gaussian process). However, most of them would still require too many BEM calculations to estimate the variance because they require the values at the locations instead of just the locations to predict the variance. The exception is Gaussian process regression, which assumes the uncertainty is Gaussian so that it does not require the values at the locations to predict the variance. The rest of this chapter explains how to minimize the variance using a Gaussian process.

§4.4 Gaussian processes

Gaussian processes¹²⁶ are statistical models for modeling continuous variables. Given a set of input variables and an output variable, a Gaussian process estimates the output by assuming that points that are close together in the input space will have a high covariance in the output space (i.e. points that have similar inputs have similar outputs). As the points get further apart, the covariance deteriorates, which leads to greater uncertainty about an estimated output value. Thus, the covariance allows a Gaussian process to also estimate the uncertainty of an output value. To understand Gaussian processes better, this section

- Introduces Gaussian processes with an example
- Describes some possible covariance functions
- Provides the equations to calculate a Gaussian process

§4.4.1 Introduction

Before getting into the details, Figure 4.2 illustrates a Gaussian process for a small dataset (blue dots) that is taken from a sine wave (red curve). The blue curve shows what a Gaussian process predicts for this data and the shaded region illustrates the standard deviation. Most interpolation methods only give the blue curve, but a Gaussian process also gives the shaded region. Furthermore the width of the shaded region is independent of the values of the data. Figure 4.2 demonstrates the following important characteristics of Gaussian processes:

- Gaussian processes can be used for both interpolation and extrapolation;
- The standard deviation increases as the distance to the closest data point increases;
- There is a maximum standard deviation, which is an input parameter;
- Far from any data point, this Gaussian process predicts 0 dB, which is also an input parameter.




Figure 4.2: An example Gaussian process

This figure gives an example Gaussian process. The red curve is a sine wave and represents the reference solution. The blue dots are points sampled from the sine wave. The blue curve is what the Gaussian process predicts, and the shaded area shows the standard deviation that the Gaussian process predicts.

§4.4.2 Covariance function

The covariance function quantifies the similarity between two points. In the literature, the most common covariance function is the squared exponential covariance function:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_y^2 \exp \left[\frac{-\|\mathbf{x} - \mathbf{x}'\|_2^2}{l^2} \right] \quad (4.10)$$

where $\|\mathbf{x}\|_2$ is the Euclidean norm (i.e. the 2-norm or the k -dimensional length of \mathbf{x}). In addition, σ_y^2 and l are the amplitude and characteristic length scale of the covariance function. They are considered hyper-parameters, which are usually found by maximizing the log marginal likelihood. The amplitude is written as σ_y^2 to indicate that it is a variance term (i.e. a squared standard deviation) and represents a maximum variance.

This covariance function works well because it is largest when the two input points are the same and smallest when they are far apart. In addition, this function is infinitely differentiable and has some flexibility due to the hyper-parameters. However, for this application the covariance function actually needs to be modified. Two points far from the complex object are likely to be more similar to each other than two points that are the same distance apart but closer to the complex object, so the characteristic length l must be a function of \mathbf{x} and \mathbf{x}' . Furthermore, the characteristic length is likely different for each direction. Keeping $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}', \mathbf{x})$ and incorporating these changes produces the following covariance function:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_y^2 \exp \left[\sum_i \frac{-(x_i - x'_i)^2}{(\mathbf{m}_i^2 x_i + \mathbf{b}_i^2)(\mathbf{m}_i^2 x'_i + \mathbf{b}_i^2)} \right] \quad (4.11)$$

where \mathbf{m} and \mathbf{b} are hyper-parameters and every component of \mathbf{x} and \mathbf{x}' are assumed to be non-negative so that the denominator is always positive. In the equation, \mathbf{m} and \mathbf{b} are squared to show that they are always non-negative and to make the optimization more efficient (i.e. fewer constraints). The exponent is a weighted distance function where the characteristic length can now vary linearly and in each direction rather than only be a constant. Depending on the situation, one might choose a different function for the characteristic length, but the linear approximation is the simplest expression that will allow the distribution of points to be non-uniform. In the 1D case, the equation simplifies slightly because there are not any vectors:

$$k(x, x') = \sigma_y^2 \exp \left[\frac{-(x - x')^2}{(m^2 x + b^2)(m^2 x' + b^2)} \right]. \quad (4.12)$$

To better understand this equation, Figure 4.3 illustrates this equation using four different combinations of m and b . When $m = 0$ as in the left-hand plots, Equation (4.12) reduces to Equation (4.10) by substituting $l = b^2$. The left-hand plots demonstrate that the characteristic length l affects the width of the diagonal of high covariance but that its width is independent of x or x' . If $m \neq 0$ as in the right-hand plots, then the width of the high covariance diagonal is a function of x and x' . Thus, the covariance function in Equation (4.12) is more flexible than in Equation (4.10).

A small width of the main diagonal means that the data points must be close together to provide a good interpolation, but a large width indicates that the data points can still produce a good interpolation even when the data points are further apart. In practice, Equation (4.10) would produce equally spaced points whereas Equation (4.12) allows data points to concentrate close to the complex object where the correction factor is likely to be most sensitive to the source/receiver positions.




Figure 4.3: The covariance function for various characteristic lengths

These plots show the covariance in Equation (4.12) for $\sigma_y = 1$ and various values of m and b . The top plots use $b^2 = 0.3$ m, and the bottom plots use $b^2 = 3$ m. The left-hand plots use $m^2 = 0$ (dimensionless), and the right-hand plots use $m^2 = 0.3$.

§4.4.3 Calculations

Since Gaussian processes are established in the literature¹²⁶, this section presents the results without derivation. Calculating a Gaussian process requires the following input parameters:

- \mathbf{X}_{data} , a column vector where each row contains the location of a data point;
- \mathbf{y}_{data} , a column of the values at the locations \mathbf{X}_{data} ;
- σ_n^2 , the variance of the noise in the \mathbf{y}_{data}
- $k(\mathbf{x}, \mathbf{x}')$, a covariance function;
- $\mathbf{X}_{\text{interp}}$, a column vector where each row contains the location of an interpolation point.

Given all of the inputs, the first step is to calculate the following covariance matrices:

$$\mathbf{K}_{\text{data},i,j} = k(\mathbf{X}_{\text{data},i}, \mathbf{X}_{\text{data},j}) + \sigma_n^2 \delta_{i,j}; \quad (4.13)$$

$$\mathbf{K}_{\text{mixed},i,j} = k(\mathbf{X}_{\text{data},i}, \mathbf{X}_{\text{interp},j}); \quad (4.14)$$

$$\mathbf{K}_{\text{interp},i,j} = k(\mathbf{X}_{\text{interp},i}, \mathbf{X}_{\text{interp},j}); \quad (4.15)$$

where $\delta_{i,j}$ is the Kronecker delta function, which equals one when $i = j$ and zero otherwise, and \mathbf{X}_i refers to the i th row of \mathbf{X} . Then, the values $\mathbf{y}_{\text{interp}}$ at the locations $\mathbf{X}_{\text{interp}}$, the variance $\mathbf{v}_{\text{interp}}$ at the locations $\mathbf{X}_{\text{interp}}$, and the log marginal likelihood [i.e. $\log p(\mathbf{y}_{\text{data}}|\mathbf{X}_{\text{data}})$] are

$$\mathbf{y}_{\text{interp}} = \mathbf{K}_{\text{mixed}}^T \mathbf{K}_{\text{data}}^{-1} \mathbf{y}_{\text{data}}; \quad (4.16)$$

$$\mathbf{v}_{\text{interp}} = \text{diag}(\mathbf{K}_{\text{interp}} - \mathbf{K}_{\text{mixed}}^T \mathbf{K}_{\text{data}}^{-1} \mathbf{K}_{\text{mixed}}); \quad (4.17)$$

$$\log p(\mathbf{y}_{\text{data}}|\mathbf{X}_{\text{data}}) = -\frac{1}{2} \mathbf{y}_{\text{data}}^T \mathbf{K}_{\text{data}}^{-1} \mathbf{y}_{\text{data}} - \frac{1}{2} \log |\mathbf{K}_{\text{data}}| - \frac{N}{2} \log 2\pi; \quad (4.18)$$

where N is the length of the vector \mathbf{y}_{data} and diag means to take only the elements from the main diagonal (i.e. where $i = j$).

To connect each of these symbols with something concrete, in Figure 4.2

- \mathbf{X}_{data} is the x -positions of the blue dots;
- \mathbf{y}_{data} is the y -values of the blue dots;
- $\mathbf{X}_{\text{interp}}$ is the x -positions of the blue curve;
- $\mathbf{y}_{\text{interp}}$ is the y -values of the blue curve;
- $\mathbf{v}_{\text{interp}}$ is related to the width of the shaded region;
- σ_n^2 is the minimum variance, which is at the data point locations;
- $\sigma_y^2 + \sigma_n^2$ is related to the maximum width of the shaded region, which is far from any data points.

The log marginal likelihood is not as easy to picture. The notation $\log p(\mathbf{y}_{\text{data}}|\mathbf{X}_{\text{data}})$ indicates that the log marginal probability is the log of the probability of the values \mathbf{y}_{data} given \mathbf{X}_{data} . Since the predicted values $\mathbf{y}_{\text{interp}}$ do not have to go through the data values \mathbf{y}_{data} , the greater the difference between the model and the data, the lower the probability and the worse the model, which is why maximizing the log marginal likelihood is a reasonable way to improve the model. Conceptually, in Equation (4.18) the first term represents how well the data fits the model where a better fit increases the log marginal likelihood. The second term represents the model's complexity where the log marginal likelihood decreases with greater complexity. Lastly, the third term normalizes the

probability. Thus, maximizing the log marginal likelihood balances model fit and complexity.

For this application, the most important aspect of these equations is that the variance $\mathbf{v}_{\text{interp}}$ is independent of \mathbf{y}_{data} , which means that it can be evaluated without BEM calculations. Some BEM computations are needed to determine the hyper-parameters the first time that $\mathbf{v}_{\text{interp}}$ is evaluated, but any other time does not require BEM calculations. Moreover, changing \mathbf{X}_{data} does not require new BEM computations.

§4.5 Determining the hyper-parameters

The last step before estimating the variance is to find the hyper-parameters. Since the hyper-parameters are generally unknown, usually one optimizes them using data. This optimization problem requires the following:

- Acquire data;
- Define the objective function;
- Specify any constraints (if any);
- Choose an optimization approach;
- Avoid small local maximum.

§4.5.1 Sample data

Optimizing the hyper-parameters requires some data, so 2.5D BEM provides the sample data. Again, the question is where to put the points. One option is to randomly sample the entire 5D domain with up to 10,000 points. The number of sample points is limited because calculating the log marginal likelihood requires inverting \mathbf{K}_{data} , which is an $O(N^3)$ operation.

To reduce the computational cost, the 5D problem can be broken down into five 1D problems, which assumes the covariance is only weakly affected by quadratic terms. This option is less expensive because fewer points are necessary to adequately sample five lines than to all of sample 5D space. Each line would hold four of the variables constant while changing the fifth variable. For example, the common point in all of the plots in Figure 4.4 is $(x_s, z_s, x_r, z_r, \Delta y) = (-2, 0, 2, 0, 0)$ m, and one at a time the variables have the following values:

- $x_s = [-35.00, -34.75, -34.50, \dots - 1.00]$ (137 points);
- $z_s = [0.00, 0.05, 0.10, \dots 4]$ (81 points);
- $x_r = [1.0, 1.5, 2.0, \dots 200]$ (399 points);
- $z_r = [0.0, 0.1, 0.2, \dots 35]$ (351 points);
- $\Delta y = [0.00, 0.25, 0.50, \dots 60]$ (241 points);

which makes a total of about 1200 points. The additional cost of computing these extra points is negligible compared to the total cost of computing over 1 million points.

An additional benefit of breaking the problem into five 1D problems is that the 1D data can easily be interpreted and checked. Figure 4.4 verifies that points far from the complex object and the ground are fairly similar (i.e. have high covariance) and that points close to the complex object and ground are less similar. Long horizontal stripes indicate that the values along that direction are very similar. Checking the data is also very important because excess numerical noise causes the optimization to converge to an unrepresentative solution. Increasing the accuracy of the numerical integrals and ensuring convergence in the 2.5D BEM calculations mitigates the numerical noise.




Figure 4.4: Sample data in each direction

Each of the plots illustrates the correction factors as a function of frequency (50 Hz – 5 kHz in third-octaves) and position. Each plot gives the correction along a line parallel to a different axis, which is given along the plot's x -axis. All of the lines intersect at $(x_s, z_s, x_r, z_r, \Delta y) = (-2, 0, 2, 0, 0)$ m.

§4.5.2 The objective function

Typically, the objective function is the log marginal likelihood, and the goal is to maximize it. This setup assumes that there is only one output variable with one associated log marginal likelihood, but the current application has 21 output variables (i.e. one for each third-octave band from 50 Hz to 5 kHz) that each has a corresponding log marginal likelihood. Instead, the objective function is the mean log marginal likelihood across all of the frequency bands.

The derivatives of the objective function with respect to the hyper-parameters are the mean derivatives across the frequency bands. For each frequency, the gradient (i.e. $\nabla_i f = \frac{\partial f}{\partial \theta_i}$) and Hessian (i.e. $\mathbf{H}_{ij}(f) = \frac{\partial^2 f}{\partial \theta_i \partial \theta_j}$) of the log marginal likelihood with respect to the hyper-parameters $\boldsymbol{\theta} = [\sigma_y \mathbf{m}^T \mathbf{b}^T]^T$ are

$$\nabla[\log p(\mathbf{y}|\mathbf{X})] = \frac{1}{2} \mathbf{y}^T \mathbf{K}^{-1} \nabla(\mathbf{K}) \mathbf{K}^{-1} \mathbf{y} - \frac{1}{2} \text{tr}[\mathbf{K}^{-1} \nabla(\mathbf{K})]; \quad (4.19)$$

$$\begin{aligned} \mathbf{H}[\log p(\mathbf{y}|\mathbf{X})] = & \frac{1}{2} \mathbf{y}^T \mathbf{K}^{-1} [\mathbf{H}(\mathbf{K}) - 2 \nabla(\mathbf{K}) \mathbf{K}^{-1} \nabla^T(\mathbf{K})] \mathbf{K}^{-1} \mathbf{y} \\ & + \frac{1}{2} \text{tr}[\mathbf{K}^{-1} \nabla(\mathbf{K}) \mathbf{K}^{-1} \nabla^T(\mathbf{K}) - \mathbf{K}^{-1} \mathbf{H}(\mathbf{K})]. \end{aligned} \quad (4.20)$$

Since the derivatives only contain the data matrices and vectors, the data subscript has been dropped from these equations for brevity (i.e. $\mathbf{y} = \mathbf{y}_{\text{data}}$, $\mathbf{X} = \mathbf{X}_{\text{data}}$, and $\mathbf{K} = \mathbf{K}_{\text{data}}$). To derive these derivatives, notice that \mathbf{K} is a symmetric matrix (i.e. $\mathbf{K} = \mathbf{K}^T$) and $(\log|\mathbf{B}|)' = \text{tr}(\mathbf{B}^{-1}\mathbf{B}')$ where \mathbf{B} is an invertible matrix, tr is short for trace and means to sum the elements along the diagonal, and the prime denotes a derivative. Also, be careful about the high level of abstraction; $\nabla(\mathbf{K})$ is a vector of matrices and $\mathbf{H}(\mathbf{K})$ is a matrix of matrices. The term $\nabla^T(\mathbf{K})$ transposes the gradient vector but all of the matrices within the gradient are not transposed.

The remaining derivatives (i.e. $\nabla(\mathbf{K})$ and $\mathbf{H}(\mathbf{K})$) depend on the covariance function and involve the chain, product, and quotient rules for derivatives. For the 1D case and only one element of \mathbf{K} , let Equation (4.12) be expressed as four functions

$$k = f_1 e^{f_2}; \quad f_1 = \sigma_y^2; \quad f_2 = \frac{f_3}{f_4}; \quad f_3 = -(x - x')^2; \quad (4.21)$$

$$f_4 = m^4 x x' + m^2 b^2 (x + x') + b^4. \quad (4.22)$$

The gradients of these functions with respect to $[\sigma_y \quad m \quad b]^T$ are

$$\nabla(k) = k \left[\frac{\nabla(f_1)}{f_1} + \nabla(f_2) \right]; \quad \nabla(f_1) = \begin{bmatrix} 2\sigma_y \\ 0 \\ 0 \end{bmatrix}; \quad \nabla(f_2) = -f_2 \frac{\nabla(f_4)}{f_4}; \quad (4.23)$$

$$\nabla(f_4) = \begin{bmatrix} 0 \\ 4m^3 x x' + 2mb^2(x + x') \\ 2m^2 b(x + x') + 4b^3 \end{bmatrix}. \quad (4.24)$$

The gradient of f_3 is not included because it (and the Hessian) is all zeroes. The Hessians of the functions are

$$\mathbf{H}(k) = \frac{\nabla(k)\nabla^T(k)}{k} + k \left[\frac{\mathbf{H}(f_1)}{f_1} - \frac{\nabla(f_1)\nabla^T(f_1)}{f_1^2} + \mathbf{H}(f_2) \right]; \quad (4.25)$$

$$\mathbf{H}(f_1) = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}; \quad \mathbf{H}(f_2) = \frac{\nabla(f_2)\nabla^T(f_2)}{f_2} + f_2 \left[\frac{\nabla(f_4)\nabla^T(f_4)}{f_4^2} - \frac{\mathbf{H}(f_4)}{f_4} \right]; \quad (4.26)$$

$$\mathbf{H}(f_4) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 12m^2 x x' + 2b^2(x + x') & 4mb(x + x') \\ 0 & 4mb(x + x') & 2m^2(x + x') + 12b^2 \end{bmatrix}. \quad (4.27)$$

§4.5.3 Constraints

The form of Equation (4.12) was chosen specifically to avoid constraints. This optimization problem does not have any equality or inequality constraints. Although the denominator of the exponent must be positive, the hyper-parameters m and b are squared so that those hyper-parameters do not need any inequality constraints. Thus, this optimization problem does not have any equality or inequality constraints.

§4.5.4 Optimization method

§4.2.3 gives three common optimization methods: steepest decent, Newton, and quasi-Newton approaches. Since the objective function (i.e. the log marginal likelihood) and its first and second derivatives are known and there are only a few hyper-parameters, the most appropriate method is Newton's method without any constraints. In Matlab, the appropriate function is `fminunc()`¹³², which uses a trust region approach¹³³.

§4.5.5 Improving the local results

Since the problem is not convex, any optimized solution might only be a local maximum. Thus, it is quite possible that the optimization would result in a very small (i.e. poor) local maximum. To mitigate this problem, the algorithm described above uses multiple randomly generated start locations. These calculations are independent so that they can be run in parallel. Any future steps use the hyper-parameters that correspond to the best objective function.

§4.6 Minimizing the variance

After determining the hyper-parameters, the next optimization problem is to minimize the variance with respect to the data point locations \mathbf{X}_{data} . Similar to the last section, this section

- Defines the objective function;
- Specifies any constraints;
- Determines appropriate optimization methods.

§4.6.1 Objective function

Since the variance is also a function of position (i.e. it has multiple values), there is not a clear best objective function because there is more than one way to combine the points. Minimizing the mean value with respect to \mathbf{X}_{data} seems to be the most natural choice, but

it requires calculating the variance, which involves a matrix inversion, at a large number of locations and does not actually produce the desired effect. Minimizing the mean produces small variance between the point that are furthest apart, which are the points that are far from the complex object and are arguably the least important points.

Instead, the objective function is the global maximum variance, which means that this optimization problem requires a global optimization problem with respect to $\mathbf{X}_{\text{interp}}$ to evaluate the objective function. Thus, before moving forward with minimizing the maximum variance, it is necessary to have a method to find the global maximum.

In general, the problem of finding a global maximum for a non-convex function (the variance is not a convex function) is very difficult. However, knowledge about the specific function can make the problem much more manageable. First, the range is set and finite, so a brute force approach would finely sample the entire range for each dimension (i.e. $x_s, z_s, x_r, z_r, \Delta y$) and pick the largest value to be the value of the objective function. Second, the variance is smallest at the data points and largest about halfway between the data points. Thus, sampling the function at the halfway points between each of the data points and choosing the largest value is a good, inexpensive first approximation. In practice, the error is typically about 5-10%, and this level of accuracy may be sufficient for some applications. The maximum is not exactly at the midpoint between the data points because the characteristic length of the covariance function is not constant and increases with the dimension, so the maximum is actually slightly less than the midpoint. This information suggest finely sampling the regions just less than halfway between each of the data points and letting the objective function value be the maximum of these values. Finally, if accuracy is very important, then another optimization could be performed for each midpoint where the midpoint is the initial guess. Thus, the objective function would be the largest maximum across all midpoints and dimensions. This is probably the most expensive option.

§4.6.2 Constraints

To be clear, the goal is to minimize with respect to \mathbf{X}_{data} the maximum of the variance with respect to $\mathbf{X}_{\text{interp}}$. The number of variables and the constraints both depend on the model for \mathbf{X}_{data} . For example, the points could be specified with a simple exponential function like Equation (4.9), which would have only one variable per dimension and no constraints, or by letting each data point position be a variable, which would have many variables with the constraints that the positions are all inside the range and no data points are at the same location. The first option is far less expensive because there is only one variable and works well with the current covariance function but provides less flexibility. The second option gives great freedom but is much more expensive. There are also many

options in between and the best model depends on the application. If more than one model is plausible, then try as many as practicable.

§4.6.3 Optimization method

§4.2.3 discusses three common optimization methods: steepest decent, Newton, and quasi-Newton approaches. In this case, the analytic derivatives are even more involved than for the log marginal likelihood case and are different for each covariance function and model for \mathbf{X}_{data} . To indicate the complexity, the first step to calculate the analytical gradient and Hessian would be differentiating Equation (4.17) with respect to \mathbf{X}_{data} [or potentially c_2 from Equation (4.9)]. Using only one interpolation point, the general forms are

$$\nabla(\mathbf{v}_{\text{interp}}) = \mathbf{K}_m^T \mathbf{K}_d^{-1} \nabla(\mathbf{K}_d) \mathbf{K}_d^{-1} \mathbf{K}_m - 2\mathbf{K}_m^T \mathbf{K}_d^{-1} \nabla(\mathbf{K}_m); \quad (4.28)$$

$$\begin{aligned} \mathbf{H}(\mathbf{v}_{\text{interp}}) = & \nabla(\mathbf{K}_m^T) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_d) \mathbf{K}_d^{-1} \mathbf{K}_m + \mathbf{K}_m^T \mathbf{K}_d^{-1} \nabla(\mathbf{K}_d) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_m) \\ & - 2\mathbf{K}_m^T \mathbf{K}_d^{-1} \nabla(\mathbf{K}_d) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_d) \mathbf{K}_d^{-1} \mathbf{K}_m + \mathbf{K}_m^T \mathbf{K}_d^{-1} \mathbf{H}(\mathbf{K}_d) \mathbf{K}_d^{-1} \mathbf{K}_m \\ & + \mathbf{K}_m^T \mathbf{K}_d^{-1} \nabla(\mathbf{K}_d) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_m) + \nabla(\mathbf{K}_m^T) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_d) \mathbf{K}_d^{-1} \mathbf{K}_m \\ & - 2\nabla(\mathbf{K}_m^T) \mathbf{K}_d^{-1} \nabla^T(\mathbf{K}_m) - 2\mathbf{K}_m^T \mathbf{K}_d^{-1} \mathbf{H}(\mathbf{K}_m); \end{aligned} \quad (4.29)$$

where again for brevity the subscripts have been shortened so that $\mathbf{K}_d = \mathbf{K}_{\text{data}}$ and $\mathbf{K}_m = \mathbf{K}_{\text{mixed}}$. The majority of the work still remains to calculate $\nabla(\mathbf{K}_{\text{data}})$, $\nabla(\mathbf{K}_{\text{mixed}})$, $\mathbf{H}(\mathbf{K}_{\text{data}})$, and $\mathbf{H}(\mathbf{K}_{\text{mixed}})$.

Thus, instead of calculating the gradient and Hessian analytically, the gradient is approximated using finite differences and the Hessian is approximated using the BFGS^{134–137} quasi-Newton approach. Since there are no constraints, the appropriate Matlab function is again `fminunc()`. If there are constraints then use `fmincon()`.

§4.7 Choosing N

The final unknown that must be optimized is the number of points for each dimension (i.e. $x_s, z_s, x_r, z_r, \Delta y$). The algorithm has the following steps:

1. Guess N for every dimension
2. Minimize the maximum variance for every dimension
3. Add one to the N associated with dimension that has the largest maximum variance
4. Repeat steps 2-3 until the product of all the N s is over some specified limit

For step 1, the optimization process requires an initial guess for the number of points N for each dimension (i.e. $x_s, z_s, x_r, z_r, \Delta y$). Since the points are on a grid, the total number of points is the product of all of the N s. The initial value for each N is not very important because the optimization process updates this number on each iteration. However, all of the values should be less than the expected optimal values because the optimization procedure can only increase N for each dimension. For example, since the product of all of the optimized N s should be about 1 million in the current application, then N for each dimension could start at five because each dimension probably needs more than five points to minimize the maximum variance.

For step 2, §4.6 discusses minimizing the maximum variance in detail, but there are a couple ways to decrease the computation time for multiple similar evaluations. For instance, the first time through the steps, every dimension's variance must be optimized and the maximum must be found. However, since only one dimension is modified on each iteration, this information should be stored so that the same work does not have to be repeated. In addition, in some cases the optimal values are not changed dramatically by adding one additional point; thus, the previous optimal values should also be stored to use as initial guesses in the next iteration.

For step 3, find the dimension with the largest maximum variance and add one to the associated N value. This step compares the final objective function values from the previous step and updates the appropriate N .

For step 4, the algorithm is terminated when the total number of points (i.e. the product of all of the N s) reaches a certain threshold based on available memory. It's also wise to look at how the objective function (i.e. the maximum variance) is changing as the threshold is being reached. On the one hand, if the objective function only modestly improves for each additional point, then perhaps the additional points are not worth the extra storage costs. On the other hand, if the objective function is improving rapidly with each additional point, then perhaps the limit should be slightly increased to take advantage of the reduced maximum variance. In addition, there are other reasonable stopping criteria; for example, the algorithm could terminate when the maximum variance falls below a preset threshold.

§4.8 Comparison

Finally, this section implements a simplified version of the whole optimization process to demonstrate how the optimization process works and the expected level of improvement. Instead of doing the whole 5D problem, this section goes through a 2D problem, which makes 1) visualizing the results easier and 2) calculating densely spaced reference results

feasible. In addition, two other non-optimized solutions (linear and exponential spacing) are used for comparison purposes. This section takes the same format as the sections in the previous validation chapter with description, results, and analysis subsections.

§4.8.1 Description

For the current problem, the complex and simplified objects are the same T-barrier and I-barrier from §3.2, and the parameters are the same as the hard ground case in §3.3 except for the source and receiver positions. Since this example only evaluates the table of corrections and does not even use EM, the source and receiver positions are not expressed in the global coordinates where the source and receiver each have three coordinates but rather in the local coordinates of the T-barrier. To simplify the problem, three of the five local variables are fixed: $x_s = -3.0$ m, $z_s = 0.3$ m, and $z_r = 1.5$ m; and the other two variables have the following ranges: $x_r = [1, 200]$ m and $\Delta y = [0, 60]$ m.

For these two dimensions, the data locations are determined using linear, exponential, and optimized distributions. The linear distribution uses Equation (4.5) with the range and number of points along each dimension to choose the point locations. The only unknown is the number of points for each dimension. According to §2.6.1, each dimension can contain about 19 points. However, using 19 points for both dimensions means that the points will be much closer together along Δy than along x_r because Δy has a much smaller range. At this point, there is no reason to expect the correction factors to change more rapidly in the Δy -direction than the x_r -direction, so the number of points for each dimension should be proportional to the range. Thus, keeping the total number of points less than $19^2 = 361$ and making the number of points for each dimension proportional to the range yields 9 points for Δy and is 30 points for x_r (i.e. $30/9 \approx 199/60$ and $9 \times 30 = 270 < 361$). The number of points could be increased slightly, but these numbers are sufficient for a simple comparison. Using (4.5) yields the linear distribution of points in Table 4.1.

Next, the exponential distribution uses Equation (4.7) with the range and number of points per dimension. To keep the comparison as fair as possible, all of the methods use the same number of points per dimension, so the points follow from Equation (4.7) and are in Table 4.1.

Table 4.1: Data positions for comparison

These are the data locations for the linear, exponential, and optimized distributions. Figure 4.5 illustrates the data locations.

n	Linear		Exponential		Optimized	
	x_r (m)	Δy (m)	x_r (m)	Δy (m)	x_r (m)	Δy (m)
1	1.00	0.00	1.00	0.00	1.00	0.00
2	7.86	7.50	1.20	0.67	1.34	2.50
3	14.72	15.00	1.44	1.79	1.73	5.75
4	21.59	22.50	1.73	3.67	2.19	9.98
5	28.45	30.00	2.08	6.81	2.74	15.49
6	35.31	37.50	2.49	12.06	3.37	22.67
7	42.17	45.00	2.99	20.83	4.12	32.01
8	49.03	52.50	3.59	35.49	5.00	44.17
9	55.90	60.00	4.31	60.00	6.03	60.00
10	62.76		5.18		7.24	
11	69.62		6.22		8.66	
12	76.48		7.46		10.32	
13	83.34		8.96		12.28	
14	90.21		10.75		14.57	
15	97.07		12.91		17.26	
16	103.93		15.49		20.42	
17	110.79		18.60		24.13	
18	117.66		22.33		28.49	
19	124.52		26.81		33.59	
20	131.38		32.18		39.59	
21	138.24		38.63		46.63	
22	145.10		46.37		54.89	
23	151.97		55.67		64.58	
24	158.83		66.83		75.95	
25	165.69		80.22		89.30	
26	172.55		96.30		104.97	
27	179.41		115.61		123.36	
28	186.28		138.78		144.94	
29	193.14		166.60		170.27	
30	200.00		200.00		200.00	

Using the optimization procedure discussed in this chapter produces the optimized locations. For simplicity, the 2D problem is broken down into two 1D problems. The distribution function is Equation (4.9) where the coefficient c_2 must be found for each dimension (i.e. x_r and Δy). For each 1D problem, the covariance function is given by Equation (4.12). Then, 2.5D BEM calculates the x_r and Δy data in Figure 4.4. Using the optimization procedure in §4.5 with $\sigma_n = 0.05$ dB produces the hyper-parameters for each dimension. Specifically, minimizing the log marginal likelihood yields the hyper-parameters. For this case, Table 4.2 gives the hyper-parameters (σ_y , m , and b) and the optimized objective function [i.e. $\ln p(\mathbf{y}|\mathbf{X})$].

Table 4.2: Optimization values

These are the intermediate values for calculating the optimized values. The first step of the optimization finds the hyper-parameters by maximizing the log marginal likelihood, and the second step determines the coefficient c_2 by minimizing the maximum variance.

	x_r	Δy
σ_y	4.71	2.92
m	0.654	0.247
b	0.623	0.709
$\ln p(\mathbf{y} \mathbf{X})$	664	-519
c_2	1.38	2.87
$\max(\mathbf{v}_{\text{interp}})$	0.004	8.52

Recall that m^2 is the slope and b^2 is the intercept of the characteristic length, and remember that a shorter characteristic length indicates more rapid fluctuations in the data. The values of m and b suggest that close to the origin the characteristic length is similar along both dimensions, but moving away from the origin the characteristic length grows faster in the x_r -direction. Looking at Figure 4.4 confirms this result because the correction factors do not vary much for $x_r > 30$ m but they do throughout the entire range of Δy .

When the goal is to minimize the objective function, the much smaller value of the objective function [i.e. $\ln p(\mathbf{y}|\mathbf{X})$] for Δy than for x_r indicates that the Gaussian process does not fit the Δy data as well as the x_r data. Looking closely at the data in Figure 4.4 reveals that there is more numerical noise in the data along Δy than along x_r , especially above 1 kHz. This noise comes from numerically approximating integrals in 2.5D BEM. The best approach would be to recalculate the data using a high precision to mitigate the noise. Another option would be to apply a smoothing filter like a moving average to reduce the noise. If the noise cannot be reduced, then the standard deviation of the noise (σ_n) could be made larger or even made a hyper-parameter. Right now, the model sets

$\sigma_n = 0.05$ dB because that is the estimated level of uncertainty that is expected from the 2.5D BEM calculations.

Finally, minimizing the maximum variance (§4.6) determines the coefficient c_2 from Equation (4.9), and Table 4.2 gives the results. Recall from Figure 4.1, that c_2 is an inverse measure of the concentration of the point. The points are approximately linearly spaced for large values of c_2 and exponentially spaced when $c_2 = 1$. Since both of the values of c_2 are slightly larger than one, both dimensions have point distributions between linear and exponential but much closer to exponential. Since c_2 is larger for Δy than for x_r , the point distribution for Δy is slightly more linear than for x_r .

Since the number of points in each direction are already set (i.e. 9 for Δy and 30 for x_r), the optimization procedure does not have to choose how many point should be used. However, the much smaller maximum variance [i.e. $\max(\mathbf{v}_{\text{interp}})$] for x_r than for Δy indicates that the number of points along Δy should be increased and the number of points along x_r should be decreased. When the number of points in each direction is optimized the maximum variance in each direction tends to be approximately equal instead of being multiple orders of magnitude different. The algorithm wants more points along Δy to be able to model the numerical noise in that direction, but recalculating the data with less noise would reduce the need for points in the Δy -direction. Using a smoothing filter is another option if the noise cannot be eliminated at a reasonable cost.

Finally, plugging the values of c_2 into Equation (4.9) produces the optimized locations, which are given in Table 4.1. Overall, the values in Table 4.1 are very similar for exponential and optimized distributions, so the optimized results are not expected to be much better than the exponential results. However, this conclusion could not have been known before performing the optimization. In addition, the optimization method provides a numerical way to choose the number of points for each dimension, which was not used here.

§4.8.2 Results

After determining the locations of the data points, the next step is to calculate the values at the locations. 2.5D BEM provides the data for the tables of corrections. Then, to evaluate the interpolation error, 2.5D BEM calculated the levels on a fine grid of receivers over the entire domain of x_r and Δy with 25 cm between each point in both dimensions. Now, the actual values calculated directly with 2.5D BEM on a fine grid can be compared with the values calculated by interpolating the 2.5D BEM results for the linear, exponential, and optimized distributions. Figure 4.5 gives these results. All three distributions perform fairly well with the linear distribution doing the worst and the optimized distribution only doing slightly better than the exponential distribution.

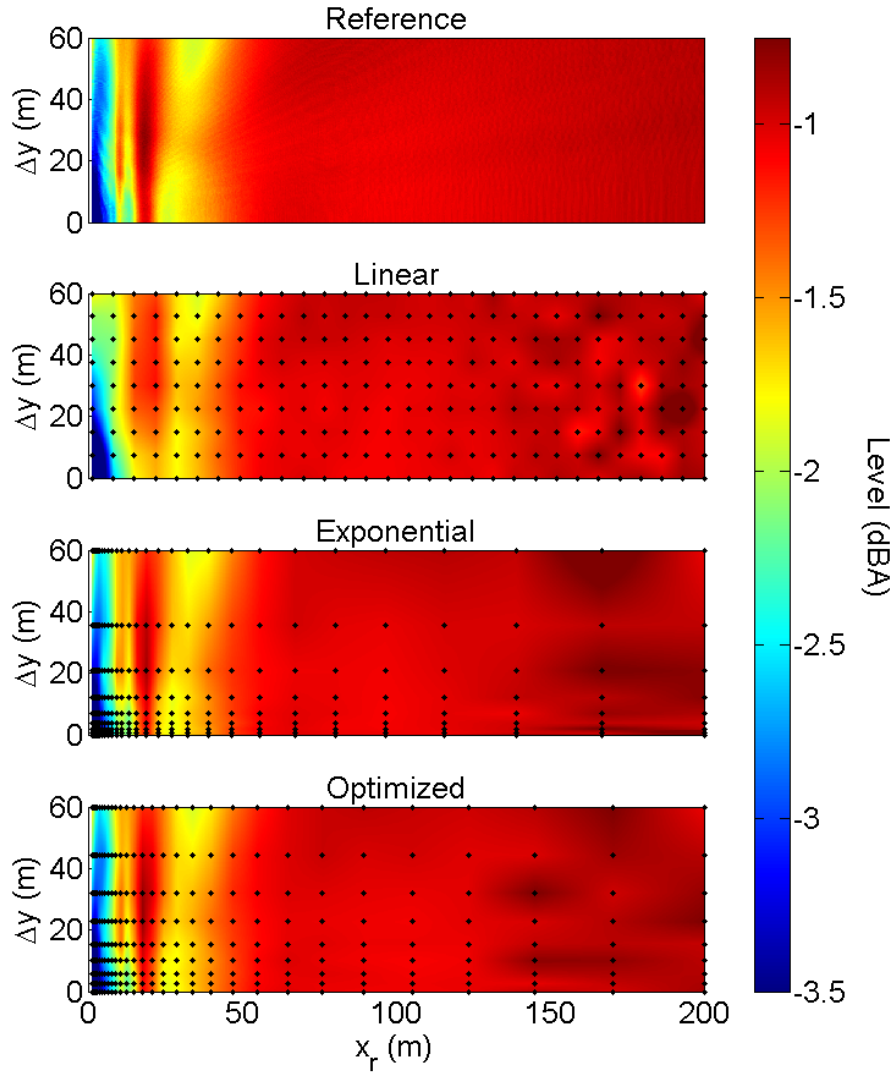


Figure 4.5: T-top correction using different methods

The T-top correction is calculated four different ways: 1) directly with 2.5D BEM called the reference result, 2) interpolating the linearly distributed data, 3) interpolating the exponentially distributed data, and 4) interpolating the optimized data distribution. The black dots show where the data points are for the interpolated data. The levels shown are a weighted average over 50-5000 Hz (A-weighted and weighted to look like a car spectrum).

In addition, Figure 4.6 gives the error between each of the interpolated distributions and 2.5D BEM. The error is the interpolated value minus the reference value. This plot demonstrates more clearly that the exponential and optimized distributions have lower error than the linear distribution, especially for $x_r < 30$ m. Although, the optimized results are only slightly better than the exponential results. For all of the methods, the interpolation error in Figure 4.6 is significantly smaller than the actual correction factor in Figure 4.5 (notice the different scale).




Figure 4.6: Error levels for the three distributions

These plots give the error of the linear, exponential, and optimized distributions compared to the reference 2.5D BEM computations. Notice the much smaller scale in this figure than in Figure 4.5.

§4.8.3 Analysis

Finally, taking the data from Figure 4.6 and calculating some important statistical measures produces Figure 4.7. Linear, exponential, and optimized distributions all produce a mean error that is close to zero. Comparing the spread of the data demonstrates that the linear distribution has the largest range, standard deviation, and RMS error, and the optimized distribution has the smallest.

While the optimized distribution does perform the best, the small improvement makes justifying using the optimized method over the exponential distribution difficult because the optimized method is significantly more complicated and expensive. In addition, the optimized method has shown to be sensitive to numerical noise in the 2.5D BEM data. However, what has been demonstrated here is a very restricted optimization. Allowing the points to be arbitrarily placed along each dimension would improve the result, and using a different covariance function may also help.

The optimization method also chooses the number of points for each dimension in a numerical way, which is a benefit that is not demonstrated in this example. For example,

instead of using 9 points for Δy and 30 points for x_r , the example could have used 11 points for Δy and 25 points for x_r , which keeps the total number of points about the same. Calculating both grids and evaluating the interpolation error in the full 5D setting would require too many BEM computations, so the optimization procedure provides an efficient way to choose the number of points along each dimension.




Figure 4.7: Error comparison for the three distributions

The horizontal lines for each box and whisker plot from bottom to top are the smallest value, first quartile, second quartile, third quartile, and largest value, and the values below each box are the mean (μ), standard deviation (σ), and root mean squared error (RMS). The left box represents the interpolation error using the linear distribution. The center box represents the interpolation using the exponential distribution. The right box represents the interpolation using the optimized distribution.

§4.9 Summary

This chapter discusses the issue of where to put the data points so that the interpolation error is minimized. This problem occurs frequently across many disciplines. For example, where should weather stations be located to obtain the greatest accuracy from a limited number of stations? Where should a finite set of noise monitors be put to best approximate the noise around an airport? Since actually calculating the mean error for many different point distributions is infeasible, this chapter chooses the locations that minimize the variance of the interpolated points, which is estimated using a Gaussian process. The variance is a measure of the possible spread or uncertainty, so as the variance decreases the errors should also decrease. This change probably introduces some error (i.e. the optimized locations may not be the exact locations that minimize the error the best), but it also makes the problem feasible.

The optimization procedure has the following major steps:

1. Choose a distribution function
2. Select a covariance function
3. Obtain sample data using 2.5D BEM
4. Optimize the hyper-parameters (i.e. the covariance function parameters) by maximizing the log marginal likelihood
5. Optimize the distribution function parameters by minimizing the maximum variance

Chapter 5 Conclusion

§5.1 Objectives

The primary objectives of this chapter are to

- synthesize the major results from the previous chapters
- highlight the novelty and importance of the work
- present a perspective on future work in this area

§5.2 Principal results

The most important result of this dissertation is the hybrid method that extends EMs to more complicated geometries while being far less expensive than 3D reference methods. The increased complexity compared to EMs results from incorporating 2.5D BEM calculations so that objects can have an arbitrary cross-section. The improved speed compared to 3D BEM stems from assuming that the attenuation of a complex object compared to a reference object is mostly independent of its surrounding so that the extra attenuation can be approximated using 2.5D BEM. This assumption is similar to the EM assumption that attenuation from geometrical divergence is independent of the attenuation from atmospheric absorption. Furthermore, computational efficiency is also improved by storing a large number of 2.5D BEM results in a table that can easily be searched and interpolated whenever EM needs the information instead of performing a 2.5D BEM computation for every propagation path. Several additional aspects of the hybrid method are important for its efficiency and accuracy.

First, the hybrid method uses EMs as the efficient method. Certainly, other methods could be used instead, but EMs are particularly fitting because of their efficiency and

modeling procedure. Specifically, the correction factors are easy to apply because EMs combine the ray paths in a vertical plane through the source and receiver and separate the paths that travel over and around the complex object. Both of these aspects are important for approximating a 3D object using a 2.5D approach.

Second, the hybrid method uses 2.5D BEM as the reference method. As a 2.5D method, it efficiently models arbitrary cross-sections and point sources, and as a boundary element approach, it efficiently models hard flat ground and infinite domains. All four of these aspects are critical for the reference method. Other methods could be used, but 2.5D BEM is very appropriate.

Third, the hybrid method uses six Cartesian coordinates. Using fewer than six variables would dramatically decrease the accuracy of the hybrid method because Δy would have to be excluded and the extra attenuation varies just as much with Δy as the other variables. Using more than six variables would also decrease the accuracy of the hybrid method because there would not be enough data points to accurately represent how the extra attenuation changes with respect to each variable. The hybrid method uses Cartesian coordinates because the extra attenuation of complex objects does not have any nice symmetry for all complex objects.

Fourth, the hybrid method usually models the complex object on hard, flat ground to calculate the correction factors even if the ground is not hard. This approach sacrifices a little accuracy for huge reductions in computation time and memory. Specifically, 2.5D BEM is faster using hard ground than impedant ground, and more importantly the number of tables of corrections is dramatically reduced. Instead of a new table for every ground type, the hybrid method only needs one.

Fifth, the hybrid method models the effect of other objects in the propagation path on the extra attenuation by changing the apparent source and/or receiver positions for the complex object. Specifically, multiple diffractions, reflections, and lateral diffractions change the apparent source and/or receiver positions for the complex object. This choice is why the extra attenuation is not completely independent of the surrounding geometry and greatly improves the accuracy of the method for negligible cost.

Sixth, the hybrid method linearly interpolates data on a non-uniform grid to obtain the extra attenuation from the table of corrections. Using higher order interpolation methods like cubic interpolation is much more complicated and expensive because the interpolation is 5D. Using more flexible data structures (e.g. scattered or quad-tree) are also too computationally expensive to implement because of the large number of data points.

Seventh, the hybrid method can also extrapolate the correction factors. However, extrapolation is inherently less accurate than interpolation, so interpolation should be

used whenever possible. If many points must be extrapolated and they are affecting the accuracy of the results, then the range of the source and receiver positions should be increased.

Eighth, the hybrid method can benefit from optimizing the source and receiver positions to reduce the interpolation error. The optimization procedure minimizes the maximum variance, which is modeled with a Gaussian process. Effectively, the source and receiver positions are concentrated close to the complex object and the ground where the extra attenuation fluctuates the most rapidly.

Finally, this dissertation also models four cases using the hybrid method. The hard ground case verifies that the interpolation error is small regardless of the data points being far apart compared to a wavelength. The soft ground case demonstrates that extra attenuation term is fairly invariant with respect to the ground type so that only 2.5D BEM results using hard ground are necessary. The buildings case indicates that the reflection, multiple diffraction, and lateral diffraction approximations are sufficiently accurate. The fourth case establishes that the hybrid method can be applied to situations that are far too expensive for any 3D reference method. All of the cases indicate that the hybrid method makes a substantial contribution to EMs.

§5.3 Novelty and importance

This dissertation contains two major novelties:

- Hybrid method;
- Location optimization applied in acoustics.

The idea of combining EMs and BEM using a table of corrections is not new, but to the best of the author's knowledge, no one has fully specified a way to combine them. Conceptually, the biggest gap in the method was how to do anything more complicated than a simple diffraction over a complex object. No one had specified how the rest of the geometry (e.g. other buildings) would impact the extra attenuation term. This omission makes the method unusable in urban settings. In addition, the development of the hybrid method revealed many choices that have never been justified. From small decisions like choosing the number of variables to big ones like choosing the interpolation method and data structure, previous literature did not justify or even specify in many cases the approach. Chapter 2 seeks to give the reader this information.

Furthermore, to the best of the author's knowledge, no one has fully verified an approach combining EMs and BEM. In contrast, Chapter 3 gives multiple 3D validations against 2.5D BEM and FM-BEM. Certainly, the hybrid method can be further validated using

different geometries and measurements, but the current level of validation is absent from the literature for an equivalent hybrid method.

Finally, optimizing data locations by minimizing the variance is not new to those in the optimization field. However, this type of technique is new to those in acoustics and more specifically to urban outdoor sound propagation. In applying this technique to this application, two difficulties arose. First, the number of spatial dimensions is considerably higher. Typically, there are at most three spatial dimensions, but this application has five. Chapter 4 discusses how to keep the problem manageable by breaking the problem into five 1D problems. Second, other applications usually have only one output variable whereas in acoustics the output is a spectrum across frequency. Chapter 4 also explains how to handle that subtlety.

§5.4 Future work

As with any large project, there are still some ways that the hybrid method can be improved. This section discusses potential opportunities starting with the most important.

First, the hybrid method has been developed mostly only considering flat ground. While the current method can handle non-flat ground, the 2.5D BEM computations will be significantly more expensive. The hybrid method would be improved by finding the most efficient way to model non-flat ground in 2.5D BEM.

Second, the method can be further validated using different geometries and measurements. The hardest part of the validation is finding a reference method that can yield results for the whole frequency range. Moreover, measurements can be costly and time consuming but would further validate the method.

Third, the interpolation method and method to predict the variance are different, but finding a method that could do both would be better for two reasons. First, the optimization depends on the values predicted by the Gaussian process and not the linear spline. This difference likely causes a slightly suboptimal solution, which could be slightly improved if the methods were the same. Second, the current method does not provide any indication of the uncertainty of the extra attenuation. If the interpolation method could also provide the variance, then the user would have a better idea of the uncertainty of the value. The difficulty is that linear interpolation does not predict the variance and Gaussian processes are too expensive to use on the entire dataset because they require a dense matrix inversion.

Forth, the current implementation of the hybrid method uses an I-barrier that is 20 cm thick as the simplified barrier because that thickness corresponds to the T-barrier. However, since EM and 2.5D BEM do not predict the same levels for the I-barrier, an improvement might be possible if the width of the I-barrier in 2.5D BEM was optimized to predict the levels closer to EM. This suggestion has not been implemented because it would probably not provide a meaningful improvement.

Finally, the extrapolation function might be improved by calculating about 5,000 random samples that are close together but far from the ground and the complex object. This option would require additional 2.5D BEM computations, but might give better results. This idea was not implemented because the small potential improvement did not seem worth the cost.

§5.5 Summary

The hybrid method enables EMs to model complex geometries. To develop the hybrid method, this dissertation

- Elucidates and justifies the important decision required to implement the hybrid method
- Validates the hybrid method against reference computations
- Optimizes data locations to minimize the interpolation error

Extended summary

Environmental noise (i.e. road, rail, aircraft, and industrial noise) is ubiquitous in urban society and is associated with negative health outcomes. According to the World Health Organization, high noise exposure can lead to increased annoyance, sleep disturbance, cardiovascular disease, and cognitive impairment in children. Thus, noise is a public health concern, and its negative effects should be mitigated.

Noise can be reduced by decreasing the level of the source or by changing the propagation path. Depending on the situation, the source level could be decreased by advancing technology (e.g. quieter engine, tire, or muffler design in cars), changing the source (e.g. traveling by bike instead of by car), or eliminating the source (e.g. car-free pedestrian zones). The propagation path could be altered by changing the distance (e.g. changing flight paths to avoid flying over cities), geometry (e.g. a noise barrier between a highway and neighboring residents), or surface materials (e.g. using grass instead of cement) between the source and the receiver. Outdoor sound propagation models are important because they quantify the impact of each of these changes.

However, efficiently and accurately modeling urban sound propagation is difficult because cities are both large and complex. Because of the size of cities, reference approaches like the boundary element, finite element, and time-domain finite-difference methods are prohibitively expensive, and because of the complexity of cities, geometrical approaches like the engineering methods, ray tracing, and beam tracing would require too many orders of diffraction to obtain accurate results.

A hybrid approach is potentially more accurate than the geometrical methods alone and more computationally efficient than reference methods alone, but current hybrid approaches are still too expensive because they still need too many reference calculations for high frequencies in three dimensional space. In particular, most hybrid methods use a fictitious boundary to separate the domain into a complex region and a simple region. A

reference method accurately approximates the pressure/velocity on the fictitious surface, and then an efficient method propagates those values into the simple region for little cost. This approach becomes too expensive because the number of points on the surface is proportional to the frequency squared because it requires multiple points per wavelength in two dimensions.

This dissertation develops and validates a more efficient hybrid method that combines an engineering method (e.g. CNOSSOS-EU or Harmonoise) and the 2.5D boundary element method. The hybrid method uses an engineering method a few reasons. First, engineering methods can approximate the most important sources of attenuation (e.g. geometrical divergence, atmospheric absorption, the ground effect, reflections, and diffractions over/around simple objects). Second, engineering methods can model large domains because their cost is independent of frequency. Instead, the cost depends on the number of faces/edges and the maximum order of reflection/diffraction. Third, government agencies, consultants, and city planners all widely use engineering methods as the standard way to model outdoor sound propagation.

The hybrid method also uses the 2.5D boundary element method for a few different reasons. First, the 2.5D boundary element method accurately models arbitrarily-shaped objects with a constant cross-section, which is the principal improvement that the reference method offers the hybrid method. In addition, the boundary element method accomplishes this task very efficiently for this application. Outdoor sound propagation commonly has infinite domains with hard flat ground, which the boundary element method can model without increasing the size of the mesh. Finally, the hybrid method uses a 2.5D method instead of a 2D method because 2D methods can only model infinite coherent line sources when the engineering method uses point sources. The hybrid method also does not use a 3D method because a 3D method is much more expensive. The mesh of the 2.5D boundary element method is only 1D versus the 2D mesh of the 3D boundary element method, which substantially reduces the mesh's size and complexity. Furthermore, the 3D boundary element method cannot model the attenuation of a propagation path that only diffracts over a complex object because it does not separate the propagation paths that go over versus around the complex object.

To model outdoor sound propagation, the engineering methods start by setting up the scene. This step includes specifying the location of the sources, receivers, and any other objects in the scene, the impedances of all the surfaces, and the properties of the medium. The hybrid method modifies this step by also loading additional information about each complex object, which the next paragraph discusses in greater detail. Next, the engineering method finds the most important propagation paths between the source and the receiver. The hybrid method does not change this step. Then, the engineering method calculates all of the attenuations (e.g. geometrical divergence and atmospheric absorption). The hybrid method modifies this step by adding an extra attenuation term for

complex objects along the propagation path. The extra attenuation is the attenuation of the complex object minus the attenuation of the simplified object. Conceptually, this modification is a simple, but in practice approximating this extra attenuation is complicated and requires reference calculations. Finally, the engineering method sums all of the propagation path contributions to get the total field at the receiver.

Since performing a boundary element computation for every propagation path would be prohibitively expensive, before the engineering method begins, the hybrid method instead performs a large number of reference calculations over a wide range of source/receiver positions and frequencies. The hybrid method describes the source and receiver positions using five Cartesian coordinates that are defined relative to the complex object (i.e. the perpendicular distance of the source/receiver from the object, the height of the source/receiver above the base of the object, and the parallel distance between the source and receiver). For efficiency, the source/receiver positions are constrained to a non-uniform grid, which allows a point to be found using $O(\log n)$ operations. The boundary element method usually puts the complex object on hard flat ground because it efficiently approximates a large number of real situations. The hybrid method converts multiple single frequency boundary element method calculations to octave or third-octave results and then subtracts the simplified object level from the complex object level. Finally, the hybrid method sorts and tabulates the results for quick access when the engineering method needs them. This table is the additional information that is loaded when the scene is setup in the engineering method.

When the engineering method calculates the attenuations for each path, the hybrid method first checks the path to determine if the path diffracts over the complex object. If the complex object is not in the propagation path, then the extra attenuation is zero. If the complex object is in the propagation path, the hybrid method calculates the path length differences of each diffracting object in the propagation path. If the complex object does not have the largest path length difference, then the extra attenuation is zero. If the complex object has the largest path length difference, then the hybrid method calculates the local source/receiver positions. The table of corrections does not use the same coordinates as the engineering method because the engineering method does not have a fixed origin. Instead, in the engineering method the user is free to choose the origin that best suits the current application. Thus, these global coordinates in the engineering method must be converted to the local coordinates in the table of corrections. In addition, the location of the source/receiver for the table of correction is the apparent position of the source/receiver. In particular, reflections from vertical surfaces use the image source/receiver, and lateral diffractions use a source/receiver that appears to be coming from the diffraction edge but keeps the original path length.

After converting the source/receiver coordinates, the exact position is probably not in the table of corrections. If the point is inside the data's convex hull, then the data is linearly

interpolated using the points of the smallest hyper-rectangle that contains the interpolation point. Thus, the interpolation has two distinct steps: 1) find the smallest containing hyper-rectangle and 2) use the local data to approximate the value. Using points on a grid makes searching for the appropriate hyper-rectangle much more efficient. If the point is outside the data's convex hull, then the data is extrapolated so that extrapolation points close to the dataset are similar in value to the closest data point and extrapolation points far from the dataset decay to zero. If many points are outside of the table of corrections, then range of the input variables might need to be expanded.

As an initial validation, the hybrid method models a T-barrier with hard ground. The complex barrier is a T-barrier that is 3 m tall and 0.2 m thick and has a 1 m wide top, and the simplified barrier is an I-barrier that is the same as the T-barrier without the top. For comparison, the 2.5D boundary element method independently models the same scene for a source that is 0.3 m high and 3 m from the barrier and a dense horizontal grid of receivers that are 1.5 m high.

Overall, the two methods produce very similar results. For the T-barrier, the hybrid method reduces the RMS error by 1.5 dBA compared to the engineering method using the I-barrier. The largest source of error is the difference between the engineering method and the 2.5D boundary element method for the simplified I-barrier. This difference stems from fundamentally different assumptions in the two models. Specifically, the boundary element method sums pressures using the phase information, and the engineering method ignores the phase information because in real-world situations turbulence and surface imperfections alter the phase. In addition, the hybrid method has some interpolation error for the T-barrier, which could be decreased by using more points or by choosing better locations for the points.

Next, the hybrid method uses the same table of corrections for a T-barrier with soft ground similar to grass. For the T-barrier, the hybrid method reduces the RMS error by 1.1 dBA compared to the engineering method using the I-barrier. This case demonstrates that a table of corrections that is calculated using hard ground can be reasonably applied to other types of grounds. Using a different ground type does slightly increase the error, but the large reduction in computation time and computer memory is likely sufficient compensation. Ultimately, the user can decide based on the required accuracy.

As a final test case, the hybrid method models a T-barrier with two buildings. This case still uses that same table of the correction as the previous two cases, but the reference method must be changed to the fast-multipole boundary element method because this situation is fully three-dimensional and the regular boundary element method can only model up to 400 Hz for this scene. In contrast, the fast-multipole boundary element method can reach up to 1.6 kHz. This case demonstrates that the hybrid method still performs well even when the scene contains other objects. For the T-barrier, the hybrid

method reduces the RMS error by 2.0 dBA compared to the engineering method using the I-barrier.

As a proof of concept, the hybrid method finally models a T-barrier in a realistic, large scene (i.e. 180 m x 80 m) with several buildings and partially soft ground. This case showcases the hybrid method because the hybrid method produces reasonable results when the engineering method would be inaccurate and the reference methods can only model low frequencies. In contrast, the hybrid method can easily model outdoor sound propagation for the full range of frequencies for this scene and even larger scenes with more buildings.

In the hybrid method, an optimization problem naturally arises to minimize the interpolation error given a maximum number of points in the table of corrections. The distribution of the data points affects the interpolation error, so changing the distribution could decrease the interpolation error. Without optimizing, the hybrid method could use a linear or exponential distribution. Along each dimension, a linear distribution evenly spreads out the points and an exponential distribution concentrates the points in one area. Since the correction changes most rapidly close to the complex object and close to the ground, an exponential distribution is better than the linear distribution, but there are many other options that have not even been considered. If the distribution has more degrees of freedom, then the interpolation error would decrease, and the problem becomes how to determine the optimization parameters.

The most straightforward approach is to choose the values that minimize the interpolation error; however, this option is prohibitively expensive because of the required data in order to make a reference dataset. Instead of minimizing the interpolation error directly, the hybrid method minimizes the variance of the interpolated value using a Gaussian process, which indirectly minimizes the interpolation error. The hybrid method uses a Gaussian process because it is a standard way to interpolate data that also gives the variance of the interpolated value. The interpolation error should decrease as the variance of the interpolated values decreases. Again, the hybrid method uses this indirect approach to minimize the interpolation error because the direct approach requires too many boundary element method computations.

Using a Gaussian process reduces the data requirements, but it does not eliminate them. Even though calculating the variance does not directly depend on any data values, the variance does depend on some hyper-parameters, which are best approximated using data. After finding the hyper-parameters by optimizing the log marginal likelihood, a Gaussian process yields the variance as a function of the data locations and not the value at those locations, so repeatedly evaluating the objective function does not require any more boundary element computations. This reduction in boundary element computations

is the primary reason to introduce the substantial extra complexity associated with a Gaussian process rather than minimize the interpolation error directly.

This optimization procedure produced a small improvement when there was only one parameter per dimension (i.e. five parameters in total), but it would likely produce even better results with a more flexible model. In addition, optimizing the distribution gives a more rigorous way to choose the number of points along each dimension and the distribution of the points within each dimension. Knowing that a distribution is better than other nearby distributions is also a benefit even if the interpolation error does not drastically decrease. However, the optimized distribution is not necessarily optimal because the optimization problem is not convex.

In conclusion, the hybrid method provides a novel and useful tradeoff between efficiency and accuracy for modeling urban outdoor sound propagation with complex geometries. The hybrid method is more accurate than an engineering method for complicated geometries and more efficient than the reference methods. After creating a database of tables of corrections for the most common complex objects, the hybrid method only increases the computation time because the appropriate tables of corrections must be loaded into memory, which only has to be done once per scene and currently requires less than one minute per table to complete. All of the other computations are negligible compared to the rest of the computations in the engineering method.

Résumé étendu

Le bruit environnemental (i.e. routier, ferroviaire, aérien et industriel) est omniprésent en milieu urbain et est à l'origine de conséquences sanitaires néfastes. Selon l'Agence Mondiale pour la Santé (OMS), une exposition à des niveaux de bruit élevés accroît, par exemple, la gêne, les perturbations du sommeil, les maladies cardiovasculaires et les troubles cognitifs chez l'enfant. Le bruit est donc un sujet de santé publique très important, et sa réduction constitue un enjeu sociétal majeur.

Les nuisances sonores peuvent être limitées en réduisant les niveaux sonores à la source ou en agissant sur les chemins de propagation. Selon la situation, le niveau sonore émis peut être réduit au moyen de solutions techniques avancées (e.g. des moteurs plus silencieux, des pneus ou des pots d'échappement optimisés), en changeant la source (e.g. en utilisant la bicyclette plutôt que la voiture), ou en éliminant la source (e.g. zones piétonnes sans automobiles). Les chemins de propagation peuvent être, quant à eux, altérés pour réduire l'exposition au bruit, en modifiant les distances (e.g. en modifiant les trajectoires de vol afin d'éviter le survol des zones urbaines), les géométries (e.g. écrans acoustiques entre une autoroute et les habitations voisines), ou au moyen de matériaux absorbants disposés sur les surfaces de propagation (e.g. utilisation d'herbe au lieu de ciment) entre sources et récepteurs. Dans cette démarche de réduction des nuisances, les modèles de propagation sonore en milieu extérieur ont donc un rôle important, puisqu'ils permettent de quantifier l'impact des actions mises en œuvre.

Cependant, il est difficile de modéliser de façon efficace et précise la propagation sonore en zones urbaines. Les approches dites de référence (éléments de frontière, éléments finis et approche temporelle par différences finies) ne sont pas adaptées à la taille des villes considérées, du fait des temps de calcul prohibitifs. Les approches géométriques, tout comme les approches d'ingénierie, moins coûteuses en temps de calcul, sont de leur côté, inaptes à rendre compte de la complexité des villes. Ces approches, comme celles par exemple basées sur les techniques de rayon et de faisceaux sonores, demanderaient des

ordres de réflexion et diffraction trop élevés pour obtenir des résultats avec une précision suffisante.

Une approche hybride, basée sur le couplage entre une méthode de référence et une méthode d'ingénierie, est potentiellement plus précise que les approches d'ingénierie et plus efficace (en temps de calcul) que les méthodes de référence. Toutefois, celles développées à ce jour demeurent encore trop coûteuses, notamment parce qu'elles demandent un nombre de calculs 3D en hautes fréquences trop important. En effet, la plupart des approches hybrides existantes nécessite l'utilisation de frontières fictives séparant le domaine d'étude en une zone « complexe » et une zone « simple ». Dans ce cas, la méthode de référence approxime précisément la pression et la vitesse sur la frontière fictive tandis que la méthode d'ingénierie effectue la propagation de ces quantités dans la région « simple » à moindre coût. Cette approche devient malheureusement rapidement onéreuse car le nombre de points sur cette surface fictive est fonction de la longueur d'onde, et croît avec la fréquence au carré.

Cette thèse développe et valide une approche hybride alternative plus efficace, en combinant une méthode d'ingénierie (e.g. CNOSSOS-EU or Harmonoise) et un calcul BEM 2.5D. La méthode d'ingénierie utilisée permet ainsi de modéliser les principales sources d'atténuation lors de la propagation acoustique (e.g. divergence géométrique, absorption atmosphérique, effets de sol, réflexions et diffractions sur ou autour d'objets simples) sur de larges domaines, et à moindre coût, puisque le calcul est indépendant de la fréquence. Le calcul BEM 2.5D permet quant à lui de modéliser précisément les objets de forme quelconque, invariants selon une dimension, ce qui est le principal atout de l'approche hybride comparée à l'approche d'ingénierie. De plus le modèle 2.5D correspond à une géométrie invariante selon un axe ce qui correspond à de nombreux objets tels un écran ou un tramway. Le choix d'un calcul 2.5D au lieu d'un calcul 2D est cohérent avec le modèle d'ingénierie utilisé, puisque cela permet de modéliser de manière effective des sources ponctuelles, alors qu'une approche 2D ne permettrait de considérer que des sources linéiques cohérentes. Le choix d'une approche 3D aurait été, quant à lui, beaucoup trop onéreux et n'aurait pas permis de maîtriser les effets de diffraction latérale aux extrémités des objets.

Afin de modéliser la propagation extérieure, les approches d'ingénierie commencent par une saisie de la scène, en spécifiant la position des sources, des récepteurs, de tous les objets du problème, des impédances de surface et des propriétés du milieu de propagation. L'étape suivante consiste ensuite à déterminer les chemins de propagation principaux entre les sources et les récepteurs, puis à calculer toutes les atténuations associées (e.g. la divergence géométrique et l'absorption atmosphérique). L'approche hybride vient modifier l'approche classique, d'une part, en ajoutant une information additionnelle au moment de la spécification de la scène, pour signaler l'existence d'un ou plusieurs objets complexes, et d'autre part, en ajoutant un terme d'atténuation

supplémentaire (ou correction) associé à chacun de ces objets. Cette atténuation supplémentaire, calculée avec la méthode BEM, correspond à l'atténuation d'un objet complexe (i.e. qui ne peut pas être modélisé par l'approche d'ingénierie) moins l'atténuation apportée par un objet simplifié (i.e. qui peut être modélisé par l'approche d'ingénierie). Finalement, la méthode hybride, comme la méthode d'ingénierie, somme toutes les contributions associées aux différents chemins de propagation afin d'obtenir le champ total aux récepteurs.

Les calculs BEM étant trop onéreux pour être réalisés pour chaque chemin de propagation, ils sont effectués durant une étape préliminaire et permettent de constituer une table de correction. De nombreux calculs BEM de référence sont ainsi effectués pour plusieurs positions de sources/récepteurs et pour différentes fréquences, à la fois pour l'objet complexe et son équivalent simplifié. Dans la méthode hybride, les positions sont décrites au moyen d'un repère cartésien à cinq coordonnées relatives à l'objet complexe (i.e. les distances source et récepteur à l'objet, l'élévation des sources et des récepteurs et la distance relative entre sources et récepteur parallèlement à l'objet). Par soucis d'efficacité, les positions sources/récepteurs sont contraintes à une grille non uniforme, ce qui permet de retrouver facilement un point dans la grille. Les calculs BEM en bandes fines sont ensuite convertis par bande d'octave ou de tiers d'octave, afin de calculer la correction, en soustrayant les valeurs obtenues avec l'objet simple à celles obtenues avec l'objet complexe. Finalement, les facteurs de corrections sont triés et tabulés afin d'offrir un accès rapide à ses valeurs en fonction des besoins. Cette table est ensuite chargée en mémoire quand la scène est mise en place.

Lors des calculs d'atténuation pour chaque chemin de propagation, la méthode hybride commence par vérifier si ce chemin est diffracté au-dessus de l'objet complexe. Si cela n'est pas le cas, l'atténuation excédentaire est considérée comme nulle. Si cette diffraction existe, la méthode calcule la différence de chemin entre chaque objet diffractant rencontré. Si l'objet complexe ne possède pas la plus grande différence de chemins, alors l'atténuation excédentaire est nulle. Si, par contre, cette différence est la plus grande, alors l'approche hybride détermine les positions sources/récepteurs associés au chemin de propagation. En pratique, ses coordonnées ont toutefois peu de chance de correspondre exactement aux valeurs disponibles dans la table. Par conséquent, si un point est à l'intérieur de l'enveloppe convexe du jeu de données, alors les données sont linéairement interpolées en se servant des points de l'hyper-rectangle le plus petit qui contient le point d'interpolation. On notera que l'emploi de points placés sur une grille facilite grandement la recherche des hyper-rectangles. Si le point est en dehors de l'enveloppe convexe, alors les données sont extrapolées. Si de nombreux points sont en dehors des valeurs disponibles, il est toutefois plus pertinent d'étendre l'amplitude des variables d'entrée utilisées pour la construction de la table de correction.

Afin d'effectuer un premier test de validation, l'approche hybride a été appliquée au cas « complexe » d'un écran en T de hauteur de 3 m, d'épaisseur 20 cm et de sommet plat de 1 m, reposant sur un sol rigide. La forme simplifiée associée (i.e. qui peut être modélisée par l'approche d'ingénierie) est un écran droit de même hauteur et même épaisseur. Dans un premier temps, le calcul BEM 2.5D a été employé pour modéliser la scène pour une source localisée à 30 cm du sol et 3 m en amont de l'écran, et pour des récepteurs placés sur une grille horizontale à 1.5 m du sol. L'approche hybride, appliquée à l'écran en T, réduit l'erreur de 1.5 dB(A) comparé à l'approche d'ingénierie, la source d'erreur principale étant liée à la différence de modélisation entre l'approche d'ingénierie et le calcul BEM 2.5D, déjà visible dans le cas de l'écran droit. Dans un second temps, la méthode hybride a été appliquée sur la même scène, mais cette fois en présence d'un sol herbeux, tout en utilisant la même table de correction (i.e. celle obtenue avec un sol rigide). Pour l'écran en T, l'approche hybride réduit l'erreur RMS de 1.1 dB(A) comparativement à l'approche d'ingénierie utilisée pour un écran droit. Ce cas montre qu'une table obtenue pour un sol rigide peut raisonnablement être employée pour d'autres types de sol.

Par la suite, l'approche hybride a été utilisée sur une nouvelle scène, incluant l'écran en T précédant, mais en présence de deux bâtiments. La table de correction est identique à celle utilisée précédemment (sur la base de la BEM 2.5D), mais cette fois, la méthode de référence pour la validation de l'approche hybride est basée sur une approche 3D. L'emploi d'un code BEM 3D standard étant trop coûteux au-delà de 400 Hz, nous avons utilisé la méthode FMBEM (*fast-multipole boundary element method*), ce qui permet d'atteindre des fréquences de validation plus élevées (1600 Hz). Pour l'écran en T, l'approche hybride réduit l'erreur RMS de 2.0 dB(A) comparativement à l'approche d'ingénierie utilisée pour un écran droit. Ce cas montre l'avantage de l'approche hybride lorsque la scène comporte également d'autres objets que l'objet complexe.

Afin d'illustrer ce concept, l'approche hybride est finalement appliquée avec le même écran en T dans une scène plus large (i.e. 180 m x 80 m), comportant plusieurs bâtiments et un sol non rigide. Cet exemple illustre une application pratique où l'approche hybride fournit des résultats réalistes, alors que l'approche d'ingénierie donnerait des résultats peu précis et que la méthode de référence (FMBEM) serait limitée en fréquence. L'approche hybride peut ainsi facilement modéliser la propagation extérieure pour une large plage de fréquences, pour des scènes plus grandes et avec de nombreux bâtiments.

L'utilisation de l'approche hybride soulève un problème d'optimisation lorsque l'on cherche à minimiser l'erreur d'interpolation pour un nombre donné de points dans la table des corrections. La distribution de points affecte l'erreur d'interpolation et il est permis de penser qu'une optimisation de ces positions permettrait de réduire l'erreur d'interpolation. Sans optimisation, l'approche hybride peut employer une distribution linéaire ou bien exponentielle des positions. Selon chaque dimension, une distribution

linéaire répartit les points régulièrement et une distribution exponentielle concentre les points dans une zone. Puisque la correction varie plus rapidement au voisinage de l'objet complexe et à proximité du sol, une distribution exponentielle est meilleure qu'une distribution linéaire, mais il reste de nombreuses alternatives non explorées. En présence de plusieurs degrés de liberté, l'erreur d'interpolation doit diminuer et le problème posé est alors de déterminer les paramètres à optimiser.

L'approche la plus directe est de choisir les valeurs qui minimisent l'erreur d'interpolation. Cependant, cette option est beaucoup trop onéreuse du fait du nombre de données nécessaires à l'évaluation du jeu de données de référence. Plutôt que de minimiser directement l'erreur d'interpolation, la méthode hybride minimise la variance de la valeur interpolée en employant un processus Gaussien. Ceci est une façon standard d'interpoler des données en minimisant directement l'erreur d'interpolation. L'erreur d'interpolation devrait décroître lorsque la variance des valeurs interpolées décroît. Là encore l'approche hybride utilise l'approche indirecte pour minimiser l'erreur car l'approche directe demanderait trop de calcul BEM.

L'emploi d'un processus Gaussien réduit la quantité de données nécessaires, mais cette quantité n'est pas nulle. Même si le calcul de la variance ne dépend pas directement de données de référence, la variance dépend d'hyper-paramètres qui sont, dans le meilleur des cas, approchés à partir de données de calcul BEM. Après avoir déterminé ces hyper-paramètres par optimisation de la vraisemblance log-marginale, l'utilisation d'un processus Gaussien permet d'obtenir la variance en fonction des positions des données et non des valeurs en ces positions. Ainsi, une répétition de l'évaluation de la fonction « objectif » ne demande pas d'avantage de calculs BEM.

Ce processus d'optimisation entraîne déjà une légère amélioration des résultats s'il n'y a qu'un paramètre par dimension (i.e. cinq paramètres en tout), mais cela pourrait probablement donner de meilleurs résultats avec un modèle plus flexible. De plus, l'optimisation de la distribution aboutit à une façon plus rigoureuse de choisir le nombre de points pour chaque dimension ainsi que la distribution de ces points. Il est préférable de connaître cette distribution plutôt qu'une distribution voisine même si l'erreur d'interpolation ne diminue pas de façon drastique. Cependant, cette distribution n'est pas forcément optimale puisque le problème d'optimisation n'est pas convexe.

En conclusion, l'approche hybride proposée s'avère être un bon compromis entre précision et efficacité, pour modéliser la propagation sonore en milieu extérieur et en présence de géométries complexes. Cette approche est plus précise que la méthode d'ingénierie dans des situations complexes et plus efficace, en termes de coût numérique, que les méthodes de référence. Une fois la table de correction créée, l'augmentation des temps de calcul en utilisant la méthode hybride est uniquement liée au chargement de la

table en mémoire, ce qui demande moins d'une minute dans le cadre de nos cas tests. Tous les calculs supplémentaires sont négligeables comparativement au reste des calculs.

Bibliography

1. European Environment Agency, “Population exposure to environmental noise” (2017).
2. B. Berglund, T. Lindvall, and D. H. Schwela, “Guidelines for community noise” (1999).
3. J. Lekaviciute et al., “Final Report ENNAH – European Network on Noise and Health” (2013) [doi:10.2788/83694].
4. L. Fritschi et al., “Burden of disease from environmental noise: quantification of healthy life years lost in Europe” (2011).
5. S. Kephelopoulos, M. Paviotti, and F. Anfosso-Lédée, “Common noise assessment methods in Europe (CNOSSOS-EU)” (2012) [doi:10.2788/31776].
6. T. Ishizuka and K. Fujiwara, “Performance of noise barriers with various edge shapes and acoustical conditions,” *Appl. Acoust.* **65**(2), 125–141 (2004) [doi:10.1016/j.apacoust.2003.08.006].
7. C. H. Kasess, W. Kreuzer, and H. Waubke, “Deriving correction functions to model the efficiency of noise barriers with complex shapes using boundary element simulations,” *Appl. Acoust.* **102**, 88–99 (2016) [doi:10.1016/j.apacoust.2015.09.009].
8. M. Baulac, J. Defrance, and P. Jean, “Optimisation with genetic algorithm of the acoustic performance of T-shaped noise barriers with a reactive top surface,” *Appl. Acoust.* **69**(4), 332–342 (2008) [doi:10.1016/j.apacoust.2006.11.002].
9. A. Jolibois et al., “Sensitivity-based shape optimization of a rigid tramway low-height noise barrier,” in *Inter-noise*, pp. 3492–3501, Innsbruck, Austria (2013).
10. C. H. Kasess et al., “Effects of source type, position, and train structure on BEM calculations,” in *Inter-noise*, pp. 4387–4396, Hamburg, Germany (2016).

11. T. Van Renterghem and D. Botteldooren, “The importance of roof shape for road traffic noise shielding in the urban environment,” *J. Sound Vib.* **329**(9), 1422–1434 (2010) [doi:10.1016/j.jsv.2009.11.011].
12. T. Van Renterghem et al., “The potential of building envelope greening to achieve quietness,” *Build. Environ.* **61**, 34–44 (2013) [doi:10.1016/j.buildenv.2012.12.001].
13. G. B. Jónsson and F. Jacobsen, “A comparison of two engineering models for outdoor sound propagation: Harmonoise and Nord2000,” *Acta Acust. united with Acust.* **94**(2), 282–289 (2008) [doi:10.3813/AAA.918031].
14. G. B. Jónsson, “A comparison of two numerical models for outdoor sound propagation: Harmonoise and Nord2000,” Technical University of Denmark (2007).
15. B. Plovsing, “Nord2000. Comprehensive outdoor sound propagation model. part 1: propagation in an atmosphere without significant refraction” (2006).
16. B. Plovsing, “Nord2000. Comprehensive outdoor sound propagation model. part 2: propagation in an atmosphere with refraction” (2006).
17. D. van Maercke, T. Leissing, and J. Maillard, “Hosanna task 6.2: evaluation of innovative and combined mitigations” (2013).
18. J. B. Allen and D. A. Berkley, “Image method for efficiently simulating small-room acoustics,” *J. Acoust. Soc. Am.* **65**(4), 943–950 (1979) [doi:10.1121/1.382599].
19. K. M. Li and C. Y. C. Lai, “A note on noise propagation in street canyons,” *J. Acoust. Soc. Am.* **126**(2), 644–655 (2009) [doi:10.1121/1.3158599].
20. L. M. Kan, “Prediction, measurement and reduction of traffic noise,” The Hong Kong Polytechnic University (2008).
21. J. Kang, “Sound propagation in street canyons: Comparison between diffusely and geometrically reflecting boundaries,” *J. Acoust. Soc. Am.* **107**(3), 1394–1404 (2000) [doi:10.1121/1.428580].
22. U. R. Kristiansen, A. Krokstad, and T. Follestad, “Extending the image method to higher-order reflections,” *Appl. Acoust.* **38**, 195–206 (1993) [doi:10.1016/0003-682X(93)90051-7].
23. H. Kuttruff, *Room Acoustics*, 5th ed., Spon Press (2009) [doi:10.4324/9780203186237].
24. M. Pharr and G. Humphreys, *Physically Based Rendering*, Elsevier (2010) [doi:10.1016/B978-0-12-375079-2.50001-9].
25. J. S. Bradley, “A study of traffic noise attenuation around buildings,” *Acta Acust.*

- united with *Acust.* **38**(4), 247–252 (1977).
26. D. P. Hewett, “High frequency sound propagation in a network of interconnecting streets,” *J. Sound Vib.* **331**(25), 5537–5561 (2012) [doi:10.1016/j.jsv.2012.07.030].
 27. D. P. Hewett, “Sound propagation in an urban environment,” University of Oxford (2010).
 28. A. Muradali and K. R. Fyfe, “Accurate barrier modeling in the presence of atmospheric effects,” *Appl. Acoust.* **56**(3), 157–182 (1999) [doi:10.1016/S0003-682X(98)00023-1].
 29. Y. W. Lam, “An analytical model for turbulence scattered rays in the shadow zone for outdoor sound propagation calculation,” *J. Acoust. Soc. Am.* **125**(3), 1340–1350 (2009) [doi:10.1121/1.3076928].
 30. A. Billon and J. Embrechts, “A diffraction model for acoustical ray-tracing based on the energy flow lines concept,” *Acta Acust. united with Acust.* **99**(2), 260–267 (2013) [doi:10.3813/AAA.918608].
 31. T. Funkhouser et al., “A beam tracing method for interactive architectural acoustics,” *J. Acoust. Soc. Am.* **115**(2), 739–756 (2004) [doi:10.1121/1.1641020].
 32. A. Farina, “Validation of the pyramid tracing algorithm for sound propagation outdoors: comparison with experimental measurements and with the ISO-DIS 9613 standards,” *Adv. Eng. Softw.* **31**(4), 241–250 (2000) [doi:10.1016/S0965-9978(99)00053-8].
 33. H. Wang, Z. Yu, and M. Cai, “The 3D attenuation calculation of traffic noise among building groups by using beam tracing method,” *Procedia - Soc. Behav. Sci.* **96**, 1929–1937 (2013) [doi:10.1016/j.sbspro.2013.08.218].
 34. J. Keller, “Geometrical theory of diffraction,” *J. Opt. Soc. Am.* **52**(2), 116–130 (1962) [doi:10.1364/JOSA.52.000116].
 35. R. G. Kouyoumjian and P. H. Pathak, “A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface,” *Proc. IEEE* **62**(11), 1448–1461 (1974) [doi:10.1109/PROC.1974.9651].
 36. A. D. Pierce, “Diffraction of sound around corners and over wide barriers,” *J. Acoust. Soc. Am.* **55**(5), 941–955 (1974) [doi:10.1121/1.1914668].
 37. T. Kawai, “Sound diffraction by a many-sided barrier or pillar,” *J. Sound Vib.* **79**(2), 229–242 (1981) [doi:10.1016/0022-460X(81)90370-9].
 38. B. Jin et al., “Sound diffraction by a partially inclined noise barrier,” *Appl. Acoust.* **62**(9), 1107–1121 (2001) [doi:10.1016/S0003-682X(00)00094-3].
 39. H. Kim et al., “Sound diffraction by multiple wedges and thin screens,” *Appl.*

- Acoust. **66**(9), 1102–1119 (2005) [doi:10.1016/j.apacoust.2005.01.004].
40. N. Garg and S. Maji, “A critical review of principal traffic noise models: strategies and implications,” *Environ. Impact Assess. Rev.* **46**, 68–81 (2014) [doi:10.1016/j.eiar.2014.02.001].
 41. “ISO 9613-2:1996 Acoustics - attenuation of sound during propagation outdoors - part 2: general method of calculation” (1996).
 42. Sétra, “Road noise prediction: Part 1 - Calculating sound emissions from road traffic” (2009).
 43. Sétra, “Road noise prediction: Part 2 - Noise propagation computation method including meteorological effects (NMPB 2008)” (2009).
 44. D. van Maercke and J. Defrance, “Development of an analytical model for outdoor sound propagation within the Harmonoise project,” *Acta Acust. united with Acust.* **93**(2), 201–212 (2007).
 45. E. Salomons et al., “The Harmonoise sound propagation model,” *Acta Acust. united with Acust.* **97**(1), 62–74 (2011) [doi:10.3813/AAA.918387].
 46. J. Kang, “Sound propagation in interconnected urban streets: a parametric study,” *Environ. Plan. B Plan. Des.* **28**(2), 281–294 (2001) [doi:10.1068/b2680].
 47. J. Kang, “Numerical modelling of the sound fields in urban streets with diffusely reflecting boundaries,” *J. Sound Vib.* **258**(5), 793–813 (2002) [doi:10.1006/jsvi.2002.5150].
 48. J. Kang, “Numerical modeling of the sound fields in urban squares,” *J. Acoust. Soc. Am.* **117**(6), 3695–3706 (2005) [doi:10.1121/1.1904483].
 49. D. A. Naish, A. C. C. Tan, and F. N. Demirbilek, “Speech interference and transmission on residential balconies with road traffic noise,” *J. Acoust. Soc. Am.* **133**(1), 210–226 (2013) [doi:10.1121/1.4765075].
 50. E. Reboul, A. Le Bot, and J. Perret-Liaudet, “Radiative transfer equation for multiple diffraction,” *J. Acoust. Soc. Am.* **118**(3), 1326–1334 (2005) [doi:10.1121/1.2001467].
 51. J. Picaut, L. Simon, and J. Polack, “A mathematical model of diffuse sound field based on a diffusion equation,” *Acta Acust. united with Acust.* **83**(4), 614–621 (1997).
 52. J. Picaut, L. Simon, and J. Hardy, “Sound field modeling in streets with a diffusion equation,” *J. Acoust. Soc. Am.* **106**(5), 2638–2645 (1999) [doi:10.1121/1.428093].
 53. J. Picaut, J. Hardy, and L. Simon, “Sound propagation in urban areas: a periodic disposition of buildings,” *Phys. Rev. E* **60**(4), 4851–4859 (1999) [doi:10.1103/PhysRevE.60.4851].

54. T. Le Pollès et al., “Sound-field modeling in architectural acoustics by a transport theory: application to street canyons,” *Phys. Rev. E* **72**(4), 1–17 (2005) [doi:10.1103/PhysRevE.72.046609].
55. J. Picaut, “Numerical modeling of urban sound fields by a diffusion process,” *Appl. Acoust.* **63**(9), 965–991 (2002) [doi:10.1016/S0003-682X(02)00011-7].
56. S. M. Pasareanu, M. C. Remillieux, and R. A. Burdisso, “Energy-based method for near real-time modeling of sound field in complex urban environments,” *J. Acoust. Soc. Am.* **132**(6), 3647–3658 (2012) [doi:10.1121/1.4763552].
57. G. James et al., *An Introduction to Statistical Learning with Applications in R*, Springer (2014) [doi:10.1007/978-1-4614-7138-7].
58. T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer (2009) [doi:10.1007/b94608].
59. C. R. Hart et al., “Comparisons between physics-based, engineering, and statistical learning models for outdoor sound propagation,” *J. Acoust. Soc. Am.* **139**(5), 2640–2655 (2016) [doi:10.1121/1.4948757].
60. W. Desmet and D. Vandepitte, “Finite Element Modeling for Acoustics,” in *ISAAC13-International Seminar on Applied Acoustics*, Leuven, pp. 37–85 (2002).
61. ANSYS, “Theory Reference for the Mechanical APDL and Mechanical Applications” (2009).
62. I. Harari, “A survey of finite element methods for time-harmonic acoustics,” *Comput. Methods Appl. Mech. Eng.* **195**(13–16), 1594–1607 (2006) [doi:10.1016/j.cma.2005.05.030].
63. L. L. Thompson, “A review of finite-element methods for time-harmonic acoustics,” *J. Acoust. Soc. Am.* **119**(3), 1315–1330 (2006) [doi:10.1121/1.2164987].
64. O. C. Zienkiewicz, “Achievements and some unsolved problems of the finite element method,” *Int. J. Numer. Methods Eng.* **47**(1–3), 9–28 (2000) [doi:10.1002/(SICI)1097-0207(20000110/30)47:1/3<9::AID-NME793>3.0.CO;2-P].
65. J. Berenger, “A perfectly matched layer for the absorption of electromagnetic waves,” *J. Comput. Phys.* **114**(2), 185–200 (1994) [doi:10.1006/jcph.1994.1159].
66. K. Gerdes, “A review of infinite element methods for exterior helmholtz problems,” *J. Comput. Acoust.* **8**(1), 43–62 (2000) [doi:10.1142/S0218396X00000042].
67. G. Beer, I. Smith, and C. Duenser, *The Boundary Element Method with Programming for Engineers and Scientists*, SpringerWienNewYork (2008)

- [doi:10.1007/978-3-211-71576-5].
68. C. Pozrikidis, *A Practical Guide to Boundary Element Methods with the Software Library BEMLIB*, Chapman & Hall/CRC (2002).
 69. S. Kirkup, *The Boundary Element Method in Acoustics*, 2nd ed., Integrated Sound Software (2007).
 70. D. C. Hothersall, S. N. Chandler-Wilde, and M. N. Hajmirzae, “Efficiency of single noise barriers,” *J. Sound Vib.* **146**(2), 303–322 (1991) [doi:10.1016/0022-460X(91)90765-C].
 71. M. Baulac et al., “Efficiency of noise protections in urban areas: predictions and scale model measurements,” *Acta Acust. united with Acust.* **92**(4), 530–539 (2006).
 72. D. C. Hothersall, K. V. Horoshenkov, and S. E. Mercy, “Numerical modelling of the sound field near a tall building with balconies near a road,” *J. Sound Vib.* **198**(4), 507–515 (1996) [doi:10.1006/jsvi.1996.0584].
 73. T. Sakuma and Y. Yasuda, “Fast multipole boundary element method for large-scale steady-state sound field analysis. part I: setup and validation,” *Acta Acust. united with Acust.* **88**(4), 513–525 (2002).
 74. M. Fischer and L. Gaul, “Application of the fast multipole BEM for structural-acoustic simulations,” *J. Comput. Acoust.* **13**(1), 87–98 (2005) [doi:10.1142/S0218396X05002578].
 75. R. Gunda, “Boundary element acoustics and the fast multipole method (FMM),” *Sound Vib.* **March**, 12–16 (2008).
 76. J. Huijssen et al., “Numerical evaluation of source–receiver transfer functions with the Fast Multipole Boundary Element Method for predicting pass-by noise levels of automotive vehicles,” *J. Sound Vib.* **331**(9), 2080–2096 (2012) [doi:10.1016/j.jsv.2011.11.030].
 77. X. Vuylsteke, “Development of a reference method based on the fast multipole boundary element method for sound propagation problems in urban environments: formalism, optimizations & applications,” Université Paris-Est (2014).
 78. M. J. White and K. E. Gilbert, “Application of the parabolic equation to the outdoor propagation of sound,” *Appl. Acoust.* **27**(3), 227–238 (1989) [doi:10.1016/0003-682X(89)90062-5].
 79. D. Lee, A. D. Pierce, and E. Shang, “Parabolic equation development in the twentieth century,” *J. Comput. Acoust.* **8**(4), 527–637 (2000) [doi:10.1142/S0218396X00000388].
 80. T. Van Renterghem, D. Botteldooren, and P. Lercher, “Comparison of measurements and predictions of sound propagation in a valley-slope

- configuration in an inhomogeneous atmosphere,” *J. Acoust. Soc. Am.* **121**(5), 2522–2533 (2007) [doi:10.1121/1.2717765].
81. E. M. Salomons, “Noise barriers in a refracting atmosphere,” *Appl. Acoust.* **47**(3), 217–238 (1996) [doi:10.1016/0003-682X(95)00047-D].
82. N. Barriere and Y. Gabillet, “Sound propagation over a barrier with realistic wind gradients. Comparison of wind tunnel experiments with GFPE computations,” *Acust. united with Acta Acust.* **85**(3), 325–335 (1999).
83. E. M. Salomons, “Reduction of the performance of a noise screen due to screen-induced wind-speed gradients. Numerical computations and wind-tunnel experiments,” *J. Acoust. Soc. Am.* **105**(4), 2287–2293 (1999) [doi:10.1121/1.426835].
84. J. Forssén, “Calculation of sound reduction by a screen in a turbulent atmosphere using the parabolic equation method,” *Acust. united with Acta Acust.* **84**(4), 599–606 (1998).
85. J. Forssén, “Calculation of noise barrier performance in a turbulent atmosphere by using substitute sources above the barrier,” *Acta Acust. united with Acust.* **86**(2), 269–275 (2000).
86. F. Aballéa and J. Defrance, “Single and multiple reflections in plane obstacle using the parabolic equation method with a complementary Kirchhoff approximation,” *Acta Acust. united with Acust.* **93**(1), 22–30 (2007).
87. M. Hornikx and J. Forssén, “The 2.5-dimensional equivalent sources method for directly exposed and shielded urban canyons,” *J. Acoust. Soc. Am.* **122**(5), 2532–2541 (2007) [doi:10.1121/1.2783197].
88. M. Hornikx, “Numerical modelling of sound propagation to closed urban courtyards,” Chalmers University of Technology (2009).
89. M. Ögren and W. Kropp, “Road traffic noise propagation between two dimensional city canyons using an equivalent sources approach,” *Acta Acust. united with Acust.* **90**(2), 293–300 (2004).
90. M. Hornikx and J. Forssén, “Noise abatement schemes for shielded canyons,” *Appl. Acoust.* **70**(2), 267–283 (2009) [doi:10.1016/j.apacoust.2008.04.002].
91. R. Mehra et al., “Wave-based sound propagation in large open scenes using an equivalent source formulation,” *ACM Trans. Graph.* **32**(2), 1–13 (2013) [doi:10.1145/2451236.2451245].
92. T. Van Renterghem, “Efficient outdoor sound propagation modeling with the finite-difference time-domain (FDTD) method: a review,” *Int. J. Aeroacoustics* **13**(5–6), 385–404 (2014) [doi:10.1260/1475-472X.13.5-6.385].
93. D. Heimann, “Three-dimensional linearised Euler model simulations of sound

- propagation in idealised urban situations with wind effects,” *Appl. Acoust.* **68**(2), 217–237 (2007) [doi:10.1016/j.apacoust.2005.10.002].
94. D. G. Albert and L. Liu, “The effect of buildings on acoustic pulse propagation in an urban environment,” *J. Acoust. Soc. Am.* **127**(3), 1335–1346 (2010) [doi:10.1121/1.3277245].
95. R. Mehra et al., “Acoustic pulse propagation in an urban environment using a three-dimensional numerical simulation,” *J. Acoust. Soc. Am.* **135**(6), 3231–3242 (2014) [doi:10.1121/1.4874495].
96. T. Oshima, T. Ishizuka, and T. Kamijo, “Three-dimensional urban acoustic simulations and scale-model measurements over real-life topography,” *J. Acoust. Soc. Am.* **135**(6), EL324-EL330 (2014) [doi:10.1121/1.4879672].
97. C. Christopoulos, *The Transmission-Line Modeling Method: TLM*, IEEE, Inc. (1995).
98. J. Hofmann and K. Heutschi, “Simulation of outdoor sound propagation with a transmission line matrix method,” *Appl. Acoust.* **68**(2), 158–172 (2007) [doi:10.1016/j.apacoust.2005.10.006].
99. G. Guillaume et al., “Time-domain impedance formulation for transmission line matrix modelling of outdoor sound propagation,” *J. Sound Vib.* **330**(26), 6467–6481 (2011) [doi:10.1016/j.jsv.2011.08.004].
100. G. Guillaume et al., “Application of the transmission line matrix method for outdoor sound propagation modelling – part 1: model presentation and evaluation,” *Appl. Acoust.* **76**, 113–118 (2014) [doi:10.1016/j.apacoust.2013.07.011].
101. P. Aumond et al., “Application of the Transmission Line Matrix method for outdoor sound propagation modelling – part 2: experimental validation using meteorological data derived from the meso-scale model Meso-NH,” *Appl. Acoust.* **76**, 107–112 (2014) [doi:10.1016/j.apacoust.2013.07.015].
102. T. Van Renterghem, E. M. Salomons, and D. Botteldooren, “Efficient FDTD-PE model for sound propagation in situations with complex obstacles and wind profiles,” *Acta Acust. united with Acust.* **91**(4), 671–679 (2005).
103. T. Van Renterghem, E. Salomons, and D. Botteldooren, “Parameter study of sound propagation between city canyons with a coupled FDTD-PE model,” *Appl. Acoust.* **67**(6), 487–510 (2006) [doi:10.1016/j.apacoust.2005.09.006].
104. T. Van Renterghem and D. Botteldooren, “Numerical evaluation of tree canopy shape near noise barriers to improve downwind shielding,” *J. Acoust. Soc. Am.* **123**(2), 648–657 (2008) [doi:10.1121/1.2828052].
105. E. Premat et al., “Coupling BEM and GFPE for complex outdoor sound propagation,” in *Euronoise*, pp. 1–6, Naples, Italy (2003).

106. J. Defrance et al., “Outdoor sound propagation reference model developed in the European Harmonoise project,” *Acta Acust. united with Acust.* **93**(2), 213–227 (2007).
107. S. Langer and S. Hampel, “Coupling boundary elements to a raytracing procedure using the singular indirect method,” *Proc. Appl. Math. Mech.* **5**, 609–610 (2005) [doi:10.1002/pamm.200510281].
108. S. Hampel, S. Langer, and A. P. Cisilino, “Coupling boundary elements to a raytracing procedure,” *Int. J. Numer. Methods Eng.* **73**(3), 427–445 (2008) [doi:10.1002/nme.2080].
109. P. Jean, “Coupling integral and geometrical representations for vibro-acoustical problems,” *J. Sound Vib.* **224**(3), 475–487 (1999) [doi:10.1006/jsvi.1999.2195].
110. P. Jean, “Coupling geometrical and integral methods for indoor and outdoor sound propagation—validation examples,” *Acta Acust. united with Acust.* **87**(2), 236–246 (2001).
111. H. Yeh et al., “Wave-ray coupling for interactive sound propagation in large complex scenes,” *ACM Trans. Graph.* **32**(6), 1–11 (2013) [doi:10.1145/2508363.2508420].
112. J. Defrance and P. Jean, “Integration of the efficiency of noise barrier caps in a 3D ray tracing method. case of a T-shaped diffracting device,” *Appl. Acoust.* **64**(8), 765–780 (2003) [doi:10.1016/S0003-682X(03)00034-3].
113. Purdue, “PC Desktop Recommendations,” Purdue University—Engineering Computer Network, <<https://engineering.purdue.edu/ECN/Support/KB/Docs/PCDesktop>> (accessed 5 April 2017).
114. C. Foy, E. David, and G. Dutilleux, “Comparison of 3 engineering methods for outdoor sound propagation,” in *DAGA*, pp. 828–829, Oldenburg, Germany (2014).
115. *Directive 2002/49/EC of the European parliament and the council of 25 June 2002 relating to the assessment and management of environmental noise* (2002).
116. P. Jean, J. Defrance, and Y. Gabillet, “The importance of source type on the assessment of noise barriers,” *J. Sound Vib.* **226**(2), 201–216 (1999) [doi:10.1006/jsvi.1999.2273].
117. P. Jean, “A variational approach for the study of outdoor sound propagation and application to railway noise,” *J. Sound Vib.* **212**(2), 275–294 (1998) [doi:10.1006/jsvi.1997.1407].
118. D. Duhamel, “Efficient calculation of the three-dimensional sound pressure field around a noise barrier,” *J. Sound Vib.* **197**(5), 547–571 (1996) [doi:10.1006/jsvi.1996.0548].

119. ISO, “ISO 9613-1:1993 Acoustics—attenuation of sound during propagation outdoors—part 1: calculation of the absorption of sound by the atmosphere” (1993).
120. J. Defrance et al., “Hosanna task 2.3: application to innovations” (2013).
121. C. B. Barber, D. P. Dobkin, and H. Huhdanpaa, “The quickhull algorithm for the convex hulls,” *ACM Trans. Math. Softw.* **22**(4), 469–483 (1996) [doi:10.1145/235815.235821].
122. C. B. Barber, D. P. Dobkin, and H. T. Huhdanpaa, “Qhull,” <<http://qhull.org/>> (accessed 1 December 2016).
123. C. Yang et al., “Twelve different interpolation methods : a case study of Surfer 8.0,” in XXth ISPRS Congress, Commission II, pp. 778–785, Istanbul, Turkey (2004).
124. G. Y. Lu and D. W. Wong, “An adaptive inverse-distance weighting spatial interpolation technique,” *Comput. Geosci.* **34**(9), 1044–1055 (2008) [doi:10.1016/j.cageo.2007.07.010].
125. T. Bobach et al., “Natural neighbor extrapolation using ghost points,” *Comput. Des.* **41**(5), 350–365 (2009) [doi:10.1016/j.cad.2008.08.007].
126. C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, MIT Press (2006).
127. M. C. Mariani and K. Basu, “Spline interpolation techniques applied to the study of geophysical data,” *Phys. A Stat. Mech. its Appl.* **428**, 68–79 (2015) [doi:10.1016/j.physa.2015.02.014].
128. A. Donini, “Advanced turbulent combustion modeling for gas turbine application,” Technische Universiteit Eindhoven (2014) [doi:10.6100/IR773140].
129. M. E. Delany and E. N. Bazley, “Acoustical properties of fibrous absorbent materials,” *Appl. Acoust.* **3**(2), 105–116 (1970) [doi:10.1016/0003-682X(70)90031-9].
130. “le cnam - structural mechanics and couple systems laboratory,” <<http://www.lmssc.cnam.fr/en>> (accessed 13 July 2017).
131. J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., Springer (2006) [doi:10.1007/978-0-387-40065-5].
132. MathWorks, “fminunc,” 2017, <www.mathworks.com/help/optim/ug/fminunc.html> (accessed 1 January 2017).
133. T. F. Coleman and Y. Li, “An interior trust region approach for nonlinear minimization subject to bounds,” *SIAM J. Optim.* **6**(2), 418–445 (1996) [doi:10.1137/0806023].
134. C. G. Broyden, “The convergence of a class of double-rank minimization

- algorithms 1. general considerations,” *IMA J. Appl. Math.* **6**(1), 76–90 (1970) [doi:10.1093/imamat/6.1.76].
135. R. Fletcher, “A new approach to variable metric algorithms,” *Comput. J.* **13**(3), 317–322 (1970) [doi:10.1093/comjnl/13.3.317].
136. D. Goldfarb, “A family of variable metric updates derived by variational means,” *Math. Comput.* **24**(109), 23–26 (1970) [doi:10.1090/S0025-5718-1970-0258249-6].
137. D. F. Shanno, “Conditioning of quasi-Newton methods for function minimization,” *Math. Comput.* **24**(111), 647–656 (1970) [doi:10.1090/S0025-5718-1970-0274029-X].

Publications

Articles

Kamrath, Jean, Maillard, & Picaut

Optimizing point locations by minimizing the predictive variance of a Gaussian process. Manuscript in preparation. (2017).

Kamrath, Jean, Maillard, Picaut, & Langrenne

Extending standard urban outdoor noise propagation models to complex geometries. Manuscript submitted. (2017).

Conferences

Kamrath, Jean, Langrenne & Picaut

Enabling noise engineering methods to model complex geometries

J. Acoust. Soc. Am. 141, 3879 (2017)

<http://doi.org/10.1121/1.4988682>

Kamrath, Picaut & Jean

Augmenting road noise engineering methods using the Boundary Element Method

JJ CAB – Marseille (2016)

<https://jjcab2016.sciencesconf.org/>

Kamrath, Maillard, Jean, Van Maercke & Picaut

Augmenting road noise engineering methods using the Boundary Element Method

Inter-noise – Hamburg (2016);

<http://pub.dega-akustik.de/IN2016/data/articles/000105.pdf>

Kamrath, Maillard, Jean, Van Maercke & Picaut

Modeling outdoor sound propagation in urban environments

J. Acoust. Soc. Am. 139, 1983 (2016)

<http://dx.doi.org/10.1121/1.4949788>

Kamrath, Maillard, Jean, Van Maercke & Picaut

The initial development of a hybrid method for modeling outdoor sound propagation in urban areas

J. Acoust. Soc. Am. 138, 1731 (2015)

<http://dx.doi.org/10.1121/1.4933447>

Thèse de Doctorat

Matthew KAMRATH

Extension des modèles standards de propagation du bruit extérieur pour des géométries complexes

Extending standard outdoor noise propagation models to complex geometries

Résumé

Les méthodes d'ingénierie acoustique (e.g. ISO 9613-2 ou CNOSSOS-EU) approchent efficacement les niveaux de bruit générés par les routes, les voies ferrées et les sources industrielles en milieu urbain. Cependant, ces approches d'ingénierie sont limitées à des géométries de forme simple, le plus souvent de section rectangulaire. Ce mémoire développe donc, et valide, une approche hybride permettant l'extension des méthodes d'ingénierie à des formes plus complexes, en introduisant un terme d'atténuation supplémentaire qui représente l'effet d'un objet réel comparé à un objet simple.

Le calcul de cette atténuation supplémentaire nécessite des calculs de référence, permettant de quantifier la différence entre objets simple et complexe. Dans la mesure, où il est trop onéreux, numériquement, d'effectuer ce calcul pour tous les chemins de propagation, l'atténuation supplémentaire est obtenue par interpolation de données stockées dans un tableau et évaluées pour un large jeu de positions de sources, de récepteurs et de fréquences. Dans notre approche, les calculs de référence utilisent la méthode BEM en 2.5D, et permet ainsi de produire les niveaux de référence pour les géométries simple et complexe, tout en tabulant leur écart. Sur le principe, d'autres approches de référence pourraient être utilisées.

Ce travail valide cette approche hybride pour un écran en forme de T avec un sol rigide, un sol absorbant et un cas avec bâtiments. Ces trois cas démontrent que l'approche hybride est plus précise que l'approche d'ingénierie standard dans des cas complexes.

Mots clés

Propagation sonore en milieu urbain, méthodes d'ingénierie acoustique, BEM, méthode hybride, processus Gaussien de régression, acoustique numérique, optimisation

Abstract

Noise engineering methods (e.g. ISO 9613-2 or CNOSSOS-EU) efficiently approximate sound levels from roads, railways, and industrial sources in cities. However, engineering methods are limited to only simple box-shaped geometries. This dissertation develops and validates a hybrid method to extend the engineering methods to more complicated geometries by introducing an extra attenuation term that represents the influence of a real object compared to a simplified object.

Calculating the extra attenuation term requires reference calculations to quantify the difference between the complex and simplified objects. Since performing a reference computation for each path is too computationally expensive, the extra attenuation term is linearly interpolated from a data table containing the corrections for many source and receiver positions and frequencies. The 2.5D boundary element method produces the levels for the real complex geometry and a simplified geometry, and subtracting these levels yields the corrections in the table.

This dissertation validates this hybrid method for a T-barrier with hard ground, soft ground, and buildings. All three cases demonstrate that the hybrid method is more accurate than standard engineering methods for complex cases.

Key Words

Urban outdoor sound propagation, noise engineering methods, boundary element method, hybrid method, Gaussian process regression, numerical acoustics, optimization