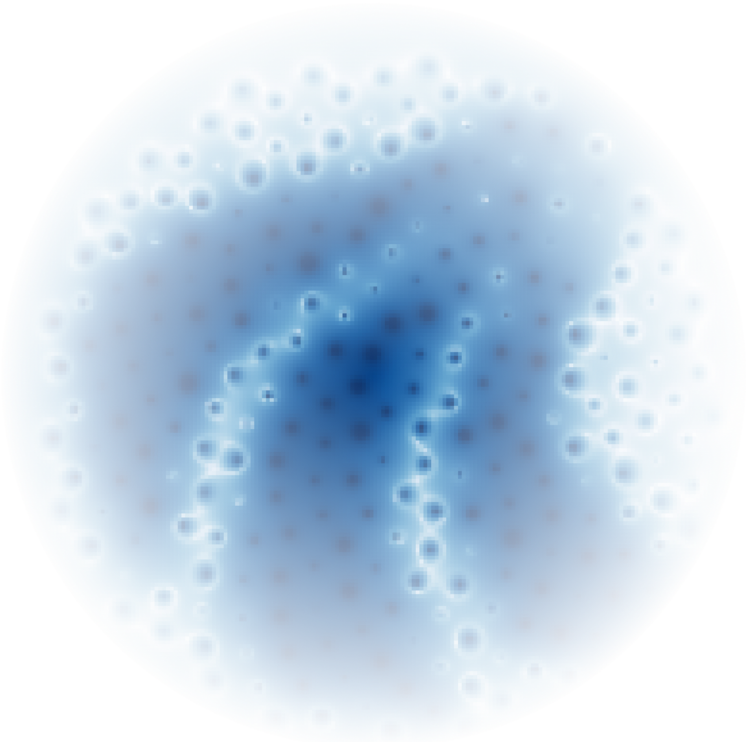

On Multiple Sound Zones for Wideband Signals

Asger Heidemann Andersen

Master's Thesis
M.Sc., Signal Processing and Computing
School of Information and Communication Technology
Spring 2014, Aalborg University

AALBORG UNIVERSITY
DENMARK

School of Information and Communication Technology
http://www.sict.aau.dk
esn@esn.aau.dk

Title:
On Multiple Sound Zones for Wideband Signals

Theme:
Digital Signal Processing

Project Period:
2014-02-01 to 2014-06-04

ECTS-points:
30

Project Group:
14gr1011

Group Members:
Asger Heidemann Andersen

Supervisors:
Jesper Kjær Nielsen, AAU
Tobias Lindstrøm Jensen, AAU
Torben Larsen, AAU

External Supervisor:
Martin Bo Møller, Bang & Olufsen

Number of pages:
74 (excluding appendix)

Attachments:
CD containing various relevant material

Abstract:

The concept of sound zones has recently gained much attention in the acoustics community. The idea of sound zones is to separate a space into multiple zones for which the audio content can be individually controlled. This is done by use of an array of loudspeakers. The most commonly studied objective is to play sound in one part of a space while simultaneously attempting to keep another part of the space quiet. There are several known methods for steering an array in order to obtain this result. An issue is that these methods typically assume the input to be sinusoidal with fixed frequency. This is not a realistic assumption for common audio signals such as speech or music. The problem of creating sound zones for such signals has not been extensively studied in the literature. In this thesis we investigate some possibilities for creating sound zones for wideband signals of this type. To facilitate the investigation of sound zone methods, a software framework for room simulation with the mirror image model has been developed. Among the discussed methods are both methods which have previously been discussed in the literature and methods which have not. We show that the task of extending sound zone methods to wideband signals introduces complications which are not seen for simple sinusoidal signals. We also discuss means of overcoming these complications. The main outcome of the presented investigations is a sound zone method which performs well for wideband signals.

Preface

The present thesis documents the work carried out during my master's project on Aalborg University. The project investigates the rather young research topic that is sound zone technology. The overall aim of the sound zone technology is to devise audio systems which are able to reproduce different audio channels at different locations in a space simultaneously. While this is an appealing concept it has has not been investigated much until recently. This may be explained in part by scientific and technological limitations, but may also in part be due to more straightforward factors. In the process of carrying out the project, I have discussed the idea of sound zones with many people, both engineers and non-engineers. My experience has been that most people from both groups receive the idea with severe disbelief. Our everyday experience with audio seems to contradict the possibility of separating a room into multiple independent sound zones. This psychological factor may have played a part in delaying research in the field. As is shown throughout in the present thesis, the concept of sound zones is not nearly as far fetched as our intuition seems to suggest. In the last decade, several interesting approaches for obtaining the desired result has been proposed. There is, however, still much work left to do before sound zone technology can be considered for consumer products. Almost all of the proposed methods suffer from one severe limitation; they consider sinusoidal input signals of a fixed frequency. The systems are therefore not very useful for typical audio signals such as speech or music. This project takes the character of an exploratory investigation of possibilities for extending sound zone technology to such signals.

The project is carried out as a mixture of analytical derivations and computer simulations. Due to the exploratory character of the project, as well as the limited time frame, it has been chosen not to carry out actual acoustical measurements. Instead, a framework for acoustical simulations has been developed and used. This is used to underline relevant theoretical points and to benchmark the discussed methods. Along with the figures of the thesis, a number of animated plots and audio demonstrations are also available. These are found on the accompanying CD along with a digital copy of the report as well as a

copy of all the necessary code to generate the supplied material.
The project has been carried out in collaboration with Bang & Olufsen which has kindly offered much help and support throughout the entire period. I would like to thank all of the supervisors who have been involved in the project, those from Aalborg University as well as Martin Bo Møller from Bang & Olufsen. All of them have provided invaluable and much appreciated feedback and ideas.
Lastly, a note should be attached to the image on the front page of the thesis. The image shows the logo of Aalborg University recreated as a simulated sound-field. This is produced by the Pressure Matching (PM) method, one of the methods discussed in the thesis. The code for generating the image is as well supplied on the CD.

Asger Heidemann Andersen
2014-06-03

Glossary

ACC	Acoustic Contrast Control
CG	Conjugate Gradient
CG-PM	Conjugate Gradient PM
DFT	Discrete Fourier Transform
DTFT	Discrete Time Fourier Transform
FFT	Fast Fourier Transformation
FIR	Finite Impulse Response
IDFT	Inverse DFT
MIM	Mirror Image Model
MF-ACC	Multi Frequency ACC
MF-PM	Multi Frequency PM
PM	Pressure Matching
SQAM	Sound Quality Assessment Material
TD-ACC	Time Domain ACC

Notation and Symbols

We provide a small reference of the notation used throughout the thesis.

Linear Algebra

Scalars, vectors and matrices are referenced with lowercase italic, lowercase bold and uppercase bold letters, respectively:

x	A scalar.
$\mathbf{x}$	A column-vector.
$\mathbf{X}$	A matrix.
$(\cdot)^{\mathrm{T}}$	Transpose.
$(\cdot)^{\mathrm{H}}$	Hermitian transpose.
$(\cdot)^{*}$	Complex conjugation (without transpose).
$\widetilde{(\cdot)}$	Block vector or matrix.
$\det(\cdot)$	Matrix determinant.

Indexing of vector- and matrix elements starts from 0 and we may choose from the following two types of notation depending on circumstances:

$$\begin{aligned}\mathbf{x} &= [\,x_0,\, x_1,\, \dots\,,\, x_{N-1}]^{\mathrm{T}}\\ &= [\,\mathbf{x}[0],\, \mathbf{x}[1],\, \dots\,,\, \mathbf{x}[N-1]\,]^{\mathrm{T}}.\end{aligned}$$

If the indexed vector is obtained though an expression, we may denote it as illustrated by the following example:

$$\mathbf{A}\mathbf{x} = [\,(\mathbf{A}\mathbf{x})[0],\, (\mathbf{A}\mathbf{x})[1],\, \dots\,,\, (\mathbf{A}\mathbf{x})[N-1]\,]^{\mathrm{T}},$$

i.e. $(\mathbf{Ax})[n]$ is the n'th element of the vector obtained by applying $\mathbf{A}$ to $\mathbf{x}$. Matrix elements are indexed similarly to vector elements:

$$\mathbf{A} = \begin{bmatrix} A_{0,0} & \dots & A_{0,N-1} \\ \vdots & \ddots & \vdots \\ A_{M-1,0} & \dots & A_{M-1,N-1} \end{bmatrix} = \begin{bmatrix} \mathbf{A}[0,0] & \dots & \mathbf{A}[0,N-1] \\ \vdots & \ddots & \vdots \\ \mathbf{A}[M-1,0] & \dots & \mathbf{A}[M-1,N-1] \end{bmatrix}.$$

If a matrix is obtained by an expression, we may use e.g. $(\mathbf{AB})[m,n]$ to refer to the m,n'th element of the matrix $\mathbf{AB}$.

Reserved Symbols

The following letters are reserved for particular purposes throughout the thesis:

N, n	N is the number of acoustic sources in the system, $n = 0, \dots N-1$ is used to index these whenever possible.
M, m	M is the number of microphones in a sound zone, $m = 0, \dots M-1$ is used to index these whenever possible.
K, k	K is the length of the impulse responses from sources to microphones, $k = 0, \dots K-1$ is used to index these whenever possible.
L, l	L is the length of the input filters, $l = 0, \dots L-1$ is used to index these whenever possible.
$\mathbf{G}$	Room response matrix.
$\mathbf{B}$	Response matrix of the bright zone.
$\mathbf{D}$	Response matrix of the dark zone.
$\mathbf{R}$	Covariance matrix of a signal.
$\mathbf{d}$	Desired response.
$\mathbf{q}$	Response of all input filters at one frequency.
$\mathbf{h}$	Time domain response of an input filter.

General functions and Operators:

$\mathbf{I}$	The identity operator.
$\boldsymbol{\mathcal{F}}$	The discrete Fourier transform (DFT).
$\boldsymbol{\mathcal{F}}_{\{K,L\}}$	Discrete Fourier transformation for length L signals with zero padding to length $K \geq L$.
$\delta(t)$	The Dirac delta impulse. While this is not a well-defined function, we allow ourselves to treat is as such when this causes no problems.
$\delta[n]$	The discrete impulse function.
$\|\|\mathbf{x}\|\|$	The ℓ^2-norm of vector $\mathbf{x}$.
$\lceil x \rceil$	The ceiling function (rounding smallest integer larger than x).
$\lfloor x \rfloor$	The floor function (rounding largest integer smaller than x).
$\lfloor x \rceil$	The rounding function (rounding to nearest integer).
$\mathbf{x} \circ \mathbf{y}$	Entry-wise multiplication of vectors $\mathbf{x}$ and $\mathbf{y}$ (vectors must be of the same length).
$\mathbf{x} * \mathbf{y}$	Linear convolution of vectors $\mathbf{x}$ and $\mathbf{y}$.
$\mathbf{x} \circledast \mathbf{y}$	Circular convolution of vectors $\mathbf{x}$ and $\mathbf{y}$ (vectors must be of the same length).

Contents

Chapter 1

Introduction

The ability to reproduce recorded or artificially generated audio has transformed the world we live in. Today, many of us are almost never out of range of an audio reproduction system. Many of the rooms in which people spend their time, both at work and at home, have audio systems installed. Furthermore, we consider audio systems to be a natural part of cars, phones, televisions, etc. With such a large basis for exposure to audio, there is a great interest in understanding and controlling the impact that audio has on humans. Several studies have indicated that exposure to different types of audio has profound implications on productivity and cognitive functioning (e.g. [1] which contains a rather thorough bibliography on the subject). The massive use of audio reproduction equipment has also led to many developments with respect to the aim of reproducing audio faithfully. These developments have made multichannel audio systems increasingly common [2].

The project documented by this thesis takes part in the aim to improve performance of audio reproduction as well as to improve the control of exposure to audio. More specifically, the project concerns an untraditional use of multichannel audio systems which has recently gained attention. Instead of using the loudspeakers as individual sources, they are used in a collaborative effort to obtain an approximation of some sound-field with desired properties, within a room. This gives rise to intriguing possibilities with respect to both audio quality and control of exposure.

This chapter serves as an introduction to the subject of the thesis. Section 1.1 provides a short historical perspective on the reproduction of sound with a focus on multichannel audio. Section 1.2 discusses some features of possible future audio systems which could be desirable but which are currently unattainable. Section 1.3 provides an overview of the current state of the art with respect to sound zones and sound-field control. Section 1.4 establishes the central problem with which the thesis is concerned and formulates a working hypothesis for the project.

1.1 Reproduction of Sound

The history of electrical reproduction of sound started in the middle of the 19th century. In the early 1860s Johann Philipp Reis worked on one of the worlds first telephones [3]. This device used an electric loudspeaker to reproduce sound. In 1877 Edison presented the phonograph which was capable of recording and reproducing audio (albeit not with an electric loudspeaker) [3]. Almost immediately after these developments it was noted that multi channel systems could be used to provide a more realistic reproduction of sound [2]. In 1933, a patent for audio reproduction across “two acoustic paths” was granted to A. Blumlein [4]. This introduced the principle of stereo and predicted many of the ways in which stereo audio would be implemented and used [2]. In 1940, Disney Studios used a multi channel audio system to provide a sense of direction to viewers of the film Fantasia in cinemas [2]. From the late 1960s quadraphonic (4-speaker) systems began to emerge on the consumer market [2]. These, however, never managed to gain significant popularity [2]. In the 1990s, standards for consumer surround sound such as the Dolby Digital 5.1 began to emerge[1]. These have been more successful and home theatres systems with five channels and a separate low frequency unit are now common. High performance audio systems for cars have as well received much attention lately. Here, audio systems with as many as 19 loudspeakers are available[2]. Professional cinema systems currently offer as many as 64 audio channels[3]. At the same time, increasingly sophisticated signal processing methods are used to optimise the performance and increase the flexibility of the systems[4] (see e.g. [5]).

1.2 The Future of Sound Reproduction

The section above displays a trajectory towards increasingly sophisticated audio systems with an increasing number of channels. The reason for this trend may seem to be obvious; “more is always better”. On the other hand, it may be worthwhile to dwell upon what we expect to gain from the increasing complexity. The straightforward purpose of multichannel audio systems is to provide the user with a sense of direction by letting the sound come from different angles. On top of this, systems with many channels may provide some less obvious possibilities. Consider the following issues which are encountered with typical home audio systems (regardless of purpose or price range):

[1]See `http://www.dolby.com/us/en/consumer/technology/home-theater/dolby-digital.html`.

[2]Bang & Olufsen produces such a system for the Audi A8/S8 models: `http://www.bang-olufsen.com/en/car-audio/car-models/audi/A8,-,S8`.

[3]This is the case for the Dolby Atmos sound system: `http://www.dolby.com/us/en/consumer/technology/movie/dolby-atmos-details.html`.

[4]This is clearly visible on the Bang & Olufsen webpage referenced above. Another notable example is that of Steinway Lyngdorf: `http://www.steinwaylyngdorf.com/technology-and-innovation`.

- The loudspeakers should be carefully positioned to produce the optimal audio experience. Even when this is done, the room significantly influences the sound (it may be possible to partially correct for this, see e.g. [5]).
- For the best experience the user must be situated in a very specific location in the room. Users outside this area do not experience the full potential of the audio system. Therefore, if multiple users are scattered across the room, only a single person can expect to hear the system's full potential[5].
- The audio system does not take each individual users specific requirements into consideration. E.g. some users may wish to listen at a higher volume than others. Some may not wish to hear the system at all.

With the above issues in mind, we may attempt to imagine an ideal audio system, without considering the limitations of physics or technology. This can be used to outline possible directions of research. For a home audio system, we may imagine features such as the following:

- Neither the room nor the loudspeaker positioning has any influence on the perceived sound.
- All users experience the optimal system performance regardless of their locations in the room.
- Each user can control parameters such as volume without affecting the experience of the other users. A user may even choose not to hear the system at all.
- Users may choose to listen to entirely different audio programs although being located in the same room.
- Users may move around the room without this changing their audio experience.

The listed features are very desirable, but may also appear somewhat unrealistic. A system with such features must clearly have a more precise control of the sound-field than what is the case for more conventional audio systems. This must be the case to an extent where the system can separate the users acoustically from one-another (such that the system can control individually what each user listens to without affecting the others). It may be noted that this can be accomplished simply by letting each user wear a set of headphones. This is, however, often undesirable (e.g. due to issues with discomfort). Furthermore, it is not necessarily possible to produce a satisfactory surround

[5]See e.g. Dolby's recommendations for home theatre speaker placement: `http://www.dolby.com/us/en/consumer/setup/connection-guide/home-theater-speaker-guide/index.html`.

sound experience with headphones[6]. A realisation of this audio system should therefore use other means to separate the users acoustically. More generally, the system should be able split the room in multiple zones and control the sound in each zone separately. Depending on the specific application it may be necessary to control the sound-field inside the zones rather well (e.g. to make surround sound possible inside a zone).

1.3 Current Sound Zone Technology

The section above considers an ideal audio system without considering whether the system is realizable or not. It may appear counter-intuitive that an audio system should be capable of e.g. playing music to one user while being inaudible to another individual in the same room. It has, however, been pointed out that this is not impossible [7, 8]. Research in the topic began in the the 1990s (see e.g. [7, 9]). Different approaches for solving the problem has since then emerged. A common characteristic for most of these methods is that they use audio systems with multiple channels to obtain the desired results. Simple approaches consist in using directional loudspeakers or loudspeaker arrays to radiate sound towards specific users [7]. More sophisticated methods take a more unified approach and also account for the acoustics of the environment. These methods typically define some desired behaviour of the sound-field and then optimise the similarity between the behaviour of the actual sound-field and the behaviour of the desired one [8–10]. This makes it possible to define different zones with different properties. In this thesis we concern ourselves with only this type of methods. An example of a much studied type of experimental set-up is seen in Figure 1.1 (this type of set-up is studied in e.g. [8, 10–14]). This set-up considers a system with a single input signal and two zones. The objective with the set-up is to keep one zone quiet while allowing the input signal to be heard clearly in the other zone. The zones are termed, respectively, the *dark zone* and the *bright zone*. Terms such as “quiet” and “heard clearly” are of course subjective and should be replaced by measurable quantities such as sound pressure. This system is a simplification of the ideal system described above. On the other hand it is much more specific and its performance is simpler to measure. If multiple inputs are provided, it may also be attempted to reproduce surround sound in one zone while keeping the other zone quiet (see e.g. [10]).

The investigated systems consist essentially of two things. The first is an array of loudspeakers. Much of the available literature considers a circular array with a uniform distribution of loudspeakers pointing inwards (as shown in Figure 1.1). The number of loudspeakers considered ranges from only a few [15] to several hundreds [10]. The second part of the sound zone system is a signal processing unit which determines an appropriate input to each of the loudspeakers based on a single input signal; the signal which should

[6]Reproduction of surround sound in headphones is the focus of binaural audio technology [6].

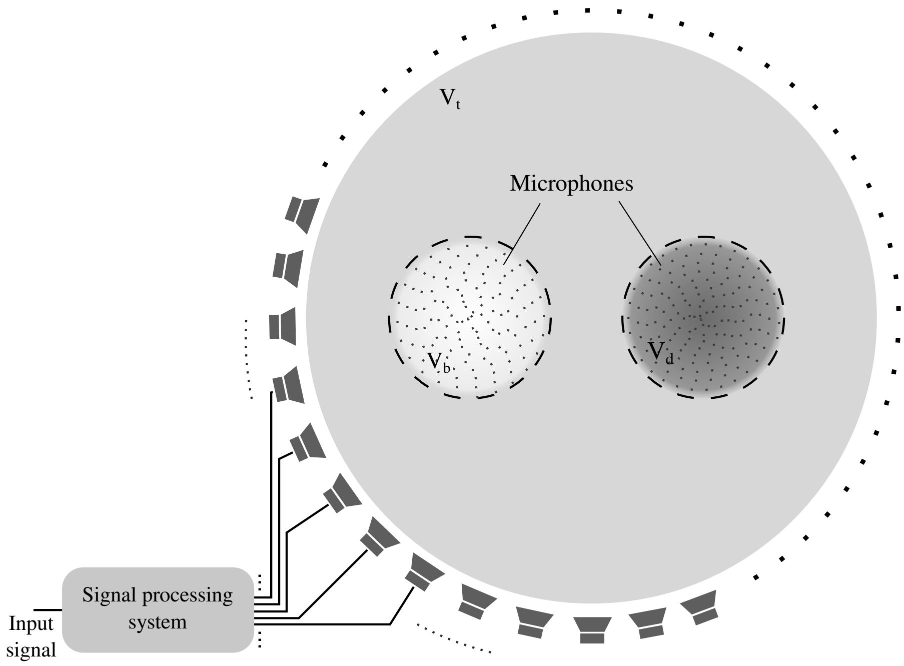

Figure 1.1: An example of a typical set-up for the generation of multiple sound zones. A volume, V_t (i.e. a room), is equipped with a circular array of loudspeakers. Furthermore, two distinct zones, V_d and V_b, are marked. The objective is to make a given acoustic signal appear brightly in V_b while keeping V_d as silent as possible. This objective is sought by use of a signal processing system which uses a single input signal to determine individual input signals to each loudspeaker.

be heard in the bright zone (one may consider multiple inputs if stereo or surround sound is desired in the bright zone). Typically, the signal for each loudspeaker is produced by delaying and scaling the system input signal (see e.g. [8]). This, however, only allows the system to work well at a single frequency. A slightly more complicated approach is to use a separate digital filter to shape the input signal to each loudspeaker [16, 17]. This can to some extent allow the system to function at a range of frequencies. Additionally, most of the investigated methods also employ an array of microphones to measure the response of the defined sound zones to the loudspeakers.

Most of the methods for producing sound zones can be lumped into two categories; energy methods and synthesis methods. The methods of both categories rely on optimisation to produce a desired sound-field. The difference lies in the choice of objective function. The two categories are briefly described below.

1.3.1 Energy Methods

The energy methods use the acoustic potential energy in each zone as a measure of the system performance. The acoustic potential energy in a zone may be estimated with an array of microphones placed in the zone [8, 11]. Typically, the methods consider set-ups like the one shown in Figure 1.1 or similar [8, 11, 12]. The most widespread method is that of Acoustic Contrast Control (ACC), in which one seeks to maximize the ratio of energy in the bright zone to energy in the dark zone [8, 12]. This can produce a significant difference in sound pressure between the two zones, both in theory and in practice [18]. An issue with these methods is the lack of control of the sound-field within the zones. It is therefore not possible to control e.g. the direction of arrival of the sound and the energy may be distributed very unevenly in the zone.

1.3.2 Synthesis Methods

The synthesis methods are methods which approach the problem by attempting to take complete control of the sound-field inside the zones [10, 15]. The methods minimise some objective function which is a measure of the degree to which the sound-field in the zones match with a desired one. The desired sound-field in the bright zone can e.g. be a plane wave with some particular direction and amplitude. This makes it possible to control the direction of arrival of sound inside the individual zones. Sound field synthesis is not by any means limited to the creation of sound zones. The methods allow completely general control of the sound-field within some space [9]. The synthesis methods require very large numbers of loudspeakers and microphones as compared to the energy methods [18].

The energy and synthesis methods produce rather different results as indicated by a comparison given in [18]. The energy methods produce large differences in sound pressure between the zones but lack the ability to control the structure of the sound-fields inside each zone. This can lead to unstructured sound-fields with large fluctuations in sound pressure. The synthesis approach avoids these issues by directly controlling the sound-field inside each zone. This, on the other hand, makes it more difficult to obtain a significant difference in sound pressure between the two zones. By these observations, the methods can be seen as two extremities. This has recently led to the consideration of hybrid methods which seek to find a midpoint between the two. One approach is to simply consider a weighted sum of the two objective functions as investigated in [13]. Other types of intermediate methods have as well been considered [14].
The discussed methods almost exclusively consider sinusoidal input signals of fixed frequency. The systems are therefore designed to work well at only a single frequency. This greatly limits the practical value of the systems and clearly indicates a need for additional research. Exceptions are e.g. [19] and [17] which shortly discusses respectively ACC and PM for multiple frequencies.
A number of scientific works have suggested various applications of the sound zone technology. Personal sound for PC users is proposed in [7] and further investigated in [12, 20]. Personal sound for mobile devices such as mobile phones is investigated in [21]. Aircraft

headrests with personal audio are considered in [16, 22]. The possibility of having independent sound zones for the front and rear seats in a car is considered by [23].

1.4 Initial Problem and Hypothesis

The preceding sections discuss some limitations of current audio systems and identifies the subject of sound zone technology as a possible mean to overcome these. A significant issue in sound zone technology is the fact that most of the currently developed methods assume the input to be a sinusoidal signal with a fixed frequency. This is obviously not a realistic assumption for signals such as speech or music. We therefore state this as an initial problem of the project:

> *Current methods in sound zone technology are typically designed to perform well at only one fixed frequency. This is not sufficient for reproduction of audio signals such as speech or music.*

An investigation of the current methods is a natural starting point for the project work. Current methods consist of optimisation problems which are solved for one particular frequency. It is believed that it is possible to extend such methods to consider wideband signals such as speech or music by solving the involved optimisation problems across a range of frequencies simultaneously. Such approaches are briefly discussed in [17, 19]. We use this as a working hypothesis for the project:

> *Current sound zone technology can be adapted to wideband signals by considering methods which solve the involved optimisation problems for a range of frequencies.*

The overall objective of the project is to investigate this hypothesis. We do this by means of analytical derivations and computer simulations. Chapter 2 introduces the simulated measurement set-up under which the investigations are carried out. Chapter 3 introduces the central energy- and synthesis methods and compares their performance with the simulated measurement set-up. Chapter 4 investigates the possibilities of extending the single-frequency methods to handle wideband signals. This problem has not been widely investigated in the scientific literature. We therefore perform an exploratory investigation of the problem, where we discuss a number of approaches for devising wideband sound zone methods. This yields insights into the complications which arise when considering wideband signals rather than single frequency signals. The discussed methods are investigated by use of simulations in Chapter 5.

1.5 Summary

This chapter provides a foundation for the remainder of the thesis. It is indicated that audio systems with increasingly many channels are becoming available, both for commercial use (e.g. cinemas) and for personal use. An audio system with many channels enables increased control of the sound-field within a room. This can in principle make it possible to separate a room in multiple sound zones with different audio content. Currently investigated methods for doing this typically only consider pure sinusoidal input signals and perform poorly for wideband signals such as speech or music. The aim of this project is to investigate the possibility of deriving methods which perform well for wideband signals.

Chapter 2

Methodology

In Chapter 1 we loosely sketched the idea of sound zones. The idea is illustrated as a very desirable future consumer product, but the typical type of set-up currently used for research is as well introduced. This chapter serves the purpose of introducing the necessary background for discussing state of the art methods for creating sound zones. Section 2.1 delimits and concretises the scope of the investigation carried out in the thesis. Section 2.2 introduces the virtual measurement set-up which is used to evaluate the discussed methods. We use the Mirror Image Model (MIM) to simulate realistic room acoustics. The principles of the MIM are covered in Section 2.3. The computational demands connected with the MIM are rather substantial. Some care therefore needs to be taken with respect to the software implementation of it. These considerations are reviewed in Section 2.4. The chapter assumes an understanding of a few concepts from basic acoustics. The necessary concepts are briefly reviewed in Appendix A.

2.1 Scope of the Investigation

Chapter 1 provides a very general discussion on many possible applications and possibilities of the sound zone technology. Before investigating the involved methods more closely, it is necessary to properly delimit the scope of the investigation. With reference to the discussions of Chapter 1, we state the following concretisations:

- We consider only single channel audio signals. I.e we do not consider recreating stereo or surround sound audio.
- We consider set-ups with a single bright zone and a single dark zone. The zones have the shape of flat disks.
- Each zone is populated by microphones which are used for sampling the internal sound-field. The number of microphones, M, is the same in both zones.

- The input signal to each loudspeaker is obtained by filtering a single system input signal with a Finite Impulse Response (FIR) filter.

- We assume that the relationship between the inputs of the loudspeakers to the outputs of the microphones can be modelled as a linear system (that is, a matrix of transfer functions).

The first two points above are made to simplify the problem. While multi-channel sound and complicated arrangements of zones may be of interest, we do not pursue such matters. This is done to keep focus as much as possible on the main objective of extending sound zones to wideband signals. Solving the problem for such simple circumstances should surely precede the investigation of more complicated ones. Populating the zones with microphones provides the necessary foundation for quantifying terms such as *bright* and *dark* by relating these to the sound-field. Using filters to shape the input signal is the natural extension of using a weighting and a delay as is the convention for single frequency systems. This essentially allows for applying a frequency dependent weighting and delay of the input signal for each loudspeaker. The use of FIR filters provides full control of the filter impulse responses and guarantees stability (see e.g. [24] for an introduction to such matters). Lastly, we assume that the entire audio system from loudspeakers (including D/A-conversion, amplification, etc.) to microphones (including amplification, A/D-conversion, etc.) can be modelled as a linear system. This includes the assumption that air acts as a linear elastic medium, which is shortly discussed in Appendix A. The assumption that all electronic and electro-acoustic components in the signal chain are linear may be reasonable or not, depending on the particular components chosen. A discussion of such matters is beyond the scope of this thesis. We refer to the fact that this assumption is very common in the sound zone literature (see e.g. [8, 10]). Furthermore, we circumvent the possible problems of non-linearities by using simple, idealised and linear components in the simulations.

The above remarks lead to a system such as the one shown in Figure 2.1. The system receives a single input signal which is filtered by N individual filters. We term these filters the “input filters”. Each resulting signal is used as input signal to one loudspeaker. The resulting acoustic field is sampled by $2M$ microphones; M in each zone. The linear relationship between loudspeaker input and microphone output is modelled by a matrix of transfer functions. The $2NM$ transfer functions used to represent this relationship are termed the “room filters”. The input filters are part of the signal processing system and are therefore entirely controllable. The room filters are given by the specifications of the loudspeakers and microphones as well as by the geometry and the physical properties of the space in which the sound zone set-up is placed. They are therefore not controllable as seen from a signal processing perspective (the filters can of course be changed by physically changing the set-up). The overall objective is essentially to design the input

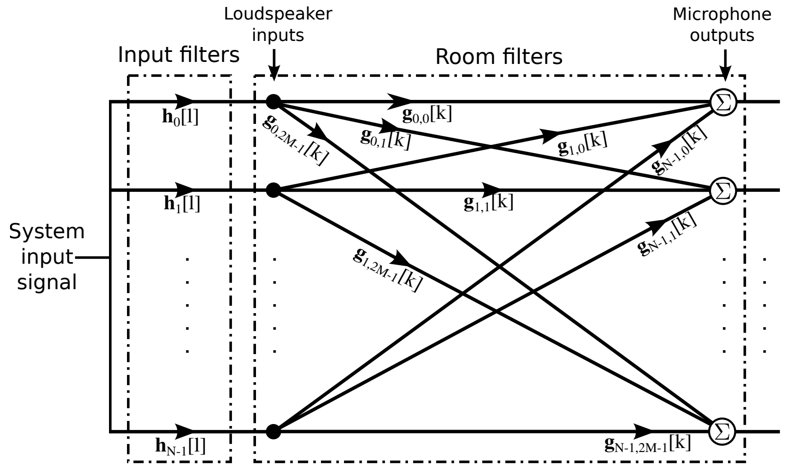

Figure 2.1: A schematic diagram of the sound zone system. A single input is provided to the system. The input signal for each of the N loudspeakers is generated by filtering this single input signal. The output at the $2M$ microphones is computed by filtering the loudspeaker inputs with a matrix of transfer functions. This matrix comprises room acoustics as well as acoustics of loudspeakers and microphones. The system contains $2M$ microphones as there are two zones, each of which contains M microphones.

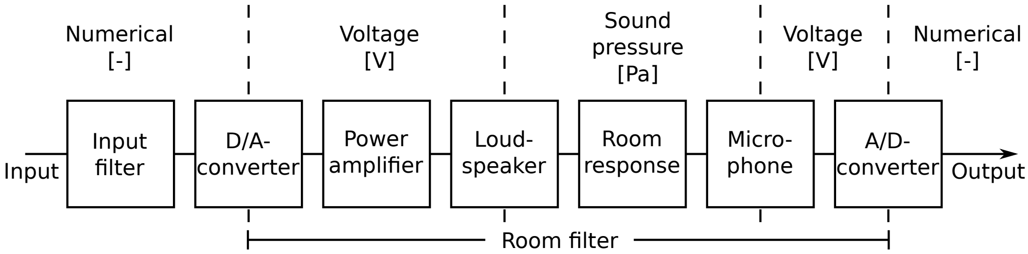

Figure 2.2: An illustration of the units through a real-world sound zone system, from the input of the signal processing system to the microphone outputs.

filters such that much sound energy is measured in the bright zone and little energy is measured in the dark zone.

2.1.1 A Note on Physical Units

There are several different physical units involved in a sound zone system. Figure 2.2 shows an example of how the signal may pass through a real-world sound zone system, with the associated change of units. The signal processing system is assumed to be digital and the input signal is supplied digitally as well. The filtered signal is converted to an analog signal which is amplified with a power amplifier. The amplified signal is used to drive a loudspeaker. The vibrating loudspeaker membrane causes an excitation of the air with a certain volume velocity. This excitation causes acoustic waves to be emitted.

These acoustic waves interact with the room and are sensed by the microphones. The waves are thereby converted back into an electrical signal. This signal is finally converted back to a digital signal.
We model everything from the input of the D/A-converter to the output of the A/D-converter with the room filters. Therefore, the room filters have no physical units. All the blocks on Figure 2.2 have complicated characteristics in practice and could be studied in great detail. To keep the focus of this thesis strictly on the design of the input filters, we make the following great simplifications:

1. All loudspeakers are modelled as simple monopole sources (the properties of such are briefly reviewed in Appendix A).

2. The source volume velocity (in $\frac{\mathrm{m}^3}{\mathrm{s}}$) is directly given by the numerical value of the input to the D/A-converter, with appropriate change of units. In other words, the volume velocity of the source is proportional to the input, with a constant of proportionality given by $1\frac{\mathrm{m}^3}{\mathrm{s}}$. This means that the outputs of the input filters directly determine the volume velocity of the sources.

3. The output of the A/D-converter is given directly by the sound pressure at the location of the associated microphone (in Pa). Thus, the A/D-converter output is proportional to the sound pressure, with a constant of proportionality given by $1\frac{1}{\mathrm{Pa}}$. Note that this, among other things, assumes a very simple and idealised model for a microphone.

With these assumptions, the free-field room filter response for loudspeaker n and microphone m is given by:

$$g_{n,m}(t) = \frac{\delta(t - \frac{d_{n,m}}{c})}{4\pi d_{n,m}}, \tag{2.1}$$

where $d_{n,m}$ is the distance between loudspeaker n and microphone m. This is seen directly from the properties of the monopole source (discussed in Appendix A). When reflections from walls are taken into account, these result in additionally delayed impulses with lower amplitude being added to the response (assuming that a reflection can be modelled as a simple attenuation).

2.2 The Studied Set-up

To allow for comparing the performance of different methods for creating sound zones, we define a specific test set-up which is used throughout the thesis. This set-up is intended to be investigated by the use of computer simulations. The specification of a test set-up therefore involves specifying all parameters which are required to perform acoustical simulations. These are:

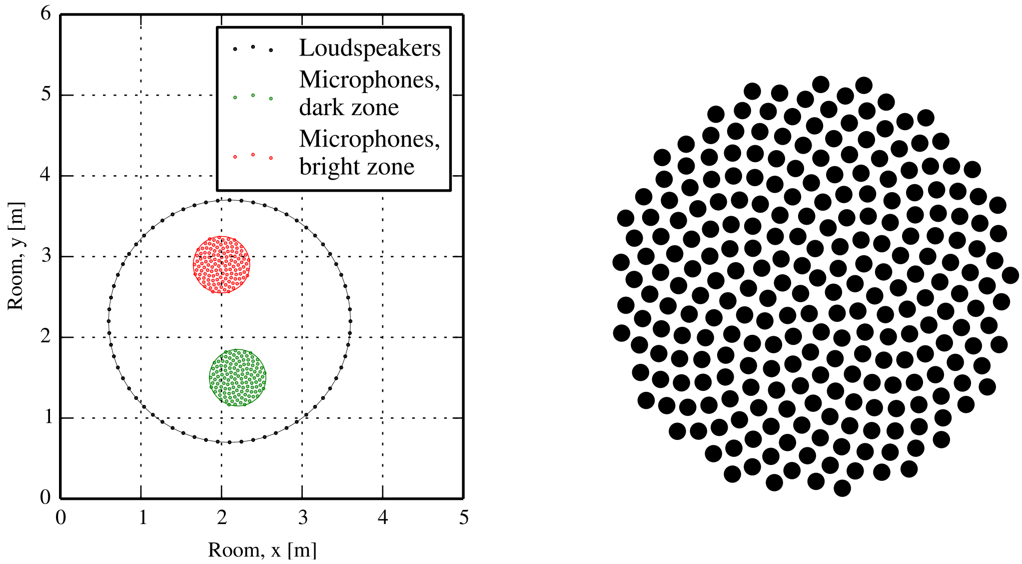

Figure 2.3: Left: The studied sound zone set-up seen from above. The set-up is entirely contained in a horizontal plane 1.6 m above the floor. The figure is generated by the script `setup_layout.py`. Right: An example of the Vogel method used for determining microphone positions. The figure is generated by the script `vogel.py`.

- The number of acoustical sources (loudspeakers), their positions and their individual properties.
- The number of microphones, their positions and their individual properties.
- The desired layout of the involved sound zones.
- The overall properties of the involved signal processing system. Since this system is the topic of investigation, some details do of course have to be left out for further investigation.
- Physical properties of the space in which the set-up is contained. This involves the speed of sound as well as the positioning of walls and their respective reflection properties.

The set-up considers a circular array of sources containing one bright zone and one dark zone. Both of the zones are disk-shaped. Such layout details are without doubt important to the resulting performance. It is, however, not the topic of investigation in this thesis and the type of layout has simply been chosen because it is the most widely studied one in sound zone literature (e.g. [10, 13, 14, 18, 25, 26]).

Speed of sound (c)	340 m/s
Room dimensions (W×L×H)	$5 \times 6 \times 3$ m
Room simulation	MIM
Length of simulated response	200 ms
Reflection coefficients (walls, floor/ceiling)	-0.5, -0.75
Sampling Frequency	10.0 kHz
# sources	48
Source array center (W×L×H)	$2.1 \times 2.2 \times 1.6$ m
Source array radius	1.5 m
# microphones, dark zone	96
# microphones, bright zone	96
Dark zone center (W×L×H)	$2.2 \times 1.5 \times 1.6$ m
Bright zone center (W×L×H)	$2.0 \times 2.9 \times 1.6$ m
Dark zone radius	0.30 m
Bright zone radius	0.30 m

Table 2.1: The parameters of the studied set-up.

The specific set-up contains a circular array of 48 monopole sources. The bright zone and the dark zone are both given by flat disks which are each populated by 96 microphones. As discussed, we use simple idealised models of loudspeakers and microphones. The entire set-up is contained in a horizontal plane 1.6 m above the floor. The two sound zones are placed symmetrically around the the centre of the source array. The set-up has deliberately been placed somewhat asymmetrically as compared to the walls of the space. This is done to avoid strong resonances. The key parameters of the set-up are given in Table 2.1. It is desired that the microphones are approximately evenly distributed. This is done by use of Vogels method [27][1]. Vogels method spreads points approximately evenly on a disk by placing them along a Fermat spiral. Assuming that the coordinate system is centred on the centre of a given sound zone, the polar coordinates of the m'th microphone in this zone are given by:

$$r_m = R\sqrt{\frac{m}{M}}, \qquad m = 0, \ldots, M-1 \tag{2.2}$$

$$\theta_m = m\phi, \qquad m = 0, \ldots, M-1 \tag{2.3}$$

where M is the number of microphones (128 in this case), R is the radius of the zone and

$$\phi = \pi(3 - \sqrt{5}) \tag{2.4}$$

[1]Vogels method was proposed in 1978 as a model for the placement of seeds on a sunflower head.

is the golden ratio. A demonstration of Vogel's method is seen in Figure 2.3. An array of this type can represent a measured sound-field accurately only up to some frequency. This frequency is to some extent determined by the spacing between the individual microphones. For an array in the shape of a rectangular lattice this is given by a 2D equivalent of the sampling theorem which states that the distance between two adjacent samples (microphones) should be less than half the wavelength of the highest frequency component contained in the wave-field [28]. A downside to Vogel's method is that the distance between neighbouring points is not entirely well-defined. Vogel presents a sequence of somewhat complicated derivations on the matter in [27]. Instead of pursuing such matters, the average of the distance between each microphone and its nearest neighbour has been computed for the particular array used. This value is approximately 5.2 cm. An array in the form of a rectangular lattice with microphones separated by such a distance can represent a sound-field accurately up to a frequency of approximately 3.25 kHz. It should be stressed that this does not necessarily apply to an array which is laid out by use of Vogel's method. This fact could speak for using a square lattice of microphones instead. However, it cannot be said to be a certain requirement that the used microphone array can accurately reconstruct a given sound-field. Therefore, we do not further discuss the issue of reconstruction at this point.
Note that the discussed set-up is contained in a plane, 1.6 m above the floor in a three-dimensional room. We therefore study a two-dimensional set-up in a three-dimensional environment.

2.3 Room Simulation

Audio systems are typically used in acoustically reflective environments and the free field assumption is therefore not deemed to be realistic for the generation of simulated room filters. Room simulation is employed to add realism to the simulations used for evaluating the various sound zone methods. What is needed is essentially a method which is able to approximate the transfer function from a source at one point in a room to a microphone at another point. Simulation of room acoustics is a rather well developed field of research and several different techniques exist. Some examples are given [29]:

- Approximate solution to the wave equation by use of e.g. finite element- or boundary element methods.
- The Mirror Image Model (MIM).
- The ray tracing method.
- Hybrid methods.

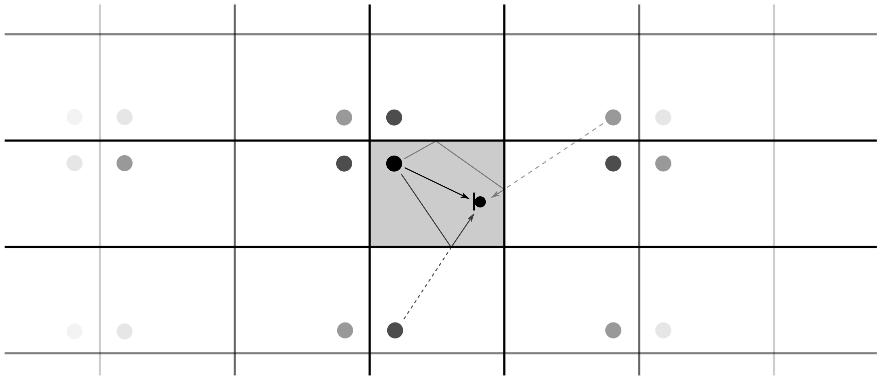

Figure 2.4: A two-dimensional illustration of the principle behind the MIM. The simulated room, which is shown as a grey rectangle, contains a source and a microphone. The reflections from the walls may be modelled by placing virtual sources outside the walls and assuming free field conditions. The direct path from source to microphone is shown with an arrow. Furthermore, two examples of reflected paths are shown. It is shown with dotted lines how the same paths can be modelled with virtual sources outside the room. Each mirror copy of the room provides one reflected path.

A comparison of the methods is given in [29]. Using finite element- or boundary element methods leads to very detailed results but the methods are typically limited to small spaces and low frequencies due to computational demands [29]. The MIM leads to very accurate results, but becomes computationally infeasible for complicated spaces with many reflecting surfaces and for spaces with long reverberation time [29]. The ray tracing method is a stochastic method which is computationally feasible for a greater range of room types than the MIM but less accurate [29]. Hybrid methods typically combine the MIM and the ray tracing method in an attempt to obtain both computational feasibility and high accuracy [29]. We focus our attention strictly on rectangular rooms for simplicity. The mirror image method is especially simple to implement for rectangular rooms and this special case is as well computationally feasible [29, 30]. The MIM is therefore chosen without much hesitation.

2.3.1 The Mirror Image Model in Continuous Time

The MIM is first introduced in [31]. Here, we closely follow an analysis of the method given in [32]. The general idea of the method is to assume free field conditions and instead model the reflections from walls as coming from additional sources placed outside the room. This principle is illustrated in Figure 2.4. The method assumes perfect specular reflections with no scattering. These assumptions may not be applicable for all rooms, but this issue is not further investigated here. When sound is reflected off a surface it loses some energy. This is modelled by assigning a reflection coefficient, β_k with $|\beta_k| < 1$, $k = 1 \ldots , 6$, to each of the six surfaces of the simulated room. The reflection coefficients represent the sound pressure of a reflected wave relative to the direct wave. Using negative reflection coefficients corresponds to the assumption of each reflection giving rise to a phase inversion. It is indicated by [32] that doing so leads to more realistic results and we therefore follow this practice. The MIM assumes that the reflection coefficients are independent of the angle of incidence as well as the frequency of a wave, which may as well not be the case in reality. Under the presented assumptions, the acoustic impulse response from the input of the source to the output of a microphone can be written as a weighted sum of delayed impulses:

$$g(t) = \sum_{b=0}^{\infty} A_b \delta(t - \tau_b), \tag{2.5}$$

where A_b is a real number which depends on the travelled distance and the number of reflections for the b'th path and τ_b is the time it takes to travel from source to microphone along the b'th path and $\delta(t)$ is Diracs delta. Let $\mathbf{x}_\mathrm{s}$, $\mathbf{x}_\mathrm{r}$ and $\mathbf{r}$ denote the source position, the microphone position and the room dimensions respectively:

$$\mathbf{x}_\mathrm{s} = [x_\mathrm{s}, y_\mathrm{s}, z_\mathrm{s}]^\mathrm{T}, \tag{2.6}$$

$$\mathbf{x}_\mathrm{r} = [x_\mathrm{r}, y_\mathrm{r}, z_\mathrm{r}]^\mathrm{T}, \tag{2.7}$$

$$\mathbf{r} = [\Delta_\mathrm{x}, \Delta_\mathrm{y}, \Delta_\mathrm{z}]^\mathrm{T}. \tag{2.8}$$

We also define the sets:

$$\mathcal{U} = \{[u_0, u_1, u_2]^\mathrm{T} \,|\, u_0, u_1, u_2 \in \{0, 1\}\}, \tag{2.9}$$

$$\mathcal{V} = \{[v_0, v_1, v_2]^\mathrm{T} \,|\, v_0, v_1, v_2 \in \mathbb{Z}\}. \tag{2.10}$$

The set $\mathcal{U}$ contains the eight possibilities for a triplet of ones and zeros. The set $\mathcal{V}$ is similar but for triplets of all integers, making the cardinality of the set infinite. We use $\mathcal{U}$ to sum over the eight possible mirror inversions of a rectangular room in three dimensions. We use $\mathcal{V}$ to sum over the mirror copies of the room. We can now rewrite (2.5) with an alternative summing structure:

$$g(t) = \sum_{\mathbf{v} \in \mathcal{V}} \sum_{\mathbf{u} \in \mathcal{U}} A(\mathbf{u}, \mathbf{v}) \delta(t - \tau(\mathbf{u}, \mathbf{v})). \tag{2.11}$$

Each term in the above sum corresponds to the contribution of one mirror image of the source. Note that each entry of $\mathcal{V}$ refers to an entire set of eight mirror inversions of the room. The sum over $\mathcal{U}$ iterates over the individual mirror inversions within such a set.

We have [32]:

$$A(\mathbf{u},\mathbf{v}) = \frac{\beta_1^{|v_0-u_0|}\beta_2^{|v_0|}\beta_3^{|v_1-u_1|}\beta_4^{|v_1|}\beta_5^{|v_2-u_2|}\beta_6^{|v_2|}}{4\pi d(\mathbf{u},\mathbf{v})}, \tag{2.12}$$

$$\tau(\mathbf{u},\mathbf{v}) = \frac{d(\mathbf{u},\mathbf{v})}{c}, \tag{2.13}$$

$$d(\mathbf{u},\mathbf{v}) = ||\mathrm{diag}([2u_0-1,\, 2u_1-1,\, 2u_2-1])\mathbf{x}_s + \mathbf{x}_r - \mathrm{diag}([2v_0, 2v_1, 2v_2])\mathbf{r}||_2. \tag{2.14}$$

The value $A(\mathbf{u},\mathbf{v})$ is the sound pressure at the microphone relative to the volume velocity of the source for the particular path given by $\mathbf{u}$ and $\mathbf{v}$. The value includes losses due to reflections and due to geometric spreading. The values $\tau(\mathbf{u},\mathbf{v})$ and $d(\mathbf{u},\mathbf{v})$ is the time and distance respectively from source to microphone along a particular reflection path. In practice it is of course necessary to limit the sum over $\mathcal{V}$ in (2.11) since this has infinitely many terms. In practice this is done by only summing over values of $(\mathbf{u},\mathbf{v})$ for which $\tau(\mathbf{u},\mathbf{v})$ is smaller than some value.

2.3.2 The Mirror Image Model in Discrete Time

The text above describes how to use the MIM to simulate the impulse response from a monopole source to an omnidirectional microphone in a rectangular room. The resulting impulse response is continuous and is given as a sum of weighted and delayed impulses. For practical purposes we need a discrete impulse response. The task of discretization involves band-limiting and sampling the continuous impulse response. The band-limiting is necessary as a continuous impulse has infinite extent in frequency. Essentially the task is to obtain a vector::

$$\mathbf{g} = [\,\mathbf{g}[0], \ldots, \mathbf{g}[K-1]\,]^{\mathrm{T}}, \tag{2.15}$$

which represents a finite-length, discrete impulse response with a sampling frequency $f_{\mathrm{s}} = \frac{1}{T_{\mathrm{s}}}$. Three different approaches for obtaining such a response have been identified. These are briefly reviewed below.

1) In the initial introduction of the MIM, [31], a very simple approach is used. Here, the computed values of $\tau(\mathbf{u},\mathbf{v})$ are rounded to the nearest multiple of the sampling period. Hereby, each reflection can be modelled by a discrete impulse. By use of this approach, we may construct a discrete representation of (2.11):

$$\mathbf{g}_1[k] = \sum_{\mathbf{v}\in\mathcal{V}}\sum_{\mathbf{u}\in\mathcal{U}} A(\mathbf{u},\mathbf{v})\delta\left[k - \lfloor \frac{\tau(\mathbf{u},\mathbf{v})}{T_{\mathrm{s}}} \rceil\right], \tag{2.16}$$

where $\lfloor\cdot\rceil$ represents rounding to nearest integer, T_{s} is the sampling period and $\delta[\cdot]$ is the discrete impulse function. This approach has the advantage of being simple and computationally inexpensive.

2) It is noted in [33] that the above approach may lead to rather significant errors. Instead, the author notes that the ideal band-limited delta-function is the sinc() function:

$$\delta(t-\tau) \xrightarrow[\text{Band-limit to } \frac{f_\mathrm{s}}{2}]{} f_\mathrm{s} \operatorname{sinc}\left(f_\mathrm{s}(t-\tau)\right) = \frac{\sin\left(\pi f_\mathrm{s}(t-\tau)\right)}{\pi(t-\tau)}, \tag{2.17}$$

where $\frac{f_\mathrm{s}}{2}$ is the frequency above which all frequency content is removed. The resulting function is band-limited such that it can be sampled with a sampling rate of f_s (or more) and reconstructed without aliasing. Band-limiting and sampling (2.11) therefore leads to the following:

$$\mathbf{g}[k] = \sum_{\mathbf{v}\in\mathcal{V}} \sum_{\mathbf{u}\in\mathcal{U}} A(\mathbf{u},\mathbf{v}) f_\mathrm{s} \operatorname{sinc}\left(k - \frac{\tau(\mathbf{u},\mathbf{v})}{T_\mathrm{s}}\right), \tag{2.18}$$

where T_s is the sampling period. Because the sinc() function extends infinitely in time, the evaluation of (2.18) is computationally infeasible. The author of [33] therefore suggests applying a Hanning-window to the sinc()-function:

$$w\left(\frac{t}{T_\mathrm{s}}\right) = \begin{cases} \frac{f_\mathrm{s}}{2}[1+\cos(\frac{2\pi t}{T_\mathrm{w} T_\mathrm{s}})] \operatorname{sinc}(f_\mathrm{s} \frac{t}{T_\mathrm{s}}), & \text{if } -\frac{T_\mathrm{w}}{2} < \frac{t}{T_\mathrm{s}} < \frac{T_\mathrm{w}}{2}, \\ 0, & \text{otherwise,} \end{cases} \tag{2.19}$$

where T_w is the window width in samples. The windowed sinc()-function, $w(\frac{t}{T_\mathrm{s}})$, is used to replace the sinc-function in (2.18):

$$\mathbf{g}_2[k] = \sum_{\mathbf{v}\in\mathcal{V}} \sum_{\mathbf{u}\in\mathcal{U}} A(\mathbf{u},\mathbf{v})\, w\left(k - \frac{\tau(\mathbf{u},\mathbf{v})}{T_\mathrm{s}}\right), \tag{2.20}$$

This is computationally feasible in contrary to (2.18) (assuming that the sum over $\mathcal{V}$ is appropriately restricted). On the negative side, it corresponds to using non-ideal band-limiting. The sampling process therefore introduces some aliasing. The trade-off between computational demands and aliasing is controlled by the window size, T_w.

3) A third option is to construct the response in the frequency domain by taking the Fourier transform of (2.11) [32]:

$$\hat{g}(f) = \sum_{\mathbf{v}\in\mathcal{V}} \sum_{\mathbf{u}\in\mathcal{U}} A(\mathbf{u},\mathbf{v}) \exp\left(-j2\pi f \tau(\mathbf{u},\mathbf{v})\right). \tag{2.21}$$

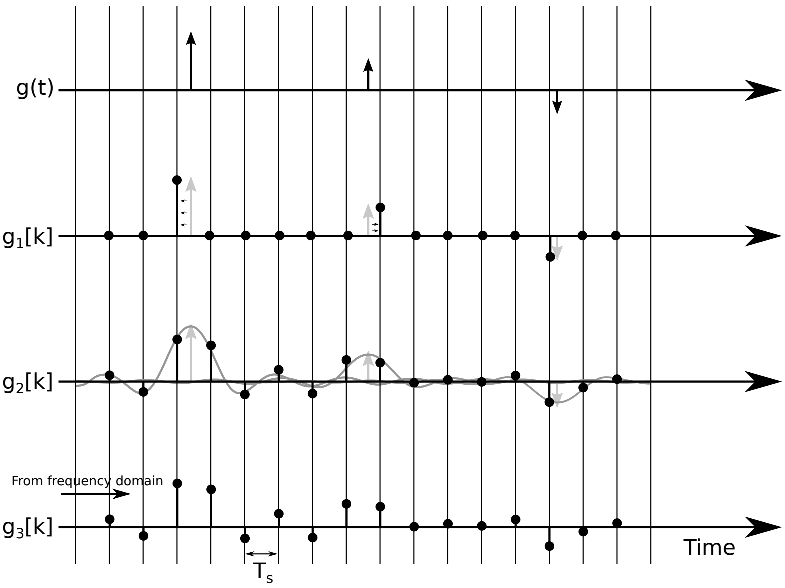

Figure 2.5: An illustration of three different methods for discretizing the response of a room. See the text for details.

We can represent this unambiguously in discrete frequency as a Fourier series, by assuming the impulse response to have finite length of T_{imp} or less[2]:

$$\hat{\mathbf{g}}[k] = \hat{g}(\frac{k}{T_{\text{imp}}}), \qquad k = -\infty, \dots, \infty. \tag{2.22}$$

We may perform band-limiting by simply setting high frequency components in (2.22) to zero. The time-domain response may thus be obtained by inverse Fourier transformation as:

$$\mathbf{g}_3[k] = \sum_{l=-\lfloor K/2 \rfloor}^{\lfloor K/2 \rfloor} \hat{g}[k] \exp{(j2\pi \frac{kl}{K})}, \tag{2.23}$$

where $\lfloor \cdot \rfloor$ represents the floor operator. High frequency components are implicitly set to zero by the limited summation. The time-domain sampling rate is given by $f_{\text{s}} = \frac{T_{\text{imp}}}{K}$. This method is expected to produce rather accurate results but is computationally expensive. The three approaches are compared in Figure 2.5.

The approximation performed in approach 1) is rather imprecise and this method should therefore only be used when the other two methods are computationally infeasible. Practical experiments with methods 2) and 3) have indicated that these are feasible to run on

[2]This can be seen as a property of the time-frequency duality. It is essentially application of the Shannon-Nyquist sampling theorem, but with time and frequency interchanged.

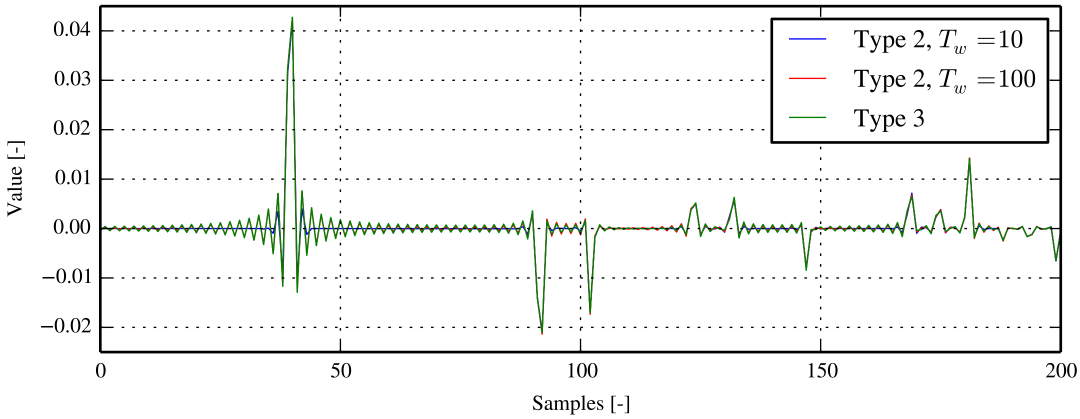

Figure 2.6: A comparison of the second and third presented approaches for obtaining the discrete room response. The second approach is displayed for two different window sizes. The largest of the window sizes is seen to produce results which are almost identical to the third approach. The figure is generated by the script `room_resp_compare.py`.

a modern PC for reasonable selections of the parameters. The first approach is therefore discarded. Methods 2) and 3) are comparable in computational demands if T_w is of the same order as the length of the simulated impulse response. Method 2), however, has the option of trading off approximation quality and computational demands. Figure 2.6 indicates that the methods produce almost identical results even for values of T_w much smaller than the length of the impulse response. The second approach can therefore produce almost the same results as the third approach, but with much less computational effort. The second approach is therefore used for the simulations presented in the thesis. The value of T_w is set to 100 samples.

2.3.3 Selection of Reflection Coefficients

The amount of reverberation in the simulated room is determined by the six reflection coefficients $\beta_1, \ldots, \beta_6$. These specify the amplitude of a reflected wave as compared to the incident wave. In acoustics, this property is often specified by an absorption coefficient, γ, instead. This specifies the fraction of sound energy which is absorbed in a reflection. The relationship between the two are given by [32]:

$$\gamma = 1 - \beta^2. \tag{2.24}$$

The reverberation time of a space is typically defined as the time it takes for an impulsive excitation of the sound-field to decay to one millionth of the initial acoustic energy (energy measured across the entire space) [34]. A very simple means for relating the absorption coefficients to the reverberation time is the Sabine equation [34]:

$$T_{60} = \frac{0.161\ \frac{\mathrm{s}}{\mathrm{m}} \cdot V}{S\xi}, \tag{2.25}$$

where T_{60} is the reverberation time, V is the volume of the space, S is the total surface area in the space and ξ is the average absorption coefficient of the surfaces in the space (weighted by area). It should be remarked that the Sabine equation provides only an estimate of the reverberation time (based on the, so called, diffuse field model of a sound-field) [34]. Here, we select reflection coefficients similar to those of a listening room with very little reverberation. The reason for doing so is twofold: 1) such conditions are deemed to be realistic for a research set-up of the type we are simulating, and 2) the computational demands connected with the simulations we carry out become immense for long room responses. For the floor and ceiling we select a reflection coefficient of 0.75 (an absorption coefficient of 0.5) and for the walls we use 0.5 (an absorption coefficient of slightly over 0.7).[3] This gives a reverberation time of approximately:

$$T_{60} = \frac{0.161\ \frac{\mathrm{s}}{\mathrm{m}} \cdot 5\ \mathrm{m} \cdot 6\ \mathrm{m} \cdot 3\ \mathrm{m}}{2 \cdot 5\ \mathrm{m} \cdot 6\ \mathrm{m} \cdot 0.5 + 2 \cdot (5\ \mathrm{m} + 6\ \mathrm{m}) \cdot 3\ \mathrm{m} \cdot 0.71} = 189\ \mathrm{ms}. \tag{2.26}$$

We therefore limit the computation in the mirror image model to 200 ms (2000 samples). That is, the sums over $\mathcal{V}$ and $\mathcal{U}$ in (2.20) are only carried out for values of $\tau(\mathbf{u}, \mathbf{v})$ less than 200 ms.

2.3.4 Changes in Room Conditions

It has been noted in the literature that sound zone systems can be very sensitive to small changes in the environment [25, 36]. We therefore include this in the investigation. This is done by evaluating the discussed methods before and after a change in the simulation parameters. Generally, the main changes that can occur in room conditions are due to: 1) atmospheric changes which influence the speed of sound, and 2) objects (e.g. furniture or human beings) moving around in the space. We imitate a change of atmospheric conditions by increasing the speed of sound by 5 $\frac{\mathrm{m}}{\mathrm{s}}$ as compared to the value given in Table 2.1. The relationship between the speed of sound and temperature in air is given by [37]:

$$v = (331\mathrm{m/s})\sqrt{1 + \frac{T_{\mathrm{c}}}{273^\circ\ \mathrm{C}}}, \tag{2.27}$$

where T_{c} is the room temperature. From this equation we can find that a change in the speed of sound from 340 m/s to 345 m/s can be caused by a change in temperature

[3] The question of whether these are realistic numbers is difficult to answer as reflection is a highly frequency dependent phenomenon in reality. Tables of absorption coefficients are given in e.g. [35] or on `http://hyperphysics.phy-astr.gsu.edu/hbase/acoustic/revmod.html#c4`.

from approximately 15° C to 23.5° C. The room simulation model we use is too simple to allow for simulating objects inside the room. Instead we mimic physical changes in the room by slightly changing the room dimensions. The room dimensions are decreased by 5 cm on all three axes as compared to the values given in Table 2.1. These two changes constitute what we refer to as changed room conditions throughout the thesis.

2.4 Software Implementation

As suggested in the above, the mirror image method presents significant computational demands[4]. Therefore, an implementation with some focus on execution speed has been deemed a necessity. This section briefly reviews the design choices underlying the implementation. Simulations and testing is carried out in Python[5]. This allows for quick development and provides access to fast implementations of many commonly used algorithms. Furthermore, the open source structure of Python is well in line with the principles of reproducible research [39]. Python is an interpreted language and this can lead to low execution speeds in some situations. This is especially the case for scripts which make heavy use of looping, as doing so introduces significant overhead in Python[6]. A naive implementation of the MIM is likely to use looping to evaluate expressions such as (2.20). There are various methods in Python for speeding up such matters. However, a thorough survey of the possibilities is beyond the scope of this thesis. Furthermore, it is not expected that any of these possibilities can lead to performance which is comparable to that of a C implementation. We therefore follow a strategy of implementing the MIM in C and providing a Python-interface to the functionality. This interface is created with Cython[7]. A total of three C functions are implemented:

- A function to compute the impulse response from a source to a microphone using approach 2) as discussed above.
- A function to compute the impulse response from a source to a microphone using approach 3) as discussed above.
- A function to compute the responses from a source to a rectangular lattice of microphones at a single frequency.

The purpose of the third function is to allow for plotting the acoustic field of some room in a two-dimensional plane (Figure 2.8 is made by use of this function). As indicated by the size of Table 2.1, quite a large number of parameters must be specified in order to carry

[4]Recent research, however, indicates that it may be possible to implement the MIM much faster under some circumstances, by use of FFT-based algorithms [38].

[5]`https://www.python.org/`.

[6]`https://wiki.python.org/moin/PythonSpeed/PerformanceTips#Loops`.

[7]`http://cython.org/`.

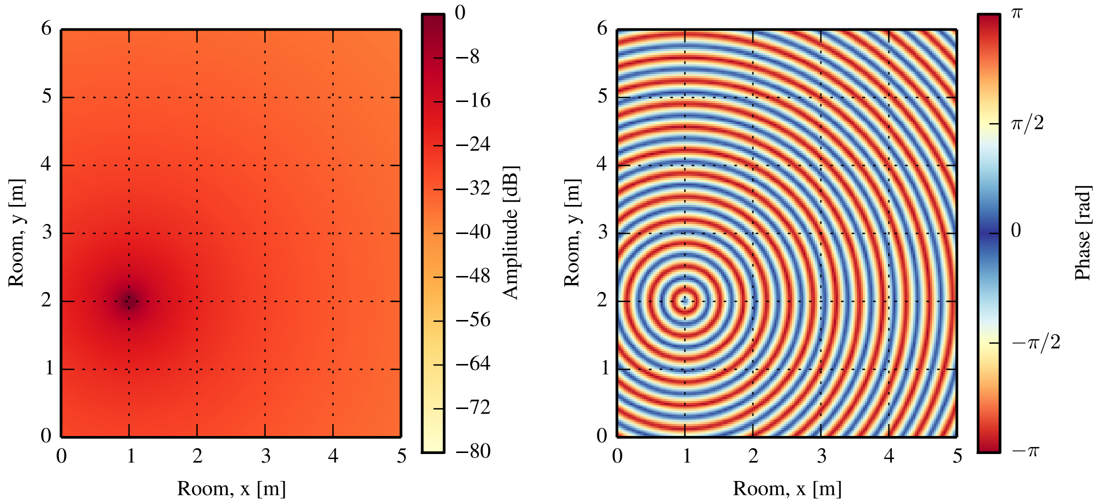

Figure 2.7: A simulated monopole emitting sound at 1000 Hz under anechoic conditions. The sound-field is visualised on the horizontal plane in which the source is located. The properties of the simulated room are those given by Table 2.1, with the exception that the reflection coefficients of all six surfaces has been set to zero. Note the simple spherical sound-field where the amplitude depends only on the distance to the source. The figure is generated by the script `room_demo.py`.

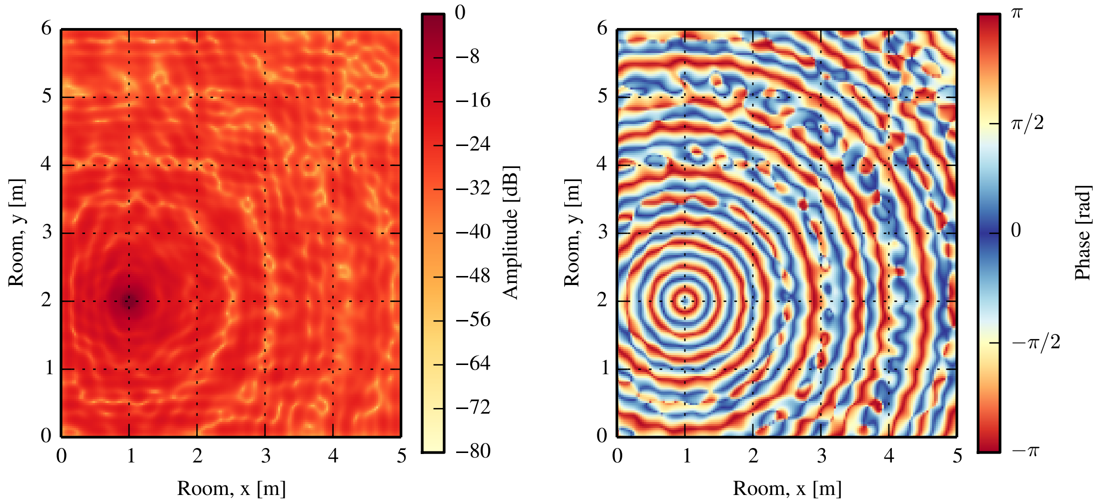

Figure 2.8: A simulated monopole emitting sound at 1000 Hz in a small room. The mirror image method has been used to simulate the properties of the room. The sound-field is visualised on the horizontal plane in which the source is located. The properties of the simulated room are those given by Table 2.1. Note the great complexity of the sound-field as compared to the anechoic case of Figure 2.7. The figure is generated by the script `room_demo.py`.

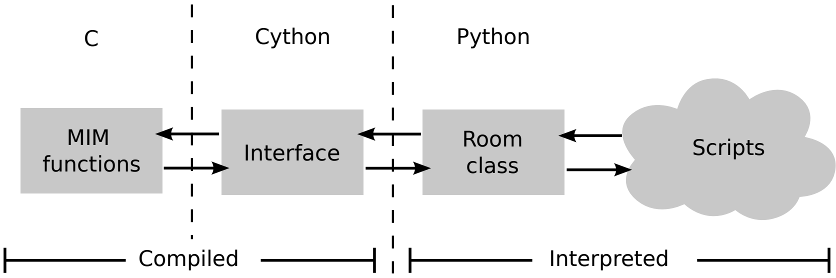

Figure 2.9: A block diagram illustrating the structure behind the implementation of the MIM. The cloud entitled "Scripts" refers to any of the scripts which make use of the functionality supplied by the room class.

out acoustical simulations. We therefore utilise the object oriented nature of Python and use a class to contain the parameters. This class holds all the parameters of Table 2.1 and contains methods which implement the necessary functionality with respect to room simulation. A block diagram of the software layout is seen in Figure 2.9. The room class comprises functionality such as that of simulating the entire matrix of transfer functions describing the relation between N sources and M microphones (Figure 2.1). This task requires NM independent calls to one of the functions for simulating the response from a source to a microphone. In such situations, the room class makes use of multiprocessing to improve execution speed on multi-processor systems. A demonstration of the capabilities of the developed software framework is seen in figures 2.6, 2.7 and 2.8.

2.5 Summary

This chapter introduces the methodology employed throughout the thesis. An acoustical test set-up is presented. The set-up features a circular array of sources surrounding a bright zone and a dark zone. The set-up is investigated purely by use of simulations. To improve the credibility of such simulations, an implementation of the Mirror Image Model (MIM) is used to simulate the acoustics of a rectangular listening room. A specific change in room conditions has as well been defined. This is used to mimic the effects of gradually changing room conditions due to physical and atmospheric changes.

Chapter 3

Single Frequency Methods

As discussed in Chapter 1, a significant amount of research has been carried out with respect to sound zone systems at a single frequency. In this chapter we survey two of the most popular approaches. Section 3.1 presents Acoustic Contrast Control (ACC) and Section 3.2 presents Pressure Matching (PM). The first of these maximises the difference in acoustic potential energy between the two zones, while the other one attempts to synthesize a particular sound-field. A comparison of some key performance measures of the two methods are given in Section 3.3.

3.1 Acoustic Contrast Control

Acoustic Contrast Control (ACC) is introduced in [8] and is a common example of an energy method (as discussed in Chapter 1). The method considers a set-up with one bright zone and one dark zone, at a single frequency, f. The input to each source is a sinusoid of frequency f and with controllable amplitude and phase. The input signal is therefore specified by a single complex scalar representing amplitude and phase. The aim of the method is to maximise the ratio of acoustic potential energy density in the bright zone to acoustic potential energy density in the dark zone.
We have N acoustic sources, the positions of which we denote $\mathbf{x}_{\mathrm{s},0}, \ldots, \mathbf{x}_{\mathrm{s},N-1} \in \mathbb{R}^{3\times 1}$ and $2M$ microphones, the positions of which we denote $\mathbf{x}_{\mathrm{r},0}, \ldots, \mathbf{x}_{\mathrm{r},2M-1} \in \mathbb{R}^{3\times 1}$. The first M microphones are located in the bright zone and the remaining M microphones are located in the dark zone. Assuming a unit amplitude input signal of zero phase, the output of a microphone at the coordinate $\mathbf{x} \in \mathbb{R}^{3\times 1}$ is given by [8]:

$$p(\mathbf{x}) = \sum_{n=0}^{N-1} g(\mathbf{x}|\mathbf{x}_{\mathrm{s},n}, f)\, q_n, \tag{3.1}$$

where $g(\mathbf{x}|\mathbf{x}_{\mathrm{s},n}, f)$ is the response of a microphone at $\mathbf{x}$ the the input of the n'th source at the frequency f. In relation to Figure 2.1 on page 11, the value of q_n is the response of the filter $\mathbf{h}_n[l]$ at the frequency f. Since we only consider the system at a single frequency, it is entirely irrelevant how the filter responds to other frequencies. Similarly, $g(\mathbf{x}_{\mathrm{r},m}|\mathbf{x}_{\mathrm{s},n}, f)$ is the response of the room filter $\mathbf{g}_{m,n}[k]$ at the frequency f. The sum over n adds the contributions of the N sources in the system. We can rewrite (3.1) as an inner product:

$$p(\mathbf{x}) = \mathbf{g}(\mathbf{x}|f)^{\mathrm{T}}\mathbf{q}, \tag{3.2}$$

where:

$$\begin{aligned}
\mathbf{g}(\mathbf{x}|f) &= [\, g(\mathbf{x}|\mathbf{x}_{\mathrm{s},1}, f),\ \ldots\ ,\ g(\mathbf{x}|\mathbf{x}_{\mathrm{s},N}, f)\,]^{\mathrm{T}}, && (3.3)\\
\mathbf{q} &= [\, q_0,\ \ldots\ ,\ q_{N-1}\,]^{\mathrm{T}}. && (3.4)
\end{aligned}$$

The vector $\mathbf{q}$ is sometimes referred to as a *steering vector*. To evaluate how well a given steering vector performs at obtaining contrast between the two sound zones, it is necessary to introduce some measure of the amount of sound energy within a zone. The average acoustic potential energy density in some volume $\mathcal{V} \subset \mathbb{R}^{3\times 1}$ is proportional to the quantity [8]:

$$\begin{aligned}
e_{\mathcal{V}} &= \int_{\mathbf{x}\in\mathcal{V}} \mathbf{p}(\mathbf{x})^{H}\mathbf{p}(\mathbf{x})\ dV\\
&= \int_{\mathbf{x}\in\mathcal{V}} \mathbf{q}^{\mathrm{H}}\mathbf{g}(\mathbf{x}|f)^{\mathrm{H}}\mathbf{g}(\mathbf{x}|f)\mathbf{q}\ dV\\
&= \mathbf{q}^{\mathrm{H}}\left(\int_{\mathbf{x}\in\mathcal{V}} \mathbf{g}(\mathbf{x}|f)^{\mathrm{H}}\mathbf{g}(\mathbf{x}|f)\ dV\right)\mathbf{q}\\
&= \mathbf{q}^{\mathrm{H}}\mathbf{R}_{\mathcal{V}}\mathbf{q},
\end{aligned} \tag{3.5}$$

where $\mathbf{R}_{\mathcal{V}} \in \mathbb{C}^{N\times N}$. The above is used as a measure of the "acoustical brightness" in $\mathcal{V}$. The particular constant of proportionality between this value and the actual value of the average acoustic potential energy density is unimportant for the derivations we consider here. Ideally, we wish to be able to evaluate (3.5) for the bright zone and for the dark zone:

$$\begin{aligned}
e_{\mathrm{b}} &= \mathbf{q}^{\mathrm{H}}\mathbf{R}_{\mathrm{b}}\mathbf{q},\\
e_{\mathrm{d}} &= \mathbf{q}^{\mathrm{H}}\mathbf{R}_{\mathrm{d}}\mathbf{q}.
\end{aligned} \tag{3.6}$$

The matrices $\mathbf{R}_{\mathrm{b}} \in \mathbb{C}^{N\times N}$ and $\mathbf{R}_{\mathrm{d}} \in \mathbb{C}^{N\times N}$ express the correlation between the input of the N sources and the acoustic potential energy in the bright and dark zone respectively. Unless very detailed information of the sound-field is available, it is not generally possible to evaluate $\mathbf{R}_{\mathrm{b}}$ and $\mathbf{R}_{\mathrm{d}}$ (as seen in (3.5), this requires the ability to perform an integral over the sound-field in each zone). Instead, we use the estimates [11]:

$$\begin{aligned}\hat{\mathbf{R}}_{\mathrm{b}} &= \mathbf{B}^{\mathrm{H}}\mathbf{B},\\ \hat{\mathbf{R}}_{\mathrm{d}} &= \mathbf{D}^{\mathrm{H}}\mathbf{D},\end{aligned} \tag{3.7}$$

where the entries of $\mathbf{B}$, $\mathbf{D} \in \mathbb{C}^{M\times N}$ are given by:

$$\begin{aligned}\mathbf{B}[m,n] &= g(\mathbf{x}_{\mathrm{r},m}|\mathbf{x}_{\mathrm{s},n}, f),\\ \mathbf{D}[m,n] &= g(\mathbf{x}_{\mathrm{r},M+m}|\mathbf{x}_{\mathrm{s},n}, f),\end{aligned} \tag{3.8}$$

where the coordinate $\mathbf{x}_{\mathrm{r},m}$ is the position of the m'th microphone and $\mathbf{x}_{\mathrm{s},n}$ is the position of the n'th microphone. That is, we use the microphones in each zone to approximate the volume integral of (3.5) (again, to within some constant of proportionality). The above approximations of $\mathbf{R}_{\mathrm{b}}$ and $\mathbf{R}_{\mathrm{d}}$ are accurate whenever the microphones are spaced closely in comparison with the wavelength of the sound [11]. It should be noted that these approximations are only accurate to within a constant of proportionality[1]. As stated previously, such constant factors are unimportant for the purposes considered here. We define the acoustical contrast between the two zones as [8]:

$$\chi = \frac{e_{\mathrm{b}}}{e_{\mathrm{d}}} = \frac{\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q}}{\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}}. \tag{3.9}$$

This expresses the amount of acoustical potential energy in the bright zone to acoustical potential energy in the dark zone. The aim is to choose $\mathbf{q}$ such as to maximise this quantity. Notice that (3.9) is invariant to scaling of $\mathbf{q}$. The maximiser is therefore not unique. The following formulation of the problem removes this additional degree of freedom:

$$\begin{aligned}\mathbf{q}_{\mathrm{ACC}} = \quad & \arg\min_{\mathbf{q}} \mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}\\ & \text{subject to } \mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q} = e_0,\end{aligned} \tag{3.10}$$

where e_0 determines the amount of acoustic potential energy in the bright zone (to within some constant of proportionality). That is, "minimise the acoustic energy density in the dark zone while keeping the energy density in the bright zone constant". It is easily seen that the solution to this problem maximises (3.9) subject to $\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q} = e_0$. To find the solution of (3.10), we compute the associated Lagrangian[2]:

$$\mathcal{L}(\mathbf{q}, \mu) = \mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q} - \mu(\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q} - e_0), \tag{3.11}$$

[1] The primary measure of solution quality is given by (3.9). Multiplying either $\hat{\mathbf{R}}_{\mathrm{b}}$ or $\hat{\mathbf{R}}_{\mathrm{d}}$ with a constant amounts simply to multiplying this equation with a constant term. If the two zones are identical, any constant factor in front of $\hat{\mathbf{R}}_{\mathrm{b}}$ and $\hat{\mathbf{R}}_{\mathrm{d}}$ cancels in (3.9). The two zones are identical for the set-up introduced in Chapter 2.

[2] See e.g. [40] for an introduction to optimisation theory.

where μ is the Lagrange multiplier. The minimiser of (3.10) satisfies [40]:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{q}} = \mathbf{0}, \tag{3.12}$$

$$\frac{\partial \mathcal{L}}{\partial \mu} = 0. \tag{3.13}$$

From (3.12) we obtain:

$$\begin{aligned}
&2\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q} = 2\mu\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q} \\
&\Updownarrow \\
&(\hat{\mathbf{R}}_{\mathrm{d}})^{-1}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q} = \mu\mathbf{q}.
\end{aligned} \tag{3.14}$$

This shows that the optimal μ is an eigenvalue of the matrix $(\hat{\mathbf{R}}_{\mathrm{d}})^{-1}\hat{\mathbf{R}}_{\mathrm{b}}$ and that the optimal steering vector, $\mathbf{q} = \mathbf{q}_{\mathrm{ACC}}$, is a corresponding eigenvector. Further rewriting (3.14) yields:

$$\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q} = \mu\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}. \tag{3.15}$$

By comparing (3.15) and (3.9) it is seen that $\mu = \chi$ for $\mathbf{q} = \mathbf{q}_{\mathrm{ACC}}$. This indicates that the optimal μ is the largest eigenvalue of $(\hat{\mathbf{R}}_{\mathrm{d}})^{-1}\hat{\mathbf{R}}_{\mathrm{b}}$ and that the minimiser of (3.10) is the corresponding eigenvector. This leaves the magnitude of $\mathbf{q}_{\mathrm{ACC}}$ undetermined. Take $\hat{\mu}$ to be the largest eigenvalue and $\mathbf{v}$ to be the corresponding normalised eigenvector. From (3.13) we have:

$$\begin{aligned}
&\frac{\partial \mathcal{L}}{\partial \mu} = \mathbf{q}_{\mathrm{ACC}}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}_{\mathrm{ACC}} - e_0 = 0 \\
&\Downarrow \\
&\mathbf{q}_{\mathrm{ACC}}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}_{\mathrm{ACC}} = e_0 \\
&\Updownarrow \quad (\text{since } \mathbf{q}_{\mathrm{ACC}} = \mathbf{v}||\mathbf{q}_{\mathrm{ACC}}||) \\
&\mathbf{v}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{v}||\mathbf{q}_{\mathrm{ACC}}||^2 = e_0 \\
&\Updownarrow \\
&||\mathbf{q}_{\mathrm{ACC}}|| = \sqrt{\frac{e_0}{\mathbf{v}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{v}}}
\end{aligned} \tag{3.16}$$

To sum up, we list the necessary steps for determining the minimiser of (3.10), $\mathbf{q}_{\mathrm{ACC}}$:

1. Form the matrices $\hat{\mathbf{R}}_{\mathrm{b}}$ and $\hat{\mathbf{R}}_{\mathrm{d}}$.
2. Find the largest eigenvalue, $\hat{\mu}$, and the corresponding normalised eigenvector, $\mathbf{v}$, of the matrix $(\hat{\mathbf{R}}_{\mathrm{d}})^{-1}\hat{\mathbf{R}}_{\mathrm{b}}$.
3. Compute $\mathbf{q}_{\mathrm{ACC}} = \mathbf{v}\sqrt{\frac{e_0}{\mathbf{v}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{v}}}$.

ACC belongs to a growing class of methods which obtain sound zones by optimising for some objective function which relates to the amount of acoustic energy in the zones. Other examples of such methods are given:

- Acoustic energy difference maximisation [11].
- Brightness control [8].
- Power control [21].

The methods are generally mathematically similar but consider slightly different objective functions and thus obtain different results.

3.2 Pressure Matching

The method presented in the section above attempts to direct much acoustic energy into the bright zone while directing little into the dark zone. The method, however, does not in any way consider how the particular sound-field inside the zones is structured. This can in principle lead to large variations in sound pressure within the individual zones. An entirely different approach to creating sound zones is that of Pressure Matching (PM). PM is an example of a synthesis method (as discussed in Chapter 1) and attempts to take complete control of the sound-field inside the two zones. The basic idea is to specify the desired amplitude and phase at each microphone and find the steering vector, $\mathbf{q}$, that approximates this as closely as possible. This method is applicable to any scenario in which it is desired to obtain control of a sound-field and is therefore not limited to sound zone technology. The method is discussed as a means to create sound zones in [10]. We again consider a set-up with N sources and two zones with M microphones each. We can write the sound pressure at the $2M$ microphones as follows [10]:

$$\mathbf{p} = \mathbf{G}\mathbf{q}, \tag{3.17}$$

where:

$$\mathbf{G} = \begin{bmatrix} \mathbf{B} \\ \mathbf{D} \end{bmatrix} \in \mathbb{C}^{2M \times N}. \tag{3.18}$$

This follows from (3.1) along with the definitions of $\mathbf{B}$ and $\mathbf{D}$, (3.8). We select a vector of desired sound pressures:

$$\mathbf{d} = \begin{bmatrix} \mathbf{d}_\mathrm{b} \\ \mathbf{d}_\mathrm{d} \end{bmatrix} \in \mathbb{C}^{2M \times 1}, \tag{3.19}$$

This vector represents the desired output of each of the $2M$ microphones (each entry representing the phase and amplitude at one microphone). The desired vector is typically obtained by sampling some desired sound-field [10]. This is elaborated in Section 3.2.1. The task of approximating this sound-field is then formulated as a regularized least squares problem [10]:

$$\mathbf{q}_{\mathrm{PM}} = \arg\min_{\mathbf{q}} ||\mathbf{G}\mathbf{q} - \mathbf{d}||^2 + \delta||\mathbf{q}||^2, \tag{3.20}$$

where δ is a regularisation parameter. The steering vector $\mathbf{q}$ describes amplitude and phase of the input signals to each source. The vector $\mathbf{G}\mathbf{q}$ describes the output at the microphones. The solution to this problem is given by [10]:

$$\mathbf{q}_{\mathrm{PM}} = (\mathbf{G}^{\mathrm{H}}\mathbf{G} + \delta\mathbf{I})^{-1}\mathbf{G}^{\mathrm{H}}\mathbf{d}. \tag{3.21}$$

For $\delta = 0$, the solution to the problem is the steering vector which minimises the distance between the actual sound pressure and the desired sound pressure at the microphones (in a least squares sense). In this case, the problem is ill-posed if $\det(\mathbf{G}^{\mathrm{H}}\mathbf{G}) \approx 0$. This issue can be remedied by choosing $\delta > 0$. A small value of δ may lead to solutions with $||\mathbf{q}_{\mathrm{PM}}||$ very large [10]. This may be unwanted as it corresponds to an array which uses much energy for obtaining the desired result. A large value of δ leads to solutions with $||\mathbf{q}_{\mathrm{PM}}||$ small, but with a resulting sound-field that resembles the desired one poorly. The parameter δ therefore introduces a trade-off between solution quality and array effort.

3.2.1 Choice of Desired Sound-Field

The PM approach requires the selection of a desired vector, $\mathbf{d} \in \mathbb{C}^{2M\times 1}$, which represents the desired complex sound pressures at the microphones in both zones. As shown in (3.19), the vector is composed of the desired pressures at the microphones in the bright zone, $\mathbf{d}_{\mathrm{b}} \in \mathbb{C}^{M\times 1}$, and the desired pressures in the dark zone, $\mathbf{d}_{\mathrm{d}} \in \mathbb{C}^{M\times 1}$. Since the dark zone should be as quiet as possible, we set $\mathbf{d}_{\mathrm{d}} = \mathbf{0}$. An approach for determining $\mathbf{d}_{\mathrm{b}}$ is to assume a virtual source to be located at some position, $\mathbf{x}_{\mathrm{vs}} \in \mathbb{R}^{3\times 1}$, and compute the pressure at each microphone, given that this source emits sound at some amplitude, A_0, under free-field conditions [10]:

$$\mathbf{d}_{\mathrm{b}}[m] = A_0 \frac{\exp(\frac{2\pi f}{c}||\mathbf{x}_{\mathrm{vs}} - \mathbf{x}_{\mathrm{b},m}||)}{||\mathbf{x}_{\mathrm{vs}} - \mathbf{x}_{\mathrm{b},m}||}, \qquad m = 0, \ldots, M-1. \tag{3.22}$$

With this approach, $\mathbf{d}_{\mathrm{b}}$ is determined through the selection of a virtual source location, $\mathbf{x}_{\mathrm{vs}}$, and a corresponding volume velocity, A_0. To obtain approximately the same sound level in the bright zone for ACC and PM we can set $e_0 = ||\mathbf{d}_{\mathrm{b}}||^2$. We do this throughout the thesis. The choice of A_0 is somewhat arbitrary for the investigations carried out in this thesis, as it corresponds simply to a scaling of the source outputs.

Some variations of the PM approach have been proposed in relation to sound zones:

- A hybrid method combining the objective functions of PM and acoustic energy difference maximisation [13],
- PM with optimal source placement by use of a LASSO-LS algorithm [19].

3.3 Comparison

The sections above introduce two fundamentally different approaches for obtaining sound zones at a single frequency. To gain additional insights into the methods, these have been applied to the simulated set-up presented in Chapter 2.

3.3.1 Performance Metrics

In order to compare the performance of the two approaches, we must define some metrics for measuring their performance. The most straightforward requirement for a good method is its ability to create a large difference in the sound pressure between the two zones. This ability is measured by the acoustic contrast, as introduced in Section 3.1:

$$\chi = \frac{\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{b}}\mathbf{q}}{\mathbf{q}^{\mathrm{H}}\hat{\mathbf{R}}_{\mathrm{d}}\mathbf{q}}. \tag{3.23}$$

It should be noted, that since χ is the value optimised by ACC, it is not possible for any other method to obtain a higher contrast than this. In spite of this, acoustic contrast is used as a primary measure of solution quality throughout the thesis. We might also use the objective function of PM, (3.20), to evaluate the performance of the methods. However, this function measures the ability of a method to recreate a specific sound-field. It is therefore not deemed to be a reasonable measure of performance for methods which do not aim to recreate any particular sound-field.

An issue which is briefly discussed in Section 3.2 is the matter of sound pressure variation within the bright zone. One could devise a number of measures for this. It should be expected that PM outperforms ACC with respect to such a measure, as PM actively seeks an even sound-field (assuming that the desired vector is chosen in a manner similar to the one discussed). Since this is not the topic of the thesis, it has been chosen not to consider the matter. It is, however, important to remark that ACC is not necessarily the superior method just because it obtains a higher acoustic contrast than PM.

3.3.2 Comparison of Performance

We compare the two discussed methods for the set-up introduced in Chapter 2 (see Table 2.1 on page 14 for the parameters of this set-up). A virtual source has been used to determine the desired vector, $\mathbf{d}$, for the PM method. This has been placed rather arbitrarily in the point $\mathbf{x}_{\mathrm{vs}} = [4.0 \text{ m}, 3.0 \text{ m}, 1.6 \text{ m}]^{\mathrm{T}}$. Some value has to be chosen for

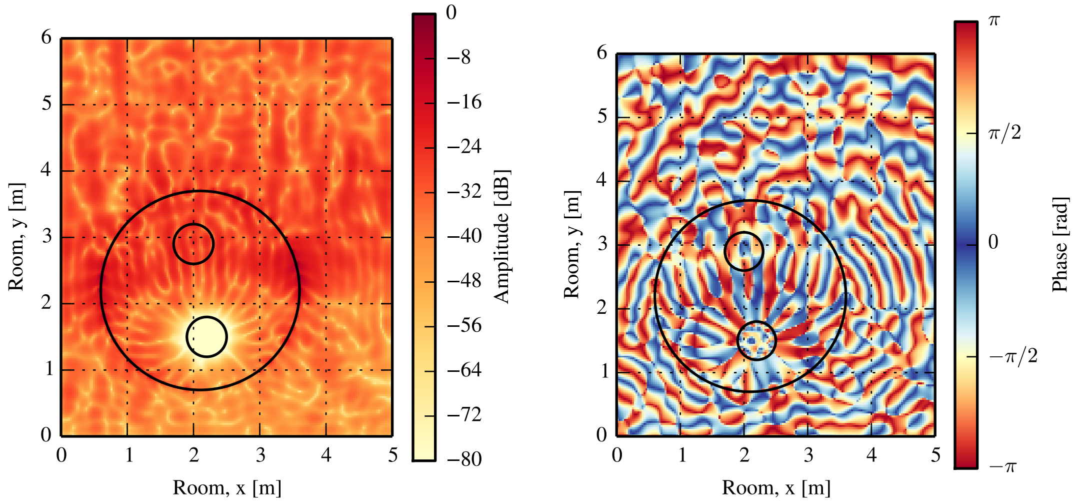

Figure 3.1: Amplitude (left) and phase (right) of the sound-field produced by ACC at 1000 Hz. The sound-field is plotted 1.6 m above the floor (in the plane where sources and microphones are placed). The used simulation environment is the one defined by Table 2.1. The bright zone is the upper one and the dark zone is the lower one. The figure is generated by the script `methods_single_frequency.py`.

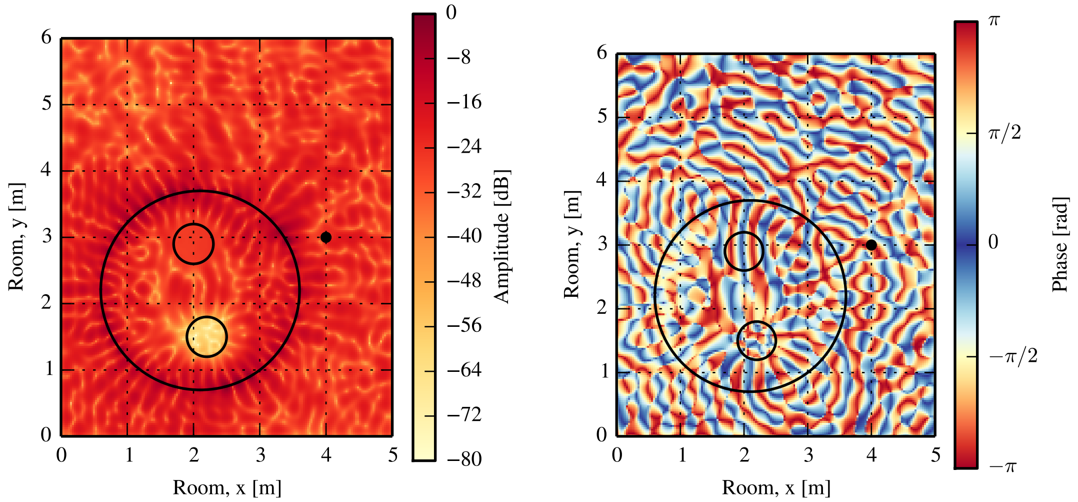

Figure 3.2: Amplitude (left) and phase (right) of the sound-field produced by PM at 1000 Hz. The sound-field in the bright zone is fitted to appear as if emanating from the virtual source marked by a black dot (the bright zone is the upper one and the dark zone is the lower one). The sound-field is plotted 1.6 m above the floor (in the plane where sources and microphones are placed). The used simulation environment is the one defined by Table 2.1. The figure is generated by the script `methods_single_frequency.py`.

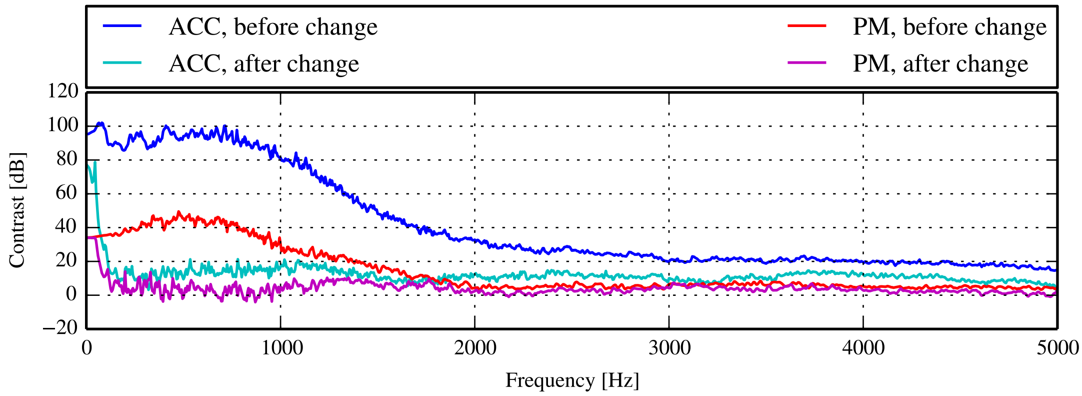

Figure 3.3: The acoustic contrast versus frequency. The contrast is shown before and after changing the parameters of the simulated room. The steering vectors are determined before changing the room conditions and are not re-computed afterwards. The figure is generated by the script `methods_single_frequency.py`.

the regularisation parameter, δ, as well. It should be stressed that this parameter is not a topic of investigation (a comparison of results for different values is given in [10]). Setting $\delta = 0$ may then seem as an obvious choice. For this case, however, the optimisation problem, (3.20), may not have a unique solution. We therefore set the parameter to a small value of 10^{-6} to guarantee uniqueness without imposing significant regularisation. Examples of the obtained sound-fields are shown in figures 3.1 and 3.2. Notice that both methods produce a significant difference in sound pressure between the bright zone and the dark zone. At the same time it is easily visible that ACC obtains a much larger contrast than PM. On the other hand, the sound pressure varies much throughout the bright zone for ACC and the phase appears somewhat disorganised. With PM, the sound-field is almost constant throughout the bright zone and the phase clearly resembles that of a wave emanating from the virtual source (as marked by a black dot).

The contrast is of course highly dependent on the frequency chosen for the system. The obtained acoustic contrast is plotted as a function of frequency in Figure 3.3. The contrast is plotted both before and after a change in the simulation environment as defined in Section 2.3.4 on page 22. As is expected, the highest contrast is obtained with ACC for all frequencies. PM on the other hand has the ability to determine the structure of the sound-field in the bright zone. It is seen that the contrast decreases significantly for both methods when the simulation environment is changed. This is not surprising, as the ACC and PM steering vectors are optimised to perform well under the initial room characteristics. It has been observed that the sensitivity of the contrast to changes in the room conditions increase with increasing reverberation time. It may be noted that the difference in contrast between the methods is much smaller with the changed conditions.

3.4 Summary

This chapter presents and investigates the two single frequency sound zone methods ACC and PM. Both of the methods take the approach of formulating the sound zone problem as an optimisation problem. ACC maximises the ratio of potential acoustical energy in the bright zone to potential acoustical energy in the dark zone (referred the acoustic contrast). PM maximises the similarity between the obtained sound-field and a desired one. The methods are compared by use of simulations. This reveals that ACC obtains the largest acoustical contrast while PM has the ability of controlling the sound-field in the bright zone.

Chapter 4

Wideband Methods

Chapter 3 introduces two of the most common approaches for realising sound zones. These are, however, designed to work well for sinusoidal inputs of one specific frequency. This is far from ideal when the task is to create a system for human listening. Almost no research has been carried out with respect to systems for wideband signals such as music or speech. In this chapter we discuss a number of approaches for creating sound zone systems for wideband signals. The presented methods are primarily based on intuitive extensions of the single frequency methods. The analysis of these extensions reveals that the problem of creating wideband sound zones is not as simple as it may appear. The analysis is carried out in an exploratory manner and leads to a variety of different results. An initial overview of the methods are given in Section 4.1. Furthermore, an overview of essential notation is given in Section 4.2. In sections 4.3 and 4.4 we investigate two simple possibilities for extending ACC to wideband signals. In Section 4.5 we investigate a simple approach for extending PM to wideband signals. In Section 4.6, a small numerical investigation of circular convolution is performed. In sections 4.7 and 4.8 we investigate an approach for improving the performance of the discussed wideband PM method.

4.1 Overview

As has been indicated in chapters 1 and 3, a rather significant amount of research has been carried out with respect to sound zones. However, the research relies almost exclusively on the assumption that the input signal is a single sinusoid with fixed and known frequency and that each input filter can therefore be described by a single complex coefficient (representing a gain and phase at that particular frequency). If a sound zone system is to be useful for human listening, it should work for wideband input signals such as speech and music. Only few high quality publications discuss the matter of sound zones for wideband signals. One proposal for applying PM to wideband signals is discussed in [19]. A very similar approach for wideband ACC is discussed in [17]. Both papers discuss the matter of wideband sound zones only as a short side note to other main topics. Due to the limited previous research on the subject, we carry out an exploratory investigation rather than a detailed review of existing methods. That is, we present a sequence of different ideas for creating wideband sound zones and discuss the properties and issues connected with these methods. This highlights some of the issues connected with extending the known sound zone methods to wideband signals.

Many of the properties which characterise a good method for generating sound zones are the same whether one discusses single frequency signals or wideband signals. One can may therefore consult Chapter 3 or relevant scientific literature for such matters [8, 10, 18]. A clear distinction, however, comes from the fact that all the properties of a method are generally frequency dependent in the wideband case. This introduces the additional complication of maintaining a desirable performance across a range of frequencies. Such matters easily lead to complicated psychoacoustic considerations. This thesis mainly focusses on the signal processing aspects of sound zones, and the topic of psychoacoustics is not as such considered. We therefore aim for sound zone methods which reproduce the input signal faithfully in mathematical terms. Most prominently, we deem that a method should not "colour" the input signal by transferring some frequencies better to the bright zone than others. To avoid complicated discussions of psychoacoustics, we continue to use acoustic contrast between the bright and dark zones as the primary means of evaluating the quality of a method.

We investigate four different approaches, two which are based on ACC and two which are based on PM. The methods have all been given names and abbreviations to ease referencing. A preliminary overview of the discussed methods are given:

- Multi Frequency ACC (MF-ACC), Section 4.3: This method attempts to treat a wideband system simply as a sequence of fully independent single frequency systems. This approach is introduced in [17].

- Time Domain ACC (TD-ACC), Section 4.4: This method takes an approach similar to ACC, but is formulated in the time domain. The section deviates slightly from the flow of the chapter, but serves as an illustration of some problems which may be encountered when attempting to devise methods for wideband sound zones. The section involves somewhat extensive derivations but can be skipped without any major loss. This approach has not previously been discussed in the literature.

- Multi Frequency PM (MF-PM), Section 4.5: This method is based on treating a wideband system as a sequence of single frequency systems in the same way as MF-ACC. The only difference is the fact that the method uses PM on each of the single frequency systems, rather than ACC. This is very similar to the approach discussed in [19].

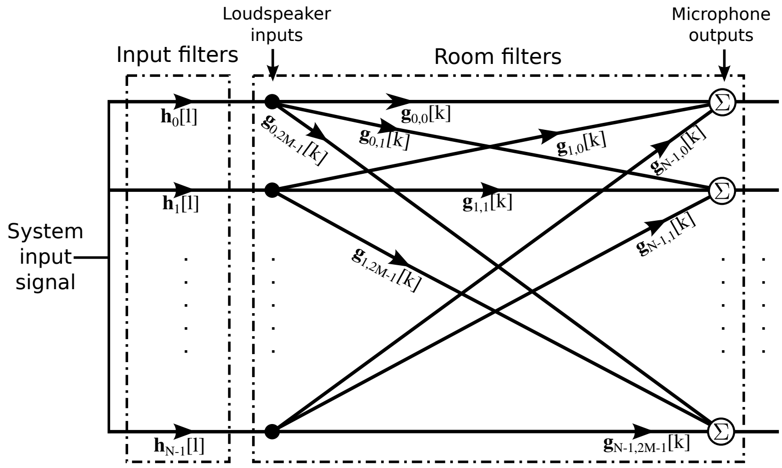

Figure 4.1: A schematic diagram of the sound zone system. A single input is provided to the system. The input signal for each of the N loudspeakers is generated by filtering this single input signal. The output at the $2M$ microphones is computed by filtering the loudspeaker inputs with a matrix of transfer functions. This matrix comprises room acoustics as well as acoustics of loudspeakers and microphones. The system contains $2M$ microphones to signify that there are two zones, each of which contains M microphones.

- Conjugate Gradient PM (CG-PM), Section 4.7 (and 4.8): This method is a modification of MF-PM. The method has a number of advantageous properties in comparison to MF-PM. The downside to the method is the fact that the involved optimisation problem is more computationally complex to solve. The method is referred to as as CG-PM due to the fact that the Conjugate Gradient (CG) Algorithm is used to solve the involved optimisation problem, rather than direct matrix inversion as is the case with MF-PM. This approach has not previously been discussed in the literature.

4.2 Notation

For convenience, we redisplay Figure 2.1 as Figure 4.1. The figure illustrates the overall sound zone system. The system consists of N input filters, which we assume to be Finite Impulse Response (FIR) filters given by:

$$\mathbf{h}_n = [\mathbf{h}_n[0], \ldots, \mathbf{h}_n[L-1]]^{\mathrm{T}}, \qquad n = 0, \ldots, N-1, \tag{4.1}$$

where L is the length of the impulse response for each input filter. A single input signal is filtered by the N input filters and the N resulting signals are further filtered by a matrix of room filters which represent the transfer functions from N source inputs to $2M$ microphone outputs, M in each sound zone. These filters are given by:

$$\mathbf{g}_{n,m} = [\mathbf{g}_{n,m}[0], \ldots, \mathbf{g}_{n,m}[K-1]]^{\mathrm{T}}, \qquad n = 0, \ldots, N-1, \;\; m = 0, \ldots, 2M-1, \tag{4.2}$$

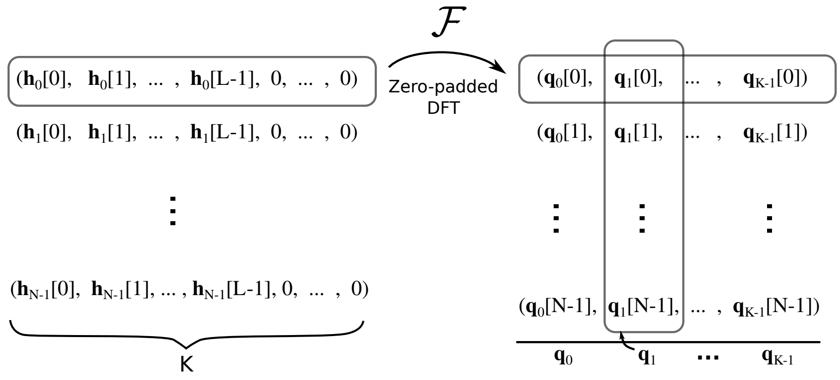

Figure 4.2: An illustration of the relationship between the input filters, $\mathbf{h}_0, \ldots, \mathbf{h}_{N-1}$, and the vectors $\mathbf{q}_0, \ldots, \mathbf{q}_{K-1}$. On the left, the filter coefficients of each input filter are arranged in rows which are zero-padded with $K-L$ zeros. By taking the DFT of a row, we obtain a corresponding row on the right. The vector $\mathbf{q}_k$ is given by the k'th columns of the structure on the right. It comprises the behaviour of all the N input filters at the k'th frequency.

where K is the length of the room filters. The responses $\mathbf{g}_{n,m}$ for $m = 0, \ldots, M-1$ represent the transfer functions into the bright zone while the responses for $m = M, \ldots, 2M-1$ represent the transfer functions into the dark zone.

We allow ourselves to assume that the input filters are at most as long as the room filters, that is $L \leq K$. This is advantageous in some of the following derivations and the results shown in Chapter 5 indicate that it is not a very restrictive limitation in practice. Since much of the work is carried out in the frequency domain, we introduce frequency domain representations of $\mathbf{h}_n$ and $\mathbf{g}_{n,m}$. We represent these in the form:

$$\mathbf{q}_k = [\mathbf{q}_k[0], \ldots, \mathbf{q}_k[N-1]]^{\mathrm{T}}, \tag{4.3}$$

$$\mathbf{B}_k = \begin{bmatrix} \mathbf{B}_k[0,0] & \cdots & \mathbf{B}_k[0,N-1] \\ \vdots & \ddots & \vdots \\ \mathbf{B}_k[M-1,0] & \cdots & \mathbf{B}_k[M-1,N-1] \end{bmatrix}, \tag{4.4}$$

$$\mathbf{D}_k = \begin{bmatrix} \mathbf{D}_k[0,0] & \cdots & \mathbf{D}_k[0,N-1] \\ \vdots & \ddots & \vdots \\ \mathbf{D}_k[M-1,0] & \cdots & \mathbf{D}_k[M-1,N-1] \end{bmatrix}, \tag{4.5}$$

where:

$$\mathbf{q}_k[n] = \frac{1}{\sqrt{K}} \sum_{\ell=0}^{L-1} \mathbf{h}_n[\ell] \exp\left(-j2\pi\frac{k\ell}{K}\right), \tag{4.6}$$

$$\mathbf{B}_k[m,n] = \frac{1}{\sqrt{K}} \sum_{\ell=0}^{K-1} \mathbf{g}_{n,m}[\ell] \exp\left(-j2\pi\frac{k\ell}{K}\right), \tag{4.7}$$

$$\mathbf{D}_k[m,n] = \frac{1}{\sqrt{K}} \sum_{\ell=0}^{K-1} \mathbf{g}_{n,(m+M)}[\ell] \exp\left(-j2\pi\frac{k\ell}{K}\right), \tag{4.8}$$

for $k = 0, \ldots, K-1$, $n = 0, \ldots, N-1$ and $m = 0, \ldots, M-1$. The vector $\mathbf{q}_k \in \mathbb{C}^{N\times 1}$ represents the frequency response of all N input filters at the k'th point of a K point DFT, that is at the frequency $f_\mathrm{s}\frac{k}{K}$ where f_s is the sampling frequency. The reason for zero padding the input filters to length K is that this makes it straightforward to compare the response of the input filters to the response of the room filters, frequency sample by frequency sample. The relationship between $\mathbf{q}_k$ and $\mathbf{h}_n$ is illustrated in Figure 4.2. The matrix $\mathbf{B}_k \in \mathbb{C}^{M\times N}$ represents the response of each microphone in the bright zone to each soruce at the k'th frequency point. Similarly, the matrix $\mathbf{D}_k \in \mathbb{C}^{M\times N}$ represents the response of each microphone in the dark zone to each source at the k'th frequency point. The total response of the bright and dark zone microphones to the input signal at the k'th frequency may be written as the following two vectors:

$$\begin{aligned}\mathbf{p}_{\mathrm{b},k} &= \mathbf{B}_k\mathbf{q}_k,\\ \mathbf{p}_{\mathrm{d},k} &= \mathbf{D}_k\mathbf{q}_k.\end{aligned} \tag{4.9}$$

The values $\mathbf{p}_{\mathrm{b},k}[0], \ldots, \mathbf{p}_{\mathrm{b},k}[M-1]$, $\mathbf{p}_{\mathrm{d},k}[0], \ldots, \mathbf{p}_{\mathrm{d},k}[M-1]$ corresponds to the outputs on the right of Figure 4.1 from top to bottom at the frequency $\frac{k}{K}f_\mathrm{s}$. This notation may appear unnecessarily complicated, but it turns out to be very useful[1].

4.3 Simple Extension of ACC to Multiple Frequencies

The first wideband sound zone method we discuss is introduced in [17]. It is a straightforward extension of single frequency ACC. The objective is to design a method which has the same properties as the single frequency ACC, but works for multiple frequencies. We have termed the method Multi Frequency ACC (MF-ACC). Before describing the method mathematically, we shortly present the underlying idea.

[1]Throughout the chapter we make heavy use of block matrix notation. It should be remarked that most of the block matrices could possibly be replaced by third order tensors, leading to a simpler notation. This has, however, been avoided as higher order tensors are rather uncommon in the signal processing literature.

In single frequency ACC (Section 3.1 on page 27), one decides a particular frequency at which the system should work well, and obtains the input filters by solving an optimisation problem. Since the system works at only a single frequency, each input filter can be specified by a single complex coefficient (representing amplitude and phase of the filter at that frequency). Solving the same problem for a different frequency yields a different set of complex filter coefficients. We may solve the problem for frequencies $\frac{0}{K}f_\mathrm{s}$, $\frac{1}{K}f_\mathrm{s}$, ... , $\frac{K-1}{K}f_\mathrm{s}$. This specifies the amplitude and phase for all the input filters for an evenly spaced grid of frequencies. This essentially specifies a discrete K-point frequency response for each of the input filters. By computing an Inverse DFT (IDFT) of each filter we obtain time domain representations of the input filters (of length K). By designing the input filters like this, the method obtains the same acoustic contrast as a single frequency system for sinusoidal input signals of frequencies $\frac{0}{K}f_\mathrm{s}$, $\frac{1}{K}f_\mathrm{s}$, ... , $\frac{K-1}{K}f_\mathrm{s}$ simultaneously. The method, however, is not guaranteed to perform well at other frequencies.

The above assumes that one can essentially decompose a wideband sound zone system into a sequence of single frequency systems. It should be noted that we have assumed that $L = K$. This assumption is mathematically necessary but otherwise undesirable. Measured or simulated room impulse responses may be very long (i.e. consist of many samples). It may not be desirable to have input filters of this length. Long FIR filter are computationally expensive to implement and there is no guarantee that long filters lead to high performance. The case of $L = K$ corresponds to the case where there are no zeros on the left hand side of Figure 4.2. The vector $\mathbf{q}_k \in \mathbb{C}^{N\times 1}$ comprises the frequency response of the N input filters at the frequency $\frac{k}{K}f_\mathrm{s}$. We solve the optimisation problem of single frequency ACC for each of the frequencies $\frac{0}{K}f_\mathrm{s}$, $\frac{1}{K}f_\mathrm{s}$, ... , $\frac{K-1}{K}f_\mathrm{s}$. This gives rise to the following K independent optimisation problems:

$$\begin{aligned}\mathbf{q}_{k,\mathrm{ACC}} = \quad & \underset{\mathbf{q}_k}{\arg\min}\ \mathbf{q}_k^\mathrm{H}\widehat{\mathbf{R}}_{\mathrm{d},k}\mathbf{q}_k \\ & \text{subject to } \mathbf{q}_k^\mathrm{H}\widehat{\mathbf{R}}_{\mathrm{b},k}\mathbf{q}_k = e_k,\end{aligned} \tag{4.10}$$

where:

$$\begin{aligned}\widehat{\mathbf{R}}_{\mathrm{b},k} &= \mathbf{B}_k^\mathrm{H}\mathbf{B}_k, \\ \widehat{\mathbf{R}}_{\mathrm{d},k} &= \mathbf{D}_k^\mathrm{H}\mathbf{D}_k.\end{aligned} \tag{4.11}$$

for $k = 0, \ldots, K-1$. The procedure for solving the problem is the same as the one discussed in Chapter 3. Note that the parameter which determines the amount of acoustic energy transferred to the bright zone, e_k, is frequency dependent. This makes it possible to introduce frequency selectivity into the system (such that some frequencies are reproduced with more energy in the bright zone than others). Here, we aim to recreate the input signal as faithfully as possible, and we therefore select a flat frequency response with $e_0 = e_1 = \ldots = e_{K-1}$. The time domain filters are obtained by IDFT of the frequency domain filter coefficients, $\mathbf{q}_{0,\mathrm{ACC}}, \ldots, \mathbf{q}_{K-1,\mathrm{ACC}}$. The overall method for determining the input filters is illustrated by Algorithm 4.1. As discussed in Chapter 3,

Algorithm 4.1 MF-ACC

input: $\mathbf{g}_{n,m} \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$ and $m = 0, \ldots, 2M-1$, $e_k \in \mathbb{R}$ for $k = 0, \ldots, K-1$
output: $\mathbf{h}_n \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$

\# Find filter responses
for $k = 0$ to $K-1$:
    Form $\mathbf{D}_k$ and $\mathbf{B}_k$ by (4.7) and (4.8)
    $\widehat{\mathbf{R}}_{\mathrm{b},k} \leftarrow \mathbf{B}_k^{\mathrm{H}}\mathbf{B}_k$
    $\widehat{\mathbf{R}}_{\mathrm{d},k} \leftarrow \mathbf{D}_k^{\mathrm{H}}\mathbf{D}_k$
    Find principal eigenvector, $\mathbf{v}$, of $(\widehat{\mathbf{R}}_{\mathrm{d},k})^{-1}\widehat{\mathbf{R}}_{\mathrm{b},k}$
    $\mathbf{q}_{k,\mathrm{ACC}} \leftarrow \mathbf{v}\sqrt{\frac{e_k}{\mathbf{v}^{\mathrm{H}}\widehat{\mathbf{R}}_{\mathrm{d},k}\mathbf{v}}}$

\# Obtain time domain filters by IDFT
for $n = 0$ to $N-1$:
    $\mathbf{h}_n \leftarrow$ IDFT of $[\,\mathbf{q}_{0,\mathrm{ACC}}[n], \mathbf{q}_{1,\mathrm{ACC}}[n], \ldots, \mathbf{q}_{K-1,\mathrm{ACC}}[n]\,]^{\mathrm{T}}$

the primary advantage of ACC is its ability to obtain a large acoustic contrast between the zones. Thus, by design, MF-ACC is able to obtain a large acoustic contrast for sinusoidal inputs at the frequencies $\frac{0}{K}f_{\mathrm{s}}, \frac{1}{K}f_{\mathrm{s}}, \ldots, \frac{K-1}{K}f_{\mathrm{s}}$. The acoustic contrast at these frequencies is computed in exactly the same manner as for the single frequency case (see (3.9) on page 29):

$$\chi\left(\frac{k}{K}f_{\mathrm{s}}\right) = \frac{\mathbf{q}_k^{\mathrm{H}}\widehat{\mathbf{R}}_{\mathrm{b},k}\mathbf{q}_k}{\mathbf{q}_k^{\mathrm{H}}\widehat{\mathbf{R}}_{\mathrm{d},k}\mathbf{q}_k}. \tag{4.12}$$

Intuitively, one may expect the method to perform approximately equally well between the frequency points. We can investigate whether this is the case by computing the acoustic contrast for a sinusoidal input of an arbitrary frequency, f. The total impulse response from the input to the m'th microphone is given by the sum of N convolution products (this is seen from Figure 4.1):

$$\mathbf{y}_m = \sum_{n=0}^{N-1} \mathbf{h}_n * \mathbf{g}_{n,m} \in \mathbb{R}^{(K+L-1)\times 1}, \tag{4.13}$$

where "$*$" denotes linear convolution. Note that in the particular case of MF-ACC, we have $L = K$, and thus $K+L-1 = 2K-1$. To compute the frequency dependent acoustic contrast without being constrained to any particular DFT-grid, we consider the Discrete

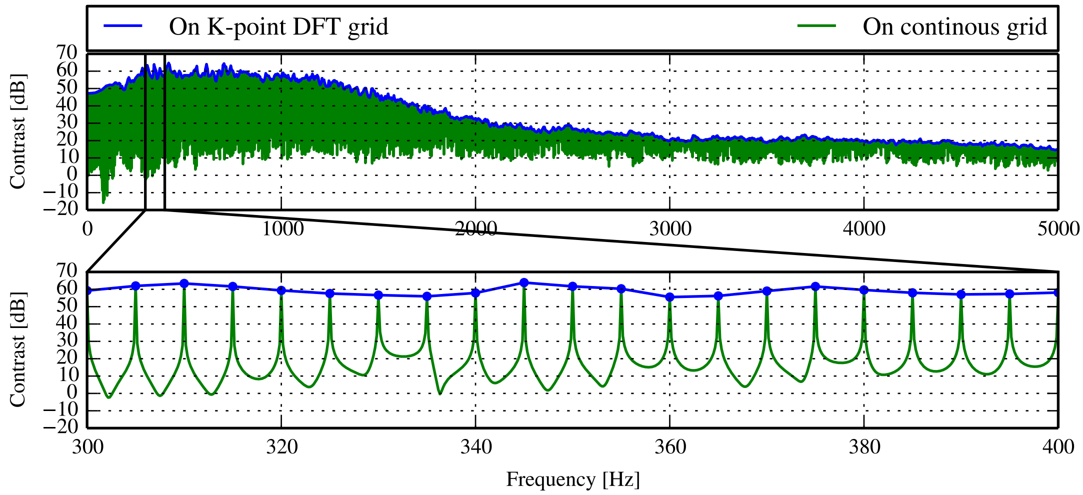

Figure 4.3: Upper: A plot of the obtained contrast for multi-frequency ACC. The results have been obtained by performing simulations on the set-up introduced in Chapter 2. The blue points show the contrast obtained on the grid where the contrast control is performed (see (4.12)). The individual frequency points have been connected by straight lines to ease viewing. The contrast is also plotted for a more dense grid of frequencies (see (4.15)). The green line appears smeared because the contrast changes rapidly with frequency. Lower: A zoom on a small region of the upper plot. The figure is generated by the script `time_domain_aliasing.py`.

Time Fourier Transform (DTFT) of the response [24]:

$$Y_m(f) = \sum_{\ell=0}^{K+L-1} \mathbf{y}_m[\ell] \exp\left(-j2\pi \frac{f}{f_\mathrm{s}} \ell\right). \tag{4.14}$$

The DTFT specifies the response of the system to sinusoidal inputs of arbitrary frequency. We can therefore use it to define the frequency dependent acoustic contrast for sinusoidal inputs of any frequncy:

$$\chi(f) = \frac{\sum_{m=0}^{M-1} |Y_m(f)|^2}{\sum_{m=M}^{2M-1} |Y_m(f)|^2}. \tag{4.15}$$

The above coincides with (4.12) for the K points on the DFT grid. This allows us to investigate the acoustical contrast for other frequencies than $\frac{0}{K} f_\mathrm{s}$, $\frac{1}{K} f_\mathrm{s}$, $\dots$, $\frac{K-1}{K} f_\mathrm{s}$.
At this point, we can investigate the method by use of simulations. Most of the simulation results are presented in Chapter 5, but we present a single plot here, as it is important to the remainder of the chapter. The MF-ACC approach has been used with the simulated set-up introduced in Chapter 2. The obtained contrast is shown in Figure 4.3. This shows that the contrast is much larger on the DFT grid than at the points in between. This is a serious downside to the method, as wideband audio signals should generally be

assumed to have frequency content across the entire spectrum, including the frequencies where the method performs poorly[2].

At this point, we see interesting but somewhat discouraging results of the MF-ACC method. We may ask ourselves if the uneven performance of the method can be traced back to any particular feature of the way in which the sound zone problem has been approached. The issue may in principle stem from a multitude of causes. E.g. 1) some unseen mathematical flaw in the method or that the method is simply an overall poor approach, 2) an error in the software implementation which has produced Figure 4.3 or 3) even possibly an error in the derivations due to the author. We cannot as such prove that our subjective dissatisfaction with the results stems from any particular cause. We can, however, indicate one somewhat elusive problem with the method.

The impulse response from input to one of the $2M$ microphones are given as a sum of N convolution products, each product being the contribution of one path from input to output (see Figure 4.1). A well known property of the DFT is the fact that multiplying two sequences in the frequency domain is equivalent to convolving them in the time domain [24]. This is essentially the property underlying (4.9) where the DFT of the output of the M microphones in one zone or the other is computed at one frequency point on the K-point DFT grid. Since the output of M microphones is computed, there are a total of NM convolution products to carry out (that, is NM multiplications per frequency point in the frequency domain). In (4.9) this computation is carried out with a matrix multiplication for each frequency point. The subtle point to notice is the fact that there are two different types of discrete convolution. When carrying out a discrete convolution by multiplication in the frequency domain, such as is done in (4.9), one obtains a circular convolution [24]. Thus, circular convolution is used as an underlying assumption in the input filter design phase. On the other hand, when the resulting contrast is evaluated (in simulation or in reality), a linear convolution is used. This introduces a difference in the assumptions underlying the method and the way in which the result is evaluated. We hypothesise that this issue is connected to the uneven performance shown in Figure 4.3. It has, however, not been possible to find a way to circumvent the problem in MF-ACC.

4.4 *A Time Domain Approach to ACC

The problem of maximising the acoustic contrast between the two zones can also be formulated in the time domain. This leads to an approach which differs slightly from the above while providing an interesting alternative view of the contrast problem. The resulting method leads to interesting insights into some issues which can occur when

[2]In fairness to [17] it must be stressed that the problem is not necessarily equally pronounced for all systems. Especially, (simulated) systems with few sources and microphones have been found to be less affected. We do not mean to question the validity of the results presented in [17].

attempting to create wideband sound zones. On the contrary, the section does not lead to a viable sound zone method and does not contribute to the overall flow of the report. The section can therefore be skipped with no major loss to the overall flow of the thesis.

We provide the system (Figure 4.1) with an input signal of length Y:

$$\mathbf{x} = [\mathbf{x}[0], \ldots, \mathbf{x}[Y-1]]^{\mathrm{T}}, \tag{4.16}$$

The output of the system at the m'th microphone is given by:

$$\mathbf{u}_m = \sum_{n=0}^{N-1} \mathbf{x} * \mathbf{h}_n * \mathbf{g}_{n,m} \in \mathbb{R}^{(Y+L+N-2)\times 1}. \tag{4.17}$$

Alternatively, we may write this as:

$$\mathbf{u}_m = \sum_{n=0}^{N-1} \mathbf{X}\boldsymbol{\Phi}_{n,m}\mathbf{h}_n = \mathbf{X}\sum_{n=0}^{N-1} \left(\boldsymbol{\Phi}_{n,m}\mathbf{h}_n\right), \tag{4.18}$$

where $\mathbf{X} \in \mathbb{R}^{(Y+L+N-2)\times(L+N-1)}$ and $\boldsymbol{\Phi}_{n,m} \in \mathbb{R}^{(L+N-1)\times L}$ are convolution operators implementing convolution with $\mathbf{x}$ and $\mathbf{g}_{n,m}$ respectively. Recall that the microphones $m = 0, \ldots, M-1$ are positioned in the bright zone while those for $m = M, \ldots, 2M-1$ are positioned in the dark zone. The acoustic potential energy in the bright and dark zone is proportional to (this is defined similarly to (3.6) on page 28, but with somewhat different notation):

$$e_{\mathrm{b}} = \sum_{m=0}^{M-1} \mathbf{u}_m^{\mathrm{T}}\mathbf{u}_m = \sum_{m=0}^{M-1}\sum_{n=0}^{N-1} \left(\mathbf{h}_n^{\mathrm{T}}\boldsymbol{\Phi}_{n,m}^{\mathrm{T}}\right)\mathbf{X}^{\mathrm{T}}\mathbf{X}\sum_{n=0}^{N-1}\left(\boldsymbol{\Phi}_{n,m}\mathbf{h}_n\right), \tag{4.19}$$

$$e_{\mathrm{d}} = \sum_{m=M}^{2M-1} \mathbf{u}_m^{\mathrm{T}}\mathbf{u}_m = \sum_{m=M}^{2M-1}\sum_{n=0}^{N-1} \left(\mathbf{h}_n^{\mathrm{T}}\boldsymbol{\Phi}_{n,m}^{\mathrm{T}}\right)\mathbf{X}^{\mathrm{T}}\mathbf{X}\sum_{n=0}^{N-1}\left(\boldsymbol{\Phi}_{n,m}\mathbf{h}_n\right). \tag{4.20}$$

It may be noted that $\mathbf{X}^{\mathrm{T}}\mathbf{X} \in \mathbb{R}^{(L+N-1)\times(L+N-1)}$ is essentially an approximation of the correlation matrix of the input signal, $\mathbf{x}$ (to within a scaling factor, that is). We therefore term this $\mathbf{R}_{\mathbf{xx}}$. Since only $\mathbf{R}_{\mathbf{xx}}$ is needed, we do not need to know the actual input signal. We can do with a stochastic model of it instead. We can rewrite the above with a block matrix structure:

$$e_{\mathrm{b}} = \tilde{\mathbf{h}}^{\mathrm{T}}\widetilde{\boldsymbol{\Phi}}_{\mathrm{b}}^{\mathrm{T}}\widetilde{\mathbf{R}}_{\mathbf{xx}}\widetilde{\boldsymbol{\Phi}}_{\mathrm{b}}\tilde{\mathbf{h}} \tag{4.21}$$

$$e_{\mathrm{d}} = \tilde{\mathbf{h}}^{\mathrm{T}}\widetilde{\boldsymbol{\Phi}}_{\mathrm{d}}^{\mathrm{T}}\widetilde{\mathbf{R}}_{\mathbf{xx}}\widetilde{\boldsymbol{\Phi}}_{\mathrm{d}}\tilde{\mathbf{h}} \tag{4.22}$$

where:

$$\widetilde{\mathbf{\Phi}}_{\mathrm{b}} = \begin{bmatrix} \mathbf{\Phi}_{0,0} & \cdots & \mathbf{\Phi}_{N-1,0} \\ \vdots & \ddots & \vdots \\ \mathbf{\Phi}_{0,M-1} & \cdots & \mathbf{\Phi}_{N-1,M-1} \end{bmatrix} \in \mathbb{R}^{(L+N-1)N \times LN}, \tag{4.23}$$

$$\widetilde{\mathbf{\Phi}}_{\mathrm{d}} = \begin{bmatrix} \mathbf{\Phi}_{0,M} & \cdots & \mathbf{\Phi}_{N-1,M} \\ \vdots & \ddots & \vdots \\ \mathbf{\Phi}_{0,2M-1} & \cdots & \mathbf{\Phi}_{N-1,2M-1} \end{bmatrix} \in \mathbb{R}^{(L+N-1)N \times LN}, \tag{4.24}$$

and:

$$\tilde{\mathbf{h}} = \begin{bmatrix} \mathbf{h}_0 \\ \vdots \\ \mathbf{h}_{N-1} \end{bmatrix} \in \mathbb{R}^{NL \times 1}, \tag{4.25}$$

$$\widetilde{\mathbf{R}}_{\mathbf{xx}} = \begin{bmatrix} \mathbf{R}_{\mathbf{xx}} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\mathbf{xx}} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_{\mathbf{xx}} \end{bmatrix} \in \mathbb{R}^{NL \times NL}. \tag{4.26}$$

These observations lead us to the following formulation of the contrast problem:

$$\begin{aligned} \tilde{\mathbf{h}}_{\mathrm{ACC}} = \quad & \arg\min_{\tilde{\mathbf{h}}} \tilde{\mathbf{h}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}} \tilde{\mathbf{h}} \\ & \text{subject to } \mathbf{h}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}} \mathbf{h} = e_0, \end{aligned} \tag{4.27}$$

where e_0 determines the total amount of acoustic energy in the bright zone (by some constant of proportionality which is determined by the physical layout of the array). Described in words, the problem is to pick the set of input filters which minimises the energy in the dark zone while keeping the energy in the bright zone at the level e_0. Contrary to the MF-ACC method, this approach includes a model of the input signal. The problem is solved in exactly the same manner as in the single frequency case. That is, $\tilde{\mathbf{h}}_{\mathrm{ACC}}$ is the principal eigenvector of $(\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}})^{-1} \widetilde{\mathbf{\Phi}}_{\mathrm{b}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}}$. It should be noted that the matrix, $\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}$, which has to be inverted, belongs to the space $\mathbb{R}^{NL \times NL}$ where N is the number of sources and L is the length of the input filters. This matrix is potentially too large to be inverted in practice. This may pose a problem for the method, but as we show below, the method suffers from a much more serious issue.

We can illustrate the type of solution obtained from (4.27) without resorting to numerical simulations. This is done by transforming the problem to the frequency domain. We do this for the particular case of the problem where $\mathbf{R}_{\mathbf{xx}} = \mathbf{I}$ and $L = K$. This is the case where the input signal is assumed to stem from a white process and the input filters have the same length as the room filters. We first look at the structure of the objective function of (4.27), in the particular case considered. This is simply the energy in the dark zone (to within some constant of proportionality):

$$e_{\mathrm{d}} = \tilde{\mathbf{h}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{R}}_{\mathbf{xx}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}} \tilde{\mathbf{h}} = \tilde{\mathbf{h}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}} \tilde{\mathbf{h}} \tag{4.28}$$

The matrix $\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}$ is a block matrix which is densely filled with circulant sub-matrices (circulant matrices are shortly discussed in Appendix B):

$$\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}} = \begin{bmatrix} \sum_m \mathbf{\Phi}_{m,0}^{\mathrm{T}} \mathbf{\Phi}_{m,0} & \cdots & \sum_m \mathbf{\Phi}_{m,N-1}^{\mathrm{T}} \mathbf{\Phi}_{m,0} \\ \vdots & \ddots & \vdots \\ \sum_m \mathbf{\Phi}_{m,N-1}^{\mathrm{T}} \mathbf{\Phi}_{m,0} & \cdots & \sum_m \mathbf{\Phi}_{m,N-1}^{\mathrm{T}} \mathbf{\Phi}_{m,N-1} \end{bmatrix} \in \mathbb{R}^{NL \times NL}, \tag{4.29}$$

where the summations run over $m = 0, \ldots, M-1$. A circulant matrix is diagonalised by the DFT [41]. To take advantage of this, we introduce the operator:

$$\tilde{\mathbf{F}} = \begin{bmatrix} \boldsymbol{\mathcal{F}} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\mathcal{F}} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \boldsymbol{\mathcal{F}} \end{bmatrix} \in \mathbb{C}^{NL \times NL}. \tag{4.30}$$

Each $\boldsymbol{\mathcal{F}}$ is an L-point DFT. Since the DFT is unitary, we have:

$$\tilde{\mathbf{F}}^{\mathrm{H}} \tilde{\mathbf{F}} = \mathbf{I}. \tag{4.31}$$

Furthermore (in the considered case of $L = K$):

$$\tilde{\mathbf{F}} \tilde{\mathbf{h}} = \begin{bmatrix} \boldsymbol{\mathcal{F}} \mathbf{h}_0 \\ \vdots \\ \boldsymbol{\mathcal{F}} \mathbf{h}_{N-1} \end{bmatrix}. \tag{4.32}$$

Notice that $\tilde{\mathbf{F}}\tilde{\mathbf{h}}$ has exactly the same entries as $\tilde{\mathbf{q}}$. The difference is that the entries in $\tilde{\mathbf{F}}\tilde{\mathbf{h}}$ are grouped according to filters (value of n) while the entries in $\tilde{\mathbf{q}}$ are grouped according to frequency (value of k).[3] Another important property of $\tilde{\mathbf{F}}$ is its ability to perform block-wise diagonalisation of $\widetilde{\mathbf{\Phi}}_{\mathrm{b}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}}$ and $\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{d}}$. These are block-circulant matrices (by (4.29)). Since circulant matrices are diagonalised by the DFT [41, 42], we have:

$$\widetilde{\mathbf{\Lambda}}_{\mathrm{b}} = \tilde{\mathbf{F}}^{\mathrm{H}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}}^{\mathrm{T}} \widetilde{\mathbf{\Phi}}_{\mathrm{b}} \tilde{\mathbf{F}} = \begin{bmatrix} \mathbf{\Lambda}_{\mathrm{b},0,0} & \cdots & \mathbf{\Lambda}_{\mathrm{b},0,N-1} \\ \vdots & \ddots & \vdots \\ \mathbf{\Lambda}_{\mathrm{b},N-1,0} & \cdots & \mathbf{\Lambda}_{\mathrm{b},N-1,N-1} \end{bmatrix}, \tag{4.33}$$

[3] For a more clear view of this, consider again Figure 4.2. The vector $\tilde{\mathbf{h}}$ contains all the elements on the left of the figure, stacked one row at a time. The application of $\tilde{\mathbf{F}}$ corresponds to row-wise DFT. Therefore, $\tilde{\mathbf{F}}\tilde{\mathbf{h}}$ contains all the elements on the right of the figure, stacked one row at a time. On the other hand, $\tilde{\mathbf{q}}$ corresponds to all the elements on the right of the figure, stacked one column at a time.

where $\mathbf{\Lambda}_{\mathrm{b},a,b} \in \mathbb{C}^{K\times K}$ are diagonal matrices. A similar matrix, $\widetilde{\mathbf{\Lambda}}_{\mathrm{d}}$, is defined for the dark zone. It can be seen that $\tilde{\mathbf{F}}^{\mathrm{H}}\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}}\widetilde{\mathbf{\Phi}}_{\mathrm{d}}\tilde{\mathbf{F}}$ has the same entries as $\widetilde{\mathbf{R}}_{\mathrm{d}}^{\mathrm{H}}\widetilde{\mathbf{R}}_{\mathrm{d}}$, but is arranged differently. By (4.31) we can rewrite (4.27) to:

$$\begin{aligned}\tilde{\mathbf{h}}_{\mathrm{ACC}} = \quad & \underset{\tilde{\mathbf{h}}}{\arg\min}\ \tilde{\mathbf{h}}^{\mathrm{T}}\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{F}}\widetilde{\mathbf{\Phi}}_{\mathrm{d}}^{\mathrm{T}}\widetilde{\mathbf{R}}_{\mathbf{xx}}\widetilde{\mathbf{\Phi}}_{\mathrm{d}}\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{F}}\tilde{\mathbf{h}} \\ & \text{subject to } \mathbf{h}^{\mathrm{T}}\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{F}}\widetilde{\mathbf{\Phi}}_{\mathrm{b}}^{\mathrm{T}}\widetilde{\mathbf{R}}_{\mathbf{xx}}\widetilde{\mathbf{\Phi}}_{\mathrm{b}}\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{F}}\mathbf{h} = e_0.\end{aligned} \tag{4.34}$$

This corresponds simply to introducing two identity operators (by (4.31)). We rewrite this:

$$\begin{aligned}\tilde{\mathbf{h}}_{\mathrm{ACC}} = \quad & \underset{\tilde{\mathbf{h}}}{\arg\min}\ (\tilde{\mathbf{F}}\tilde{\mathbf{h}})^{\mathrm{H}}\widetilde{\mathbf{\Lambda}}_{\mathrm{d}}(\tilde{\mathbf{F}}\tilde{\mathbf{h}}) \\ & \text{subject to } (\tilde{\mathbf{F}}\tilde{\mathbf{h}})^{\mathrm{H}}\widetilde{\mathbf{\Lambda}}_{\mathrm{b}}(\tilde{\mathbf{F}}\tilde{\mathbf{h}}) = e_0.\end{aligned} \tag{4.35}$$

By rearranging the ordering of rows and columns in (4.35), we obtain:

$$\begin{aligned}\tilde{\mathbf{q}}_{\mathrm{ACC}} = \quad & \underset{\tilde{\mathbf{q}}}{\arg\min}\ \tilde{\mathbf{q}}^{\mathrm{H}}\widetilde{\mathbf{R}}_{\mathrm{d}}\tilde{\mathbf{q}} \\ & \text{subject to } \tilde{\mathbf{q}}^{\mathrm{H}}\widetilde{\mathbf{R}}_{\mathrm{b}}\tilde{\mathbf{q}} = e_0,\end{aligned} \tag{4.36}$$

because the relationship between $\tilde{\mathbf{q}}$ and $\tilde{\mathbf{F}}\tilde{\mathbf{h}}$ is a simple rearrangement of elements. The same is the case for $\widetilde{\mathbf{R}}_{\mathrm{b}}$ and $\widetilde{\mathbf{\Lambda}}_{\mathrm{b}}$ as well as for $\widetilde{\mathbf{R}}_{\mathrm{d}}$ and $\widetilde{\mathbf{\Lambda}}_{\mathrm{d}}$. Notice the strong similarity between the optimisation problems (4.36) and (4.10). Where (4.10) puts constrains on the energy at each frequency, (4.36) constrains only the total amount of energy without taking into account the way in which the energy is distributed in frequency. The solution to (4.36) is a principal eigenvector of $\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$. The matrix $\widetilde{\mathbf{R}}_{\mathrm{d}}$ is block diagonal and its inverse can therefore be constructed in a block-wise manner [43]:

$$\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1} = \begin{bmatrix} \mathbf{R}_{\mathrm{d},0}^{-1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\mathrm{d},1}^{-1} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_{\mathrm{d},K-1}^{-1} \end{bmatrix}. \tag{4.37}$$

Therefore we have:

$$\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}} = \begin{bmatrix} \mathbf{R}_{\mathrm{d},0}^{-1}\mathbf{R}_{\mathrm{b},0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\mathrm{d},1}^{-1}\mathbf{R}_{\mathrm{b},1} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_{\mathrm{d},K-1}^{-1}\mathbf{R}_{\mathrm{b},K-1} \end{bmatrix}. \tag{4.38}$$

The spectrum of a block diagonal matrix is simply the union of the spectra of the diagonal elements (see Appendix B for a proof of this). One can therefore say that an eigenvalue of

$\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$ stems from one particular block matrix on its diagonal. Consider an eigenvalue, λ, of the k'th block on the diagonal of $\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$ (that is, an eigenvalue of $\mathbf{R}_{\mathrm{d},k}^{-1}\mathbf{R}_{\mathrm{b},k}$). Take the eigenvector corresponding to λ to be $\mathbf{v}$. The value λ is an eigenvalue of both $\mathbf{R}_{\mathrm{d},k}^{-1}\mathbf{R}_{\mathrm{b},k}$ and $\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$ because $\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$ is block diagonal. The corresponding eigenvector of $\widetilde{\mathbf{R}}_{\mathrm{d}}^{-1}\widetilde{\mathbf{R}}_{\mathrm{b}}$ is given by (see Appendix B for a proof of this):

$$\tilde{\mathbf{v}} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{v} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \in \mathbb{C}^{NL \times 1}, \tag{4.39}$$

where the vector $\mathbf{v} \in \mathbb{C}^{N \times 1}$ is placed in entries Nk to $N(k+1)$ of $\tilde{\mathbf{v}}$ while all other entries are zero. This is the general type of solution which is found for the problem of investigation, (4.36). That is, the resulting input filters have non-zero DFT coefficients at only a single frequency! While this result may come as a surprise, it is actually intuitively sensible when reconsidering the interpretation the optimisation problem, (4.36). The task is to design input filters to minimise the energy in the dark zone while keeping the energy in the bright zone at some level, e_0. This is equivalent to maximising the acoustical contrast between the zones. We can imagine that this problem is solved in two steps: 1) locate the frequency at which the largest acoustical contrast can be obtained, and 2) design input filters which obtain this contrast at the found frequency while setting all other input frequency components to zero. While this strategy solves the posed optimisation problem, it clearly does not yield a desirable result. In the beginning of the chapter we argued that a good wideband method for sound zones should be capable of transferring acoustic energy to the bright zone equally well at all frequencies. This is a good example of an optimisation based approach which yields "the correct answer to the wrong question". Here, we considered only a special case of the problem. Taking $\mathbf{R}_{\mathbf{xx}} \neq \mathbf{I}$ corresponds to introducing an uneven weighting of frequency content, but does not otherwise change the outcome of the method. We refer to the discussed method as Time Domain ACC (TD-ACC) but do not further investigate it.

4.5 Simple Extension of PM to Multiple Frequencies

The sections above introduce two approaches for extending ACC to multiple frequencies. Both of the approaches are shown to suffer from problems. In this section we discuss a possible way of extending PM to multiple frequencies. The method is exactly the same as the one presented in Section 4.3, but for PM rather than for ACC. The method is

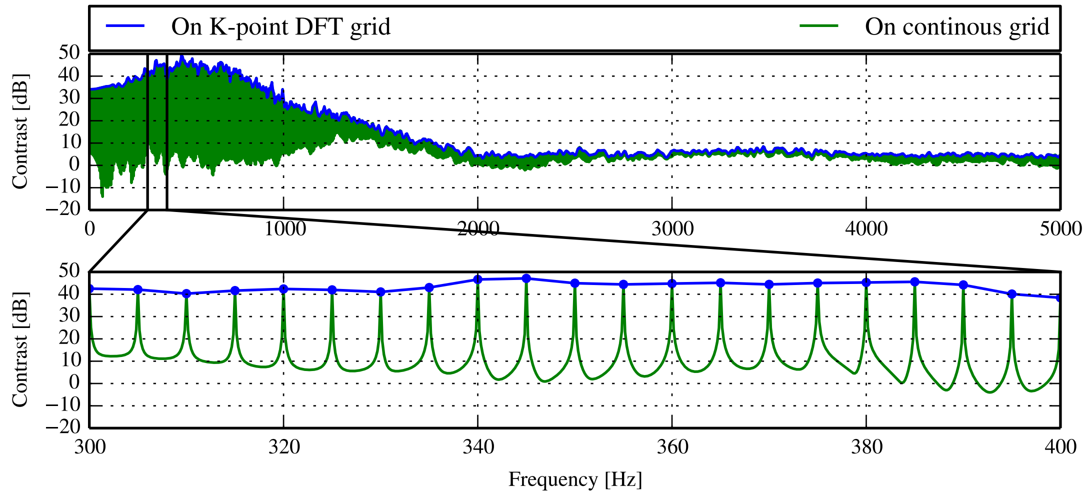

Figure 4.4: Upper: A plot of the obtained contrast for multi-frequency PM. The results have been obtained by performing simulations on the set-up introduced in Chapter 2. The blue points show the contrast obtained on the grid where the contrast control is performed (see (4.12)). The individual frequency points have been connected by straight lines to ease viewing. The contrast is also plotted for a more dense grid of frequencies (see (4.15)). The green line appears smeared because the contrast changes rapidly with frequency. Lower: A zoom on a small region of the upper plot. The figure is generated by the script `time_domain_aliasing.py`.

almost exactly the same as the method discussed in [19]. The extension to wideband signals cover exactly the same assumptions and leads to exactly the same problems as the method discussed in Section 4.3. We therefore cover the method slightly more briefly. We again consider the special case of $L = K$. For PM we can use exactly the same argument as in Section 4.3 to separate the wideband problem into K single frequency problems. The K separate problems can be written as:

$$\mathbf{q}_{k,\mathrm{PM}} = \underset{\mathbf{q}_k}{\arg\min}\; ||\mathbf{G}_k\mathbf{q}_k - \mathbf{d}_k||^2 + \delta||\mathbf{q}_k||^2, \tag{4.40}$$

for $k = 0, \ldots, K-1$, where:.

$$\mathbf{G}_k = \begin{bmatrix} \mathbf{B}_k \\ \mathbf{D}_k \end{bmatrix} \in \mathbb{C}^{2M\times N}. \tag{4.41}$$

This optimisation problem is solved in the same manner as in the single frequency case:

$$\mathbf{q}_{k,\mathrm{PM}} = (\mathbf{G}_k^{\mathrm{H}}\mathbf{G}_k + \delta\mathbf{I})^{-1}\mathbf{G}_k^{\mathrm{H}}\mathbf{d}_k. \tag{4.42}$$

This method requires a desired vector, $\mathbf{d}_k \in \mathbb{C}^{2M\times 1}$, for every frequency of the K point DFT-grid. These vectors describe the desired sound-field and there are no restrictions on how they can be chosen. A possibility is to make the desired sound-field in the bright zone a sound wave from some virtual source, while setting the desired sound-field in the dark zone equal to zero, as discussed in Section 3.2.1 on page 32. We can extend this to the wideband case in a straightforward manner:

$$\mathbf{d}_k[m] = \begin{cases} A_0 \frac{\exp(\frac{2\pi k}{c\, f_s}||\mathbf{x}_{\text{vs}}-\mathbf{x}_{\text{b},m}||)}{||\mathbf{x}_{\text{vs}}-\mathbf{x}_{\text{b},m}||}, & \text{if } m < M. \\ 0, & \text{if } m \geq M. \end{cases} \tag{4.43}$$

for $m = 0, \ldots, 2M-1$ and $k = 0, \ldots, K-1$, where $A_0 \in \mathbb{R}$ is the volume velocity of the virtual source, $\mathbf{x}_{\text{vs}} \in \mathbb{R}^{3\times 1}$ is the position of the virtual source, $\mathbf{x}_{\text{b},m} \in \mathbb{R}^{3\times 1}$ is the position of the m'th microphone and c is the speed of sound.

The presented method suffers from the same issue as demonstrated for MF-ACC. A demonstration of this is seen in Figure 4.4. We again hypothesise that the problem is caused by the underlying assumption of circular convolution. We refer to this approach as Multi Frequency PM (MF-PM). The entire procedure of applying MF-PM is shown by Algorithm 4.2.

Algorithm 4.2 MF-PM

input: $\mathbf{g}_{n,m} \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$ and $m = 0, \ldots, 2M-1$, $\mathbf{d}_k \in \mathbb{C}^{2M\times 1}$ for $k = 0, \ldots, K-1$
output: $\mathbf{h}_n \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$

\# Find filter responses
for $k = 0$ to $K-1$:
    Form $\mathbf{G}_k$ by (4.41)
    $\mathbf{q}_{k,\text{PM}} \leftarrow (\mathbf{G}_k^{\text{H}}\mathbf{G}_k + \delta\mathbf{I})^{-1}\mathbf{G}_k^{\text{H}}\mathbf{d}_k$

\# Obtain time domain filters by IDFT
for $n = 0$ to $N-1$:
    $\mathbf{h}_n \leftarrow$ IDFT of $[\mathbf{q}_{0,\text{PM}}[n], \mathbf{q}_{1,\text{PM}}[n], \ldots, \mathbf{q}_{K-1,\text{PM}}[n]]^{\text{T}}$

The K problems of (4.40) can be denoted as a single problem by utilising block matrix notation:

$$\tilde{\mathbf{q}}_{\text{PM}} = \arg\min_{\tilde{\mathbf{q}}} ||\widetilde{\mathbf{G}}\tilde{\mathbf{q}} - \tilde{\mathbf{d}}||^2 + \delta||\tilde{\mathbf{q}}||^2, \tag{4.44}$$

where:

$$\tilde{\mathbf{q}} = \begin{bmatrix} \mathbf{q}_0 \\ \vdots \\ \mathbf{q}_{K-1} \end{bmatrix} \in \mathbb{C}^{NK \times 1}, \tag{4.45}$$

$$\widetilde{\mathbf{G}} = \begin{bmatrix} \mathbf{G}_0 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_1 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{G}_{K-1} \end{bmatrix} \in \mathbb{C}^{2MK \times NK}, \tag{4.46}$$

$$\tilde{\mathbf{d}} = \begin{bmatrix} \mathbf{d}_0 \\ \vdots \\ \mathbf{d}_{K-1} \end{bmatrix} \in \mathbb{C}^{2MK \times 1}. \tag{4.47}$$

Representing the problem as a single optimisation problem is more convenient for the developments presented in Section 4.7. Note that this does not in anyway change the outcome of the method.

4.6 Linear and Circular Convolution

In sections 4.3 and 4.5 we present wideband sound zone methods which have highly frequency dependent acoustic contrast. This is essentially because the methods perform well at the frequencies for which the underlying optimisation problems are solved, but much more poorly in between these frequency points. We hypothesise that this highly frequency dependent behaviour stems from the fact that the optimisation problems implicitly assume that the convolutions between the input filters and the room filters are carried out as circular convolutions. Both in reality and in simulation they are in fact carried out as linear convolutions (as the room filters represent an actual physical system in continuous time, while circular convolution is essentially an artefact which occurs when performing discrete convolution by multiplying two DFT sequences). In this section we perform a short numerical investigation of the relationship between linear and circular convolution. This suggests a way to remedy the issues connected with circular convolution. A thorough theoretical review of the relationship between circular and linear convolution is provided in [24].

We investigate the effects of convolving two sequences with linear and circular convolution. We define the two sequences as:

$$\mathbf{x} = [\mathbf{x}[0], \dots, \mathbf{x}[K-1]]^{\mathrm{T}} \in \mathbb{R}^{K \times 1} \tag{4.48}$$

$$\mathbf{y} = [\mathbf{y}[0], \dots, \mathbf{y}[K-1]]^{\mathrm{T}} \in \mathbb{R}^{K \times 1}. \tag{4.49}$$

The linear convolution of the two sequences is written as:

$$\mathbf{z}_1 = \mathbf{x} * \mathbf{y} \ \in \mathbb{R}^{(2K-1)\times 1}. \tag{4.50}$$

We do not explicitly write a formula for the entries, as this requires rather complicated indexing and is deemed unnecessary. Importantly, the convolution product has length $2K-1$. We can convolve the two sequences in the frequency domain instead by 1) Taking the DFT of the sequences, 2) performing entry-wise multiplication of the two DFT sequences and 3) taking the IDFT of the resulting sequence [24]. We write this as:

$$\mathbf{z}_2 = \mathbf{x} \circledast \mathbf{y} = \boldsymbol{\mathcal{F}}^{-1}\left((\boldsymbol{\mathcal{F}}\mathbf{x}) \circ (\boldsymbol{\mathcal{F}}\mathbf{y})\right) \ \in \mathbb{R}^{K\times 1}, \tag{4.51}$$

where "$\circ$" is entry-wise multiplication (Hadamard product) and $\boldsymbol{\mathcal{F}}$ is the DFT. This results in a circular convolution product of length K (see [24] for a graphical illustration of the relationship between linear and circular convolution). Due to the fact that $\mathbf{z}_1$ and $\mathbf{z}_2$ belong to different spaces, they are obviously not the same. Interestingly, the differences between linear and circular convolution can be eliminated by adding appropriate zero-padding to the sequences $\mathbf{x}$ and $\mathbf{y}$ before performing convolution by (4.51) [24]. Enough zero-padding should be added for the resulting vector to be able to contain the linear convolution product. Thus, in this case we should zero-pad to length $2K-1$. We write the resulting convolution product as:

$$\mathbf{z}_3 = \mathbf{x}_{\{2K-1\}} \circledast \mathbf{y}_{\{2K-1\}} = \boldsymbol{\mathcal{F}}^{-1}\left((\boldsymbol{\mathcal{F}}_{\{2K-1,K\}}\mathbf{x}) \circ (\boldsymbol{\mathcal{F}}_{\{2K-1,K\}}\mathbf{y})\right) \ \in \mathbb{R}^{(2K-1)\times 1}, \tag{4.52}$$

where $\mathbf{x}_{\{2K-1\}}, \mathbf{y}_{\{2K-1\}} \in \mathbb{R}^{(2K-1)\times 1}$ are the same as $\mathbf{x}$ and $\mathbf{y}$ but with added zero padding. The matrix $\boldsymbol{\mathcal{F}}_{\{2K-1,K\}} \in \mathbb{C}^{(2K-1)\times K}$ represents the $2K-1$ point DFT of length K signals (carried out by padding the input with $K-1$ zeros, see Appendix B). When carrying out the circular convolution in this way, we obtain the same result as for the linear case, $\mathbf{z}_1 = \mathbf{z}_3$ [24][4].

A small numerical investigation of the two types of convolution has been carried out for the case of $K = 10$. The entries of the two vectors $\mathbf{x}$ and $\mathbf{y}$ are picked randomly from a normal distribution. The DTFTs of the resulting signals are shown in Figure 4.5. As expected, the DTFTs of $\mathbf{z}_1$ and $\mathbf{z}_3$ agree. This indicates that circular and linear convolution are equivalent when sufficient zero padding is added before carrying out

[4] The problem with circular convolution can be seen as aliasing in the time domain [24]. An exact dual of the problem exists in the time domain. If two continuous signals are appropriately sampled with respect to the sampling theorem, they may be exactly reconstructed. However, multiplying (modulating) the two sampled signals with one-another leads to a spread in frequency (by the modulation theorem [24]). This may cause the resulting signal to exceed the Nyquist frequency, introducing aliasing in the frequency domain.

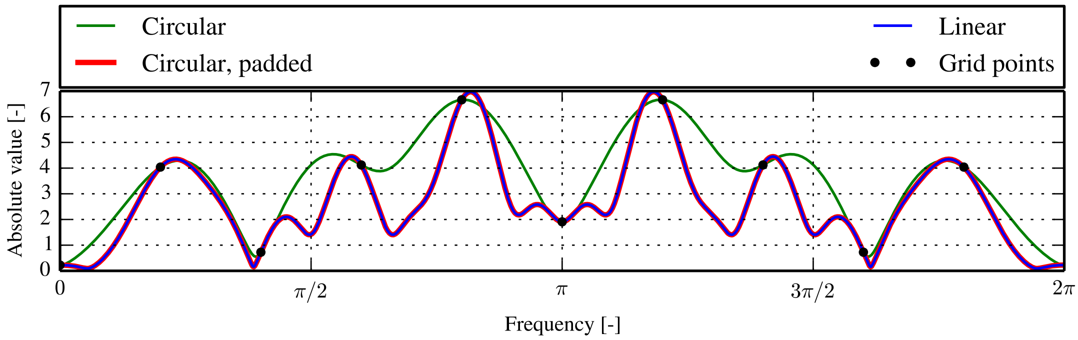

Figure 4.5: The DTFTs of the three discussed convolution products. The red and the blue lines overlap, and may therefore appear as one purple line. The black dots show the points of $(\mathcal{F}\mathbf{x}) \circ (\mathcal{F}\mathbf{y})$. It should be noted that the methods coincide at these exact points. The figure is generated by the script `convolution_demo.py`.

circular convolution. The DTFT of $\mathbf{z}_2$ is very dissimilar to the two others. It should be noted that all the three DTFTs coincide on the frequencies corresponding to a K-point DFT grid.

The example above illustrates that linear and circular convolution lead to different results unless appropriate zero padding is added before carrying out circular convolution. In sections 4.3 and 4.5 it is argued that the wideband sound zone methods MF-ACC and MF-PM assume circular convolution between the input filters and the room filters, while the convolution is actually carried out as a linear convolution when the methods are tested. Since linear and circular convolution lead to different results, this is not necessarily a reasonable assumption to make. The methods lead to high acoustic contrast on the frequencies associated with the K point DFT grid, but poor contrast at other frequencies. Figure 4.5 indicates that the DTFTs of linear and circular convolution products coincide on these particular frequencies. This loosely says that the assumption of circular convolution is non-problematic at these particular frequencies, but problematic at other frequencies. This provides a possible explanation of the highly frequency dependent behaviour of the results in MF-ACC and MF-PM. At the same time it suggests a remedy. If both the input filters and the room filters can be zero padded appropriately, circular convolution leads to the same results as linear convolution. In spite of significant effort being spent on the matter, no possible way of doing this has been found for MF-ACC. A modification of this exact type has, however, been devised for MF-PM. The method furthermore removes the restriction that the input filters must have the same length as the room filters ($L = K$). The method is described in the section below.

4.7 Wideband PM with Variable Filter Length

We present a modified version of the MF-PM method which remedies the problems associated with the underlying assumption of circular convolution and allows for the input filters to have an arbitrary length. The MF-PM method consists in solving an optimisation problem for the frequency domain coefficients of the input filters. The main idea of the derivations carried out in this section consist in transforming this optimisation problem into an equivalent problem over the time domain coefficients of the input filters. Doing so introduces the possibility of adding zero padding before carrying out circular convolution and allows for selecting the length of the input filters independently of the room filters.

We use the block-matrix formulation of the MF-PM optimisation problem, (4.44), as a starting point. Thus, the task is to express (4.44) as an optimisation problem over the time domain coefficients $\mathbf{h}_0, \ldots, \mathbf{h}_{N-1}$ rather than over the frequency domain coefficients $\mathbf{q}_0, \ldots, \mathbf{q}_{K-1}$. To avoid excessive summation over N we use block representations of these:

$$\tilde{\mathbf{q}} = \begin{bmatrix} \mathbf{q}_0 \\ \vdots \\ \mathbf{q}_{K-1} \end{bmatrix}. \in \mathbb{C}^{NK \times 1} \tag{4.53}$$

$$\tilde{\mathbf{h}} = \begin{bmatrix} \mathbf{h}_0 \\ \vdots \\ \mathbf{h}_{N-1} \end{bmatrix}. \in \mathbb{R}^{NL \times 1} \tag{4.54}$$

$$\tag{4.55}$$

The conversion into the time domain is carried out by use of the block-DFT operator, $\tilde{\mathbf{F}}$:

$$\tilde{\mathbf{F}} = \begin{bmatrix} \boldsymbol{\mathcal{F}} & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\mathcal{F}} & \ldots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \ldots & \boldsymbol{\mathcal{F}} \end{bmatrix} \in \mathbb{C}^{NK \times NK}. \tag{4.56}$$

First, we obtain the inverse DFT of $\tilde{\mathbf{q}}$ by use of $\tilde{\mathbf{F}}^{\mathrm{H}}$. The entries of $\tilde{\mathbf{q}}$ are grouped according to frequency and not according to the filter to which they belong ($\mathbf{q}$ is defined by (4.3)). It is therefore necessary to permute the entries of $\tilde{\mathbf{q}}$ before $\tilde{\mathbf{F}}^{\mathrm{H}}$ can be used to perform the IDFT (that is, $\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{q}} \neq \tilde{\mathbf{h}}$). We implement this permutation with the permutation matrix $\mathbf{P} \in \mathbb{R}^{NK \times NK}$. Since this is essentially a practicality, it is not deemed necessary to further elaborate on the particular structure of $\mathbf{P}$. As $\mathbf{P}$ is a permutation matrix it is unitary [44]:

$$\mathbf{P}^{\mathrm{H}}\mathbf{P} = \mathbf{P}\mathbf{P}^{\mathrm{H}}. \tag{4.57}$$

Permuting $\tilde{\mathbf{q}}$ with $\mathbf{P}$ produces a vector which is transformed into $\tilde{\mathbf{h}}$ by the IDFT, $\tilde{\mathbf{F}}^{\mathrm{H}}$:

$$\tilde{\mathbf{h}} = \tilde{\mathbf{F}}^{\mathrm{H}}\mathbf{P}\tilde{\mathbf{q}}. \tag{4.58}$$

where we have implicitly assumed $K = L$. Since both $\tilde{\mathbf{F}}$ and $\mathbf{P}$ are unitary, we have:

$$\mathbf{P}\tilde{\mathbf{F}}^{\mathrm{H}}\tilde{\mathbf{F}}\mathbf{P}^{\mathrm{H}} = \mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}\tilde{\mathbf{F}}^{\mathrm{H}}\mathbf{P} = \mathbf{I}. \tag{4.59}$$

We may insert this into (4.44):

$$\tilde{\mathbf{h}}_{\mathrm{PM}} = \arg\min_{\tilde{\mathbf{h}}} ||\widetilde{\mathbf{G}}\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}\tilde{\mathbf{F}}^{\mathrm{H}}\mathbf{P}\tilde{\mathbf{q}} - \tilde{\mathbf{d}}||^2 + \delta||\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}\tilde{\mathbf{F}}^{\mathrm{H}}\mathbf{P}\tilde{\mathbf{q}}||^2. \tag{4.60}$$

The above can be simplified to:

$$\tilde{\mathbf{h}}_{\mathrm{PM}} = \arg\min_{\tilde{\mathbf{h}}} ||\boldsymbol{\Psi}\tilde{\mathbf{h}} - \tilde{\mathbf{d}}||^2 + \delta||\tilde{\mathbf{h}}||^2, \tag{4.61}$$

where:

$$\boldsymbol{\Psi} = \widetilde{\mathbf{G}}\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}. \tag{4.62}$$

Thus, we have obtained the desired result of transforming (4.44) into the time domain. Do note that the transformed problem is exactly equivalent to (4.44) and therefore also includes the undesired assumption of circular convolution as well as the assumption of $K = L$. The advantage of (4.61) lies in the fact that it enables the possibility of adding zero-padding to the underlying circular convolution. To show this, we first sketch where the circular convolution is carried out. The matrix $\widetilde{\mathbf{G}}$ contains essentially the DFTs of the room filters. The input filters are given by $\tilde{\mathbf{h}}$ and the DFTs of these are obtained by application of $\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}$. Thus, the circular convolution takes place in the matrix multiplication $(\widetilde{\mathbf{G}})(\mathbf{P}^{\mathrm{H}}\tilde{\mathbf{F}}\tilde{\mathbf{h}})$. We may remedy the issues with circular convolution by adding appropriate zero padding as shown in Section 4.6. The matrix $\widetilde{\mathbf{G}}$ is computed directly from the room filters. It is therefore no problem to simply add zero padding to the room filters before computing $\widetilde{\mathbf{G}}$. We first recall that $\widetilde{\mathbf{G}}$ is a block diagonal matrix, given by (4.46). The diagonal blocks, $\mathbf{G}_0, \dots, \mathbf{G}_{K-1}$, are given by (this may be seen from (4.41) along with (4.7) and (4.8)):

$$\mathbf{G}_k[m,n] = \frac{1}{\sqrt{K}} \sum_{\ell=0}^{K-1} \mathbf{g}_{n,m}[\ell] \exp\left(-j2\pi\frac{k\ell}{K}\right), \tag{4.63}$$

with $k = 0, \dots, K-1$, $n = 0, \dots, N-1$ and $M = 0, \dots, 2M-1$. We generalise this definition to apply a ζ-point zero-padded DFT:

$$\widetilde{\mathbf{G}}_{\{\zeta\}} = \begin{bmatrix} \mathbf{G}_{0,\{\zeta\}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_{1,\{\zeta\}} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{G}_{\zeta-1,\{\zeta\}} \end{bmatrix} \in \mathbb{C}^{2M\zeta\times N\zeta}, \tag{4.64}$$

$$\tag{4.65}$$

where:

$$\mathbf{G}_{k,\{\zeta\}}[m,n] = \frac{1}{\sqrt{\zeta}} \sum_{l=0}^{K-1} \mathbf{g}_{n,m}[l] \exp\left(-j2\pi \frac{kl}{\zeta}\right), \tag{4.66}$$

where $\zeta \geq K$ is the dimension of the resulting DFT-grid (that is, a padding of $\zeta - K$ zeros are added to the room filters). This introduces the possibility of adding zero padding to the room filters. We can furthermore add zero padding to the input filters. Since $\mathbf{\Psi}$ directly includes a block DFT, $\tilde{\mathbf{F}}$, we simply have to generalise this to include zero-padding:

$$\tilde{\mathbf{F}}_{\{\zeta,L\}} = \begin{bmatrix} \boldsymbol{\mathcal{F}}_{\{\zeta,L\}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\mathcal{F}}_{\{\zeta,L\}} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \boldsymbol{\mathcal{F}}_{\{\zeta,L\}} \end{bmatrix} \in \mathbb{C}^{N\zeta \times NL}, \tag{4.67}$$

where $\boldsymbol{\mathcal{F}}_{\{\zeta,L\}} \in \mathbb{C}^{\zeta \times L}$ is the ζ point zero padded DFT for vectors of length L (discussed in Appendix B). We use this to define a zero-padded $\mathbf{\Psi}$:

$$\mathbf{\Psi}_{\{\zeta,L\}} = \tilde{\mathbf{G}}_{\{\zeta\}} \mathbf{P}^{\mathrm{H}}_{\{\zeta\}} \tilde{\mathbf{F}}_{\{\zeta,L\}} \in \mathbb{C}^{2M\zeta \times NL}. \tag{4.68}$$

where $\mathbf{P}_{\{\zeta\}} \in \mathbb{R}^{\zeta \times \zeta}$ is a necessary permutation of entries which changes the ordering of the input filter DFT (similar to $\mathbf{P}$ but for the ζ-point DFT). We may now replace $\mathbf{\Psi}$ with $\mathbf{\Psi}_{\{\zeta,L\}}$ in (4.61). We choose $\zeta = L + K - 1$ as this is the length of the convolution product of a room filter and an input filter. It is therefore a sufficient length to remove the effects of circular convolution as discussed in Section 4.6. The resulting optimisation problem is given by:

$$\tilde{\mathbf{h}}_{\mathrm{PM}} = \arg\min_{\tilde{\mathbf{h}}} ||\mathbf{\Psi}_{\{K+L-1,L\}} \tilde{\mathbf{h}} - \tilde{\mathbf{d}}_{\{K+L-1\}}||^2 + \delta ||\tilde{\mathbf{h}}||^2. \tag{4.69}$$

where $\tilde{\mathbf{d}}_{\{K+L-1\}} \in \mathbb{C}^{2M(K+L-1)\times 1}$ is defined in exactly the same manner as $\tilde{\mathbf{d}}$ but for the $(K + L - 1)$-point DFT-grid. It may be noted that the action of replacing $\mathbf{\Psi}$ with $\mathbf{\Psi}_{\{\zeta,L\}}$ removes the necessity of having $L = K$. The length of the input filters are thus not dependent on the length of the room filters. The optimisation problem is solved in the same manner as the single frequency PM:

$$\tilde{\mathbf{h}}_{\mathrm{PM}} = (\mathbf{\Psi}^{\mathrm{H}}_{\{K+L-1,L\}} \mathbf{\Psi}_{\{K+L-1,L\}} + \delta \mathbf{I})^{-1} \mathbf{\Psi}^{\mathrm{H}}_{\{K+L-1,L\}} \tilde{\mathbf{d}}_{\{K+L-1\}}. \tag{4.70}$$

It should, however, be noted that solving the problem by evaluating (4.70) may be computationally infeasible due to the large dimensions of the involved matrices. The matrix $(\mathbf{\Psi}^{\mathrm{H}}_{\{K+L-1,L\}} \mathbf{\Psi}_{\{K+L-1,L\}} + \delta \mathbf{I})$ belongs to $\mathbb{C}^{NL \times NL}$. We can simply choose L to be low enough for the inversion of this matrix not to be a problem. On the other hand this is an unfortunate and unnecessary limitation. To allow for a fair comparison between MF-ACC, MF-PM and CG-PM, it should be possible to use CG-PM for the case of $L = K$. For the simulation environment described in Chapter 2 we have $NK = 48 \cdot 2000 = 96000$. Inversion of a 96000×96000 matrix is not deemed feasible. This problem can be overcome by utilising alternative means for obtaining the solution. This is discussed in the section below.

4.8 Solution with the Conjugate Gradient Method

Solving regularised least squares problems such as that of, (4.69) is a very extensive topic which we shall not cover in detail here. The important detail to note in this particular case, is the fact that the operator $\mathbf{\Psi}_{\{K+L-1,L\}}$ is too large to store in memory. On the other hand, the action of applying it to a vector can be implemented rather efficiently by use Fast Fourier Transformation (FFT) and by taking advantage of the block structure. It is therefore important that the selected solution method does not rely on manipulating individual entries in $\mathbf{\Psi}_{\{K+L-1,L\}}$. We note that the objective function of the problem is a quadratic equation of the form:

$$\begin{aligned} f(\tilde{\mathbf{h}}) &= ||\mathbf{\Psi}_{\{K+L-1,L\}}\tilde{\mathbf{h}} - \tilde{\mathbf{d}}||^2 + \delta||\tilde{\mathbf{h}}||^2 \\ &= (\mathbf{\Psi}_{\{K+L-1,L\}}\tilde{\mathbf{h}} - \tilde{\mathbf{d}})^{\mathrm{H}}(\mathbf{\Psi}_{\{K+L-1,L\}}\tilde{\mathbf{h}} - \tilde{\mathbf{d}}) + \delta\tilde{\mathbf{h}}^{\mathrm{H}}\tilde{\mathbf{h}} \\ &= \tilde{\mathbf{h}}^{\mathrm{H}}(\mathbf{\Psi}^{\mathrm{H}}_{\{K+L-1,L\}}\mathbf{\Psi}_{\{K+L-1,L\}} + \delta\mathbf{I})\tilde{\mathbf{h}} - 2\mathrm{Re}\{\tilde{\mathbf{d}}^{\mathrm{H}}\mathbf{\Psi}_{\{K+L-1,L\}}\}\tilde{\mathbf{h}} + \tilde{\mathbf{d}}^{\mathrm{H}}\tilde{\mathbf{d}}. \end{aligned} \tag{4.71}$$

We choose to use an iterative algorithm to minimise this objective function. This allows for a trade-off between computational demands and solution accuracy. Furthermore, such algorithms typically do not require the ability to handle individual elements of the involved matrix (as opposed to e.g. algorithms which solve the problem by use of a matrix decomposition). There are many possible choices for such an algorithm (see e.g. [40] for a large selection of such). For this project, the Conjugate Gradient (CG) Algorithm has been chosen. This algorithm is only slightly more computationally complex than steepest descent and has very good convergence properties, particularly for quadratic problems [40]. The algorithm is first introduced in [45] and is summarised in [40]. Algorithm 4.3 shows the CG algorithm as it is applied to solve (4.69) (this is a rather simple adaptation compared to the version shown in [40]). Most of the involved arguments have been discussed, with the exception of the iteration count, $r_{\max}$. Several stopping criteria could have been chosen. Here, we simply run the algorithm for a fixed number of iterations, to better allow for investigating the convergence of the result. A common, and perhaps more practically sensible stopping criterion could be to stop the algorithm whenever the step-size, $||\alpha_r \mathbf{a}_r||$, reaches below some threshold [40]. We do not

Algorithm 4.3 Conjugate Gradient Method

1: **input:** $\boldsymbol{\Psi}_{\{K+L-1,L\}} \in \mathbb{C}^{M(K+L-1)\times NL}$, $\tilde{\mathbf{d}} \in \mathbb{C}^{M(K+L-1)\times 1}$, $r_{\max} \in \mathbb{N}$
2: **output:** $\tilde{\mathbf{h}}$

3: # Initialisation
4: $\tilde{\mathbf{h}}_0 \leftarrow \mathbf{0} \in \mathbb{R}^{NL\times 1}$
5: $\mathbf{b} \leftarrow 2\mathrm{Re}\{\tilde{\mathbf{d}}\boldsymbol{\Psi}_{\{K+L-1,L\}}\}^{\mathrm{H}}$
6: $\mathbf{a}_0 \leftarrow -\mathbf{b}$
7: $\mathbf{s}_0 \leftarrow \mathbf{b}$

8: # Iteration
9: **for** $r = 0$ to $r_{\max} - 1$:
10: $\quad \alpha_r \leftarrow -\frac{\mathbf{s}_r^{\mathrm{H}}\mathbf{a}_r}{2||\boldsymbol{\Psi}_{\{K+L-1,L\}}\mathbf{a}_r||^2 + 2\delta||\mathbf{a}_r||^2}$
11: $\quad \tilde{\mathbf{h}}_{r+1} \leftarrow \tilde{\mathbf{h}}_r + \alpha_r\mathbf{a}_r$
12: $\quad \mathbf{s}_{r+1} \leftarrow \mathbf{b} + 2(\boldsymbol{\Psi}^{\mathrm{H}}_{\{K+L-1,L\}}\boldsymbol{\Psi}_{\{K+L-1,L\}} + \delta\mathbf{I})\tilde{\mathbf{h}}_{r+1}$
13: $\quad \mathbf{a}_{r+1} \leftarrow -\mathbf{s}_{r+1} + \frac{||\mathbf{s}_{r+1}||^2}{||\mathbf{s}_r||^2}\mathbf{a}_r$
14: $\tilde{\mathbf{h}} \leftarrow \tilde{\mathbf{h}}_{r_{\max}}$

carry out a detailed analysis of the algorithm. A few notes on computational complexity and memory are, however, in order.

4.8.1 Computational Complexity

The complexity of the initialisation in Algorithm 4.3 is comparable to the complexity of one iteration and is therefore not particularly important (from a viewpoint of complexity). In the iteration step, the lines 11 and 13 have a computational complexity of $\mathcal{O}(NL)$. The lines 10 and 12 are dominated by the complexity of applying $\boldsymbol{\Psi}_{\{K+L-1,L\}}$. It is, however, not fair to assume that this has a complexity similar to that of a conventional matrix multiplication (which is $\mathcal{O}(NLM(K+L))$ in this case). The operator is given by (4.68). This provides a three step "recipe" for applying the operator:

1. Apply $\tilde{\mathbf{F}}_{\{K+L-1,L\}}$. This consists of zero padding, followed by N DFTs of length $K+L-1$. This can be done with a complexity of $\mathcal{O}(L(K+L)\ln(K+L))$ by using FFT [24].

2. Apply $\mathbf{P}_{\{K+L-1\}}$. This is a simple rearrangement of vector elements an does not contribute significantly to the computational demands.

3. Apply $\widetilde{\mathbf{G}}_{\{K+L-1\}}$. This is a block diagonal matrix with $K+L-1$ blocks of size $N \times 2M$. Multiplication by this matrix can therefore be implemented with a complexity of $\mathcal{O}((K+L)NM)$.

The three steps above have a combined complexity of $\mathcal{O}((K+L)(L\ln(K+L)+NM))$. Loosely speaking, the stepwise application of $\mathbf{\Psi}_{\{K+L-1,L\}}$ outperforms direct matrix multiplication whenever $(L\ln(K+L)+NM) << LNM$. For the investigated set-up we have $N=48$, $M=96$, $K=2000$. A reasonable selection of L may be somewhere in the order of 200. This gives rise to a difference of a factor of more than 100 between the compared complexities. In general it cannot be considered a realistic approach to treat $\mathbf{\Psi}_{\{K+L-1,L\}}$ directly as a matrix. We therefore use the stepwise approach. The overall complexity of each iteration of the CG algorithm therefore has a complexity of $\mathcal{O}((K+L)(L\ln(K+L)+NM))$.

4.8.2 Memory

Another significant issue is that of memory. The largest memory requirements of the algorithm are connected with the operator $\mathbf{\Psi}_{\{K+L-1,L\}} \in \mathbb{C}^{2M(K+L-1)\times NL}$. Assume that each complex entry is stored as two double precision floating points, thus taking up a total of 16 bytes. If we again assume $L=200$, the total memory requirement for the matrix is almost 65 GB. This is another reason as to why directly treating $\mathbf{\Psi}_{\{K+L-1,L\}}$ as a matrix is not appropriate. By applying the three-step method above, it is, however, not necessary to store $\mathbf{\Psi}_{\{K+L-1,L\}}$ in memory. Since steps 1 and 2 consists of nothing but taking FFTs and rearranging vector entries, we can assume that these take up no memory. The matrix $\widetilde{\mathbf{G}}_{\{K+L-1\}}$ of step 3 is very large, but we need only store the $(K+L-1)$ diagonal blocks, each of size $2M\times N$. Under the same assumptions as above, this amounts to approximately 325 MB. This can be considered an acceptable amount of memory to use for most personal computers and workstations.

4.8.3 CG-PM Overview

By solving the optimisation problem (4.69) with the CG Algorithm we obtain a computationally feasible sound zone method. In this method we have removed the effects of the underlying assumption of circular convolution from MF-PM. Furthermore, we have removed the requirement of having input filters of the same length as the room filters. We refer to this method as Conjugate Gradient PM (CG-PM). The procedure of applying the method is shown as pseudocode in Algorithm 4.4. While the matrix $\mathbf{\Psi}_{\{K+L-1,L\}}$ is required in the CG algorithm it should not be explicitly formed due to computational aspects. Matrix multiplication with $\mathbf{\Psi}_{\{K+L-1,L\}}$ can be obtained through the use of FFT as discussed in Section 4.8.1. The results of the method are compared to those of MF-ACC and MF-PM in Chapter 5.

Algorithm 4.4 CG-PM

input: $\mathbf{g}_{n,m} \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$ and $m = 0, \ldots, 2M-1$, $\widetilde{\mathbf{d}}_{\{K+L-1\}} \in \mathbb{C}^{2M(K+L-1)\times 1}$, $r_{\max} \in \mathbb{N}$
output: $\mathbf{h}_n \in \mathbb{C}^{K\times 1}$ for $n = 0, \ldots, N-1$

\# Find filter responses
Form $\widetilde{\mathbf{G}}_{K+L-1}$ by (4.64)
($\boldsymbol{\Psi}_{\{K+L-1,L\}}$ is not explicitly formed, see the text)
$\tilde{\mathbf{h}} \leftarrow$ solve (4.69) by Algorithm 4.3
Obtain $\mathbf{h}_0, \ldots, \mathbf{h}_{N-1}$ from $\tilde{\mathbf{h}}$ (see (4.54))

4.9 Summary

This chapter has discussed different methods for extending the single frequency methods, ACC and PM, to wideband signals. A straightforward extension to this approach, MF-ACC, is presented. It is shown that this method does not perform as well as one might initially expect. Another method, TD-ACC, is as well investigated. This method is shown to produce very undesirable results. A simple extension of PM to wideband signals, MF-PM, is also presented. This suffers from the same issues as MF-ACC. It is hypothesised that these problems stem from an underlying assumption of circular convolution embedded in the methods. It is shown that the MF-PM method can be converted from a frequency method to a time domain method. This leads to a method, CG-PM, which does not suffer from issues with circular convolution. The CG-PM approach involves a least-squares problem which is too large to solve by use of direct matrix inversion. The CG algorithm is presented as an alternative method to solve the optimisation problem.

Chapter 5

Results

In this chapter we present simulation results for the discussed sound zone methods. In Section 5.1 we shortly discuss the software framework used for carrying out the presented simulations, as well as the parameters chosen for the methods. The actual simulation results are presented and discussed in Section 5.2. A number of audio demonstrations for the discussed methods have as well been carried out. These are shortly discussed in Section 5.3.

5.1 Implementation Overview

Three of the four methods discussed in Chapter 4 have been implemented in Python to allow for a numerical investigation. These are the MF-ACC, MF-PM and CG-PM methods. The TD-ACC method has not been implemented as it is shown to be fundamentally flawed (Section 4.4). We do not go into details with the developed software framework, but a small overview of the structure is given.

The developed software framework consist mainly of two things: 1) A number of Python modules which implement widely used functionality and settings and 2) scripts which produce results (plots and audio files). The scripts which produce plots are listed in the captions of the figures throughout the report. All the involved files are listed in Appendix C. Here, we add a few comments to the modules which are used widely across the framework:

- `parameters.py`: Contains the overall parameters of the entire framework. Most of the scripts for generating results retrieve most of their parameters from this module. This arrangement makes it easy to change parameters across the entire framework at once. At the same time it ensures that all results are generated with the same parameters where-ever appropriate.

- `vogel.py`: Implements functionality for generating the Vogel pattern discussed in Chapter 2. Since the microphones in the simulation set-up are placed according to the Vogel pattern, this module is used across all scripts which carry out simulations on the set-up.

- `room.py`: Implements all necessary functionality with respect to room simulation. The module is used across all scripts which carry out room simulation of any type. The structure of this module is slightly more complicated than that of the others and is shortly discussed in Section 2.4 on page 23.

- `cgls.py`: Implements the CG Algorithm. This is used in all scripts which include the CG-PM method.

All matrix and vector operations are carried out with `numpy`[1] and `scipy`[2] in double precision arithmetic. A full reference of the used third party software is given in Appendix 2.4. A significant number of parameters must be set for the simulation framework. Most of these are due to the room simulation module. The chosen parameters for room simulation are discussed in Chapter 2 and are explicitly given in Table 2.1 on page 14. The remaining parameters belong to the three sound zone methods which we evaluate and to the CG Algortihm. These are shown in Table 5.1.

	Parameter	Value	Description
MF-ACC	e_0	*	Energy in the bright zone. See the text.
MF-PM/ CG-PM	δ	10^{-6}	Regularisation parameter.
	$\mathbf{x}_{\mathrm{vs}}$	$[4.0, 3.0, 1.6]^{\mathrm{T}}$	Virtual source position [m].
	A_0	0.1	Virtual source strength.
CG	$r_{\max}$	30	Number of iterations.

Table 5.1: An overview of the parameters used in the methods and the values selected for them. The choices underlying the parameters are discussed in the text.

The parameter e_0 determines the amount of acoustic energy in the bright zone for the MF-ACC method. It should be noted that changing the value corresponds to a simple scaling of the obtained input filters, and that it is therefore not very important to the results (the acoustic contrast is unaffected by e_0). The value of e_0 has been determined from the energy in the desired sound-field used in the MF-PM and CG-PM methods, in such a way that all three methods attempt to obtain the same amount of acoustic potential energy in the bright zone. This makes it more reasonable to compare the sound-fields resulting from the methods.

[1] `http://www.numpy.org/`.
[2] `http://www.scipy.org/`.

The parameter δ is a regularisation parameter for the MF-PM and CG-PM methods. For $\delta = 0$, the methods simply perform least-squares fitting of the sound-field to the desired sound-field. In this case the solution is not guaranteed to be unique. For $\delta > 0$, the methods similarly attempt to minimise the total energy in the resulting input filters. The influence of δ is investigated in [10]. Here, we have chosen δ slightly higher than zero to ensure uniqueness of the solution, while not including significant regularisation.

The position of the virtual source $\mathbf{x}_{\text{vs}}$ has been observed to be rather important to the obtained results for MF-PM and CG-PM. The source has been heuristically positioned to allow the methods to produce large acoustic contrast. The effect of changing its position is investigated as part of the presented results. The virtual source strength, A_0, has similar properties to e_0, but for MF-PM and CG-PM. It has no influence on the obtained acoustic contrast and corresponds to a pure scaling of the resulting input filters. The value of A_0 is therefore essentially an arbitrary choice.

The number of iterations to run for the CG algorithm is significant to the performance of CG-PM. Setting the value too low leads to suboptimal solutions. Setting the value too high leads to unnecessarily long computation time for the CG Algorithm. The selected value has been heuristically chosen as a low value for which the CG Algorithm has reached approximate convergence.

5.2 Simulation Results

In this section we present several plots of the performance of the three methods, MF-ACC, MF-PM and CG-PM. We especially focus on the performance of the CG-PM method as this is the main contribution of the thesis.

The main focus in the design phase of the methods has been the obtained acoustic contrast. We therefore first consider a plot of the obtained acoustic contrast as seen in Figure 5.1. The contrast is shown for the three discussed methods, and in the case of CG-PM for four different values of L (25, 100, 400 and 2000). The MF-ACC and MF-PM methods have a fixed value of L ($L = K = 2000$). The figure shows the contrast before and after a small change in the room simulation parameters (as discussed in Chapter 2). Changing the room simulation parameters is used as a simple means to illustrate how the methods react to a small change in room conditions (e.g. due to atmospheric changes or objects being moved around in the room). We consider first the results of the top plot. It is seen that the MF-ACC and MF-PM have highly frequency dependent contrast. The contrast varies so rapidly with frequency that the graphs appear smeared (see Figure 4.3 on page 44 or 4.4 on page 51 for a more detailed view of such graphs). In Chapter 4 it is hypothesised that this behaviour is caused by an underlying assumption of circular convolution. The CG-PM method has been developed in such a manner that this underlying assumption is removed. Is is seen from Figure 5.1 that the contrast of CG-PM does not

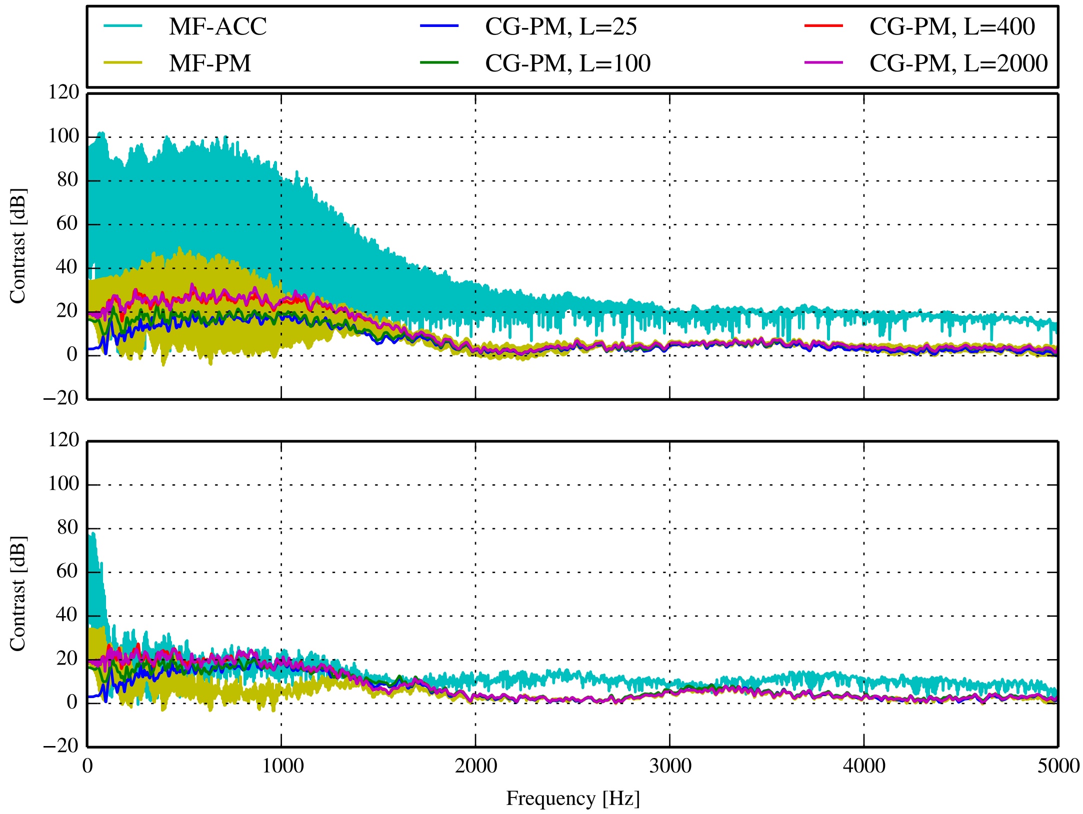

Figure 5.1: A comparison of acoustic contrast versus frequency for the discussed methods. The top plot shows the acoustic contrast obtained with the methods. The bottom plot shows the obtained contrast after slightly changing the parameters of the room simulation module as discussed in Chapter 2. The MF-ACC and MF-PM plots appear smeared because the contrast is highly dependent on frequency. The figure is generated by the script `methods_wideband.py`.

have this highly frequency dependent type of behaviour. This supports the hypothesis posed in Chapter 4. The contrast obtained for CG-PM is in between the upper and lower peaks of the MF-ACC and MF-PM methods. For all the methods, the obtained contrast decreases with increasing frequency. This is not surprising as the number of speakers and microphones necessary for controlling a sound-field increases with frequency. Similar results are observed in e.g. [10]. Another interesting facet of the plot is the dependence of L for the CG-PM solutions. Larger filter lengths provide more degrees of freedom and, not surprisingly, this leads to increased performance and thus increased contrast. Interestingly, there is almost no difference in contrast between the solution for $L = 400$ and that for $L = 2000$.

The lower plot of Figure 5.1 shows the contrast after having slightly changed the room simulation parameters. The input filters are designed to perform optimally, in some sense,

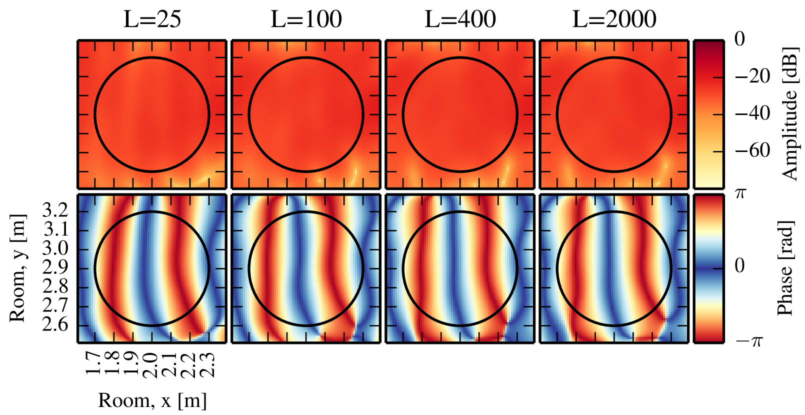

Figure 5.2: A plot of the sound-field for a sinusoidal input of 1000 Hz in the bright zone for CG-PM with different filter lengths. The axes are the same for all the plots. The upper row shows the amplitude of the sound-field for four different filter lengths while the lower row shows the corresponding phase. The figure is generated by the script `methods_wideband.py`.

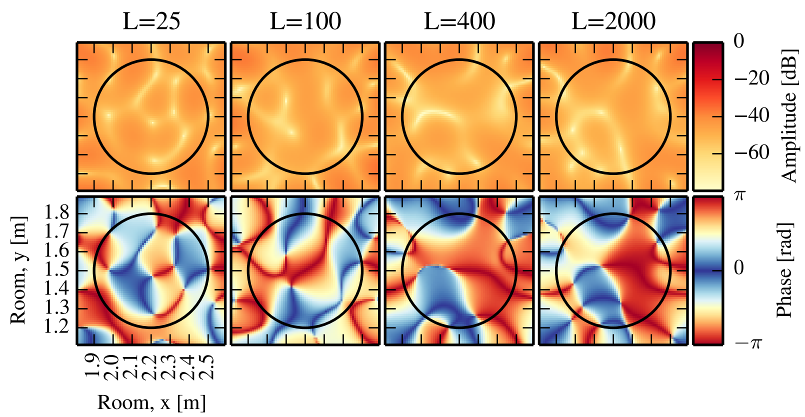

Figure 5.3: A plot of the sound-field for a sinusoidal input of 1000 Hz in the dark zone for CG-PM with different filter lengths. The axes are the same for all the plots. The upper row shows the amplitude of the sound-field for four different filter lengths while the lower row shows the corresponding phase. The figure is generated by the script `methods_wideband.py`.

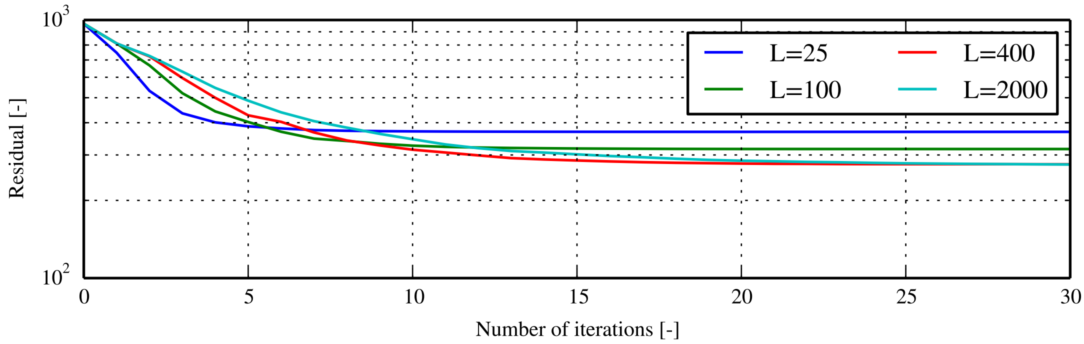

Figure 5.4: The objective function value versus CG Algorithm iteration count for CG-PM. The figure is generated by the script `methods_wideband.py`.

for the room represented by the room filters. When the room is changed after the input filters are designed, the input filters are no longer optimal in any sense. We should therefore expect to see a lower overall performance, including a lower acoustic contrast. This is clearly the case as seen in the figure. Especially the MF-ACC and MF-PM methods are severely affected. The narrow spikes of high contrast associated with these methods are very sensitive to small perturbations in the acoustic properties of the environment. The CG-PM solutions for $L = 400$ and for $L = 2000$ show a significant decrease in contrast as well. The CG-PM solutions for $L = 25$ and for $L = 100$ appear to be almost unaffected by the change.

The CG-PM method attempts to recreate a desired sound-field in the bright zone while keeping the dark zone quiet. The desired sound-field used in this project is the field produced by an ideal monopole under free-field conditions. The sound-field in the bright zone for a sinusoidal input of 1000 Hz is shown in Figure 5.2 for different filter lengths. The amplitude of the sound-field appears even across the zone for all the filters. The phase appears as a wave emanating from a source to the right of the zone, in accordance with the position of the virtual source. The corresponding sound-fields in the dark zone are seen in Figure 5.3. The sound-field in the dark zone appears unstructured in comparison to the field in the bright zone. It is difficult to see any dependence on the filter length in the plots. Close inspection of Figure 5.2 does, however, reveal slight variations in the amplitude of the sound-field throughout the bright zone for the shortest filter length.

An important part of the CG-PM method is the CG Algorithm. The CG Algortihm is an iterative algorithm which gradually approaches the solution to a least squares problem. The algorithm is guaranteed to converge in NL iterations [40] (not accounting for the effects of finite precision arithmetic). In practice it has, however, been found that the algorithm reaches approximate convergence in significantly fewer iterations. The objective function value versus the iteration count is plotted in Figure 5.4. The

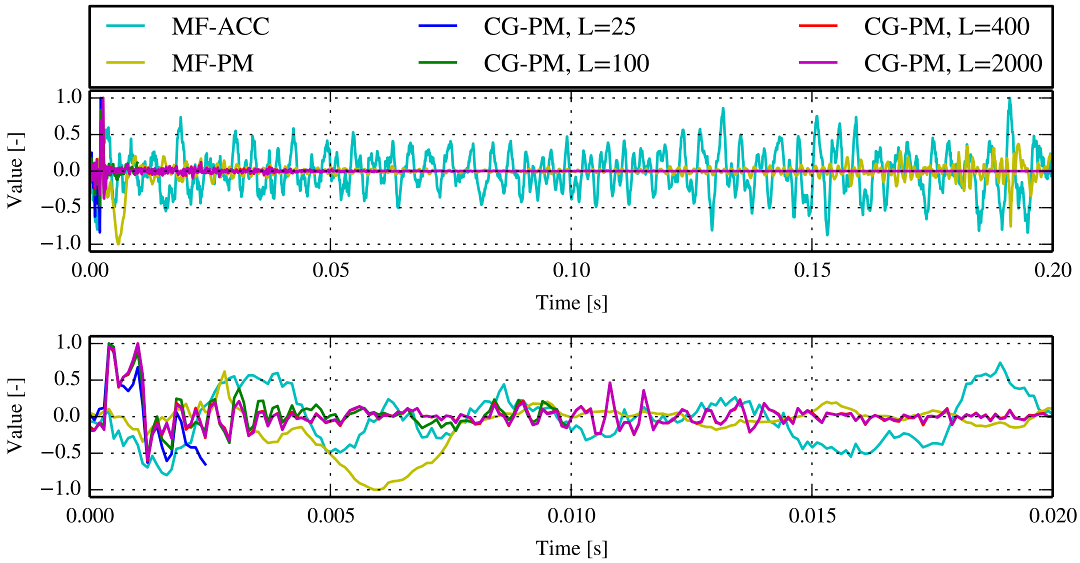

Figure 5.5: One of the input filters for each of the compared methods. The filters have all been normalised to the interval $[-1, 1]$ to ease comparison. The filters are plotted as lines connecting consecutive samples. The two plots differ only by the scale on the time axis. Note that the lower plot covers only 200 samples. The figure is generated by the script `methods_wideband.py`.

objective function of CG-PM is given by (4.69) on page 58. It is seen that approximate convergence is reached after 30 iterations. The figure also suggests that shorter filters converge in fewer iterations. Furthermore, not surprisingly, longer filters converge to a smaller objective function value.

Interesting insights can as well be gained from considering the obtained filters in the time domain. In Figure 5.5, one arbitrary input filter has been plotted for each of the compared methods (that is, the filter corresponding to $n = 7$ has been plotted for the different methods). Since only one filter from each method is compared we can of course not draw entirely general conclusions from the plot. The energy of the CG-PM filters is mainly focused in the first few samples where-after the filter coefficients converge to zero. This is well in line with the observation that there is little difference between the performance of the filter with $L = 400$ and the one with $L = 2000$. If the last 1600 coefficients of the long filter are very close to zero, the performance of the two filters should not differ much. The MF-ACC and MF-PM methods distribute the energy more evenly across the coefficients. For MF-ACC the energy appears to be spread approximately uniformly across the coefficients, while the MF-PM focuses the energy near the end of the filters.

Figure 5.6 shows the obtained acoustic contrast at 1000 Hz versus filter length for the CG-PM method. The plot shows that the filter length has little influence on the

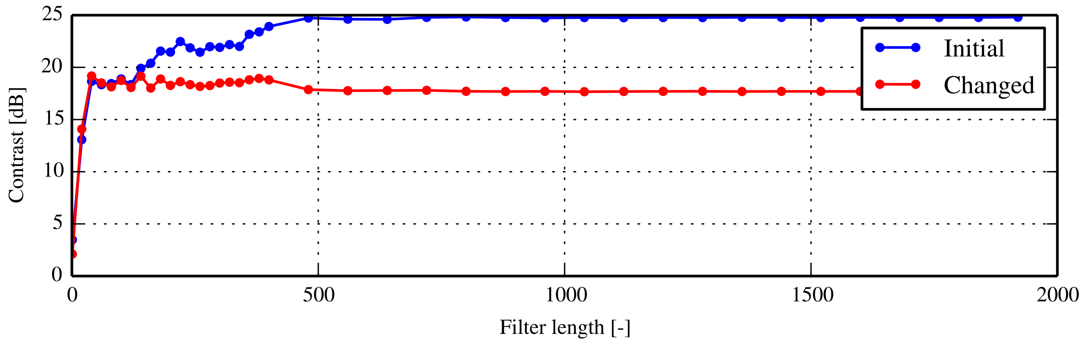

Figure 5.6: The acoustic contrast for CG-PM at 1000 Hz versus the filters length (L). The contrast is plotted both before and after introducing a slight change in the parameters of the room simulation set-up as previously discussed. The points for which the contrast has been computed are marked and are connected by straight lines for ease of viewing. The figure is generated by the script `filter_length.py`.

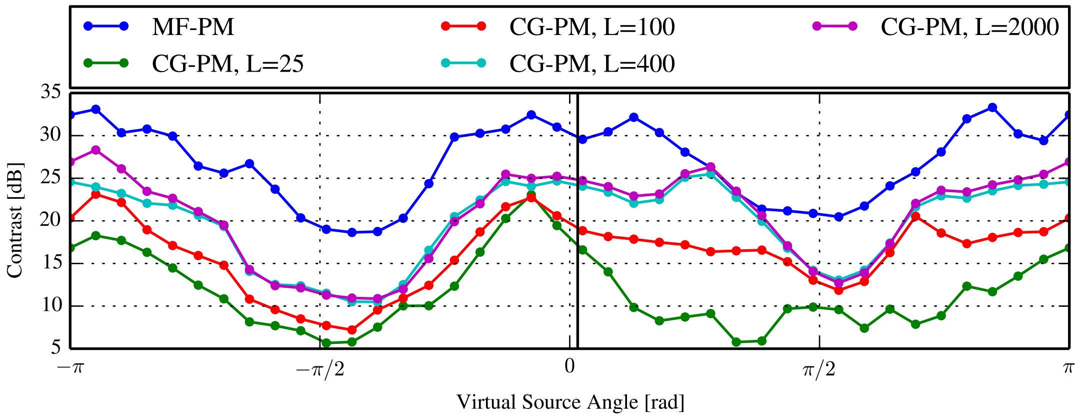

Figure 5.7: Acoustic contrast at 1000 Hz versus the angle of the virtual source in relation to the centre of the bright zone. It should be noted that 1000 Hz corresponds to a peak in acoustic contrast for MF-PM. The points for which the contrast has been computed are marked and are connected by straight lines for ease of viewing. The vertical line indicates the angle at which the virtual source is located for the other presented results. The figure is generated by the script `vsource_position.py`.

obtained contrast for filter lengths above 500, both with and without a change in the room simulation parameters. This is well in line with the other observations. For filter lengths in the range 100-400, the contrast with changed simulation parameters is slightly increased in comparison to the contrast for longer filters. This suggests that shorter filters lead to slightly increased robustness towards changes in the room conditions.

The MF-PM and CG-PM methods attempt to recreate some desired sound-field. Throughout the thesis we have chosen this sound-field to be a spherical field emanating from an ideal monopole source under free-field conditions. The position of this virtual source has a very significant impact on the obtained acoustic contrast (also investigated in e.g. [10]). To showcase this impact here, the contrast has been simulated for different virtual source positions. The virtual source has been rotated around the centre of the bright sound zone in the horizontal plane, while maintaining the same distance to it as is the case for the other simulation results. The contrast versus angle is seen in Figure 5.7. The obtained contrast is clearly highly dependent on the location of the virtual source. The methods show two points of very low contrast. One corresponds to the case where the virtual source is located such that the dark zone is in between the bright zone and the virtual source. The other corresponds to the case where the virtual source is located such that the bright zone is in between the dark zone and the virtual source[3].

5.3 Audio Demonstrations

A number of audio demonstrations have been made with the compared methods. The room simulation framework has been used to compute the response from each source to a microphone at the centre of each of the two zones (for the set-up discussed in Chapter 2). This, together with the generated input filters, allows for computing the combined impulse response from the system input and to a microphone at the centre of each of the zones. By convolving an audio signal with the obtained impulse response, we can obtain the audio signal as recorded in the middle of each of the zones. This provides a possibility of comparing the subjective qualities of the methods. Audio demonstrations have been produced out for MF-ACC, MF-PM and CG-PM (L = 25, 100, 400, 2000), both with and without changed room conditions, and for five different audio samples. The audio demonstrations are found on the CD (see Appendix C for more details). Since this thesis mainly focuses on the signal processing of sound zones, we do not carry out a thorough investigation of the subjective qualities of the methods. We do, however, provide a few observations:

[3]The CD contains an animation of the simulated sound-field of MF-PM at 1000 Hz while rotating the virtual source in the described manner. The animation is simulated under free-field conditions, in contrary to the plots shown in the thesis. The animation is found in `<CD>/Software/Output/gif_vsource_position/animation_pm_rotate_virtual_source.gif`.

- The audio heard in the bright zone for the MF-ACC and MF-PM methods sounds highly distorted. The sound heard in the dark zone is not as loud as that in the bright zone, but is clearly audible.
- The audio heard in the bright zone for the CG-PM does not appear distorted to the same extend as is the case for the two other methods. For all the four compared filter lengths, the audio in the dark zone is audible but not nearly as loud as in the bright zone. The audio sounds very similar for filter lengths 400 and 2000. The audio heard in the dark sound is, however, clearly less loud than for the two shorter filter lengths.
- It is generally clearly audible that all the methods produce larger acoustic contrast for low frequencies. The audio heard in the dark zone sounds as if it has been high-pass filtered.

The above points are based on the authors own subjective judgements and we suggest that the reader draws his or her own conclusions based on the audio demonstrations supplied on the CD.

5.4 Summary

In this chapter we have discussed a number of simulation results for the wideband sound zone methods MF-ACC, MF-PM and CG-PM. The MF-ACC, MF-PM methods are seen to suffer from highly frequency dependent behaviour. The CG-PM does not suffer from this type of behaviour. This is well in line with the fact that the CG-PM method has been developed specifically to avoid it. In addition to this result, several results on the general behaviour of CG-PM in relation MF-ACC and MF-PM have been discussed.

Chapter 6

Conclusions and Future Perspectives

The present master's thesis has investigated the possibilities for creating sound zones for wideband signals. The topic of sound zones has been extensively investigated in previous literature, but mainly for single tone signals.

The work presented in the thesis has used two single tone methods, Acoustic Contrast Control (ACC) and Pressure Matching (PM), as a staring point. These methods have been developed to work for a grid of discrete frequencies rather than just one frequency. We have referred to these methods as Multi Frequency ACC (MF-ACC) and Multi Frequency PM (MF-PM). The steps taken to allow for this have previously been investigated in the literature to some extent. The performance of the resulting methods have been investigated by use of simulations. The primary measure used for carrying out this performance evaluation is the acoustic contrast. The methods are shown to lead to highly frequency dependent acoustic contrast. It is verified through audio demonstrations that this behaviour is highly problematic to the subjectively experienced audio quality. It is hypothesised that the problem is caused by an underlying assumption of circular convolution in the design process. We have therefore proposed a modification of the MF-PM method which eliminates this underlying assumption of circular convolution. This method is referred to as Conjugate Gradient PM (CG-PM) because it involves the Conjugate Gradient Algortihm. The CG-PM method has as well been investigated by use of simulations. The acoustic contrast obtained through this method is not frequency dependent to the same extent as the two other methods. It is verified through audio demonstrations that this leads to significant reductions in audible artefacts as compared to the MF-ACC and MF-PM methods (based on the authors subjective judgements).

The discussed methods for obtaining sound zones have been investigated by use of a simulation framework which has been developed for the project. The framework carries out simulations of room acoustics based on the mirror image model. This has provided a firm foundation for the investigation of the developed methods by allowing controlled, robust and repeatable performance evaluations. It is, however, also important to be aware of the limitations of the framework. The underlying model assumes simple idealised representations of microphones, loudspeakers, reflective surfaces and the air though which the sound propagates. It furthermore assumes complete absence of measurement noise and other disturbances. It is therefore likely that several additional insights can be gained by comparing the discussed methods by use of actual measurements.

In the present thesis, we have investigated wideband sound zones mostly from a signal processing perspective. It is worth stressing that sound zone systems are intended for human listening. Therefore, it is likely that the obtained listening experience can be significantly improved by including elements from psychoacoustics into the discussed methods. Another matter which we have only briefly discussed is the computational aspects of the discussed methods. Especially the CG-PM method involves a significant computational load. An investigation of such matters may reveal possibilities for significant improvements with respect to execution time and memory requirements.

An important topic of future research is implicitly posed by the methods discussed in the thesis. We have discussed two simple wideband methods, MF-ACC and MF-PM. Both are shown shown to have undesirable features. We have shown a modification of MF-PM which does not suffer from these issues. It has, however, not been possible to devise a similar modification of MF-ACC. This should by no means be taken as an indication that such a thing is impossible. If found, such a method may have very advantageous properties.

The overall objective of this thesis has been to investigate the hypothesis that single frequency sound zone methods can be adapted to wideband signals. We have shown that this is indeed the case by investigating a number of approaches for doing so. In the process, we have revealed that the wideband case of sound zones includes multiple pitfalls which are not seen in the single frequency case. We have, however, also shown that these pitfalls can be avoided through appropriate steps. While sound zones for wideband signals is definitely a possibility, it is also clear that much work remains before sound zones can be considered for consumer products such as the hypothetical system discussed in the introduction.

Bibliography

[1] G. Cassidy and R. Macdonald, "The Effect of Background Music and Background Noise on the Task Performance of Introverts and Extroverts," *Psychology of Music*, vol. 35, pp. 517–537, 2007.

[2] E. Torick, "Highlights in the History of Multichannel Audio," *The Journal of the Audio Engineering Society*, vol. 46, pp. 27–31, Jan. 1998.

[3] D. J. Steffen, *From Edison to Marconi - The First Thirty Years of Recorded Music*. McFarland & Company, 2005.

[4] A. D. Blumlein, "Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing Systems," U.K. Patent, Jun. 14, 1933.

[5] J. Abildgaard Pedersen and K. Thomsen, "Fully Automatic Loudspeaker-Room Adaptation - The RoomPerfect System," in *Audio Engineering Society Conference: 32nd International Conference: DSP For Loudspeakers*, Sep. 2007.

[6] H. Møller, "Fundamentals of Binaural Technology," *Applied Acoustics*, vol. 36, pp. 171–218, 1992.

[7] W. F. Druyvesteyn and J. Garas, "Personal Sound," *Journal of the Audio Engineering Society*, vol. 45, no. 9, pp. 685–701, 1997.

[8] J. W. Choi and Y. H. Kim, "Generation of an Acoustically Bright Zone with an Illuminated Region Using Multiple Sources," *The Journal of the Acoustical Society of America*, vol. 111, no. 4, pp. 1695–1700, 2002.

[9] O. Kirkeby, P. A. Nelson, F. Orduna-Bustamante, and H. Hamada, "Local Sound Field Reproduction Using Digital Signal Processing," *The Journal of the Acoustical Society of America*, vol. 100, no. 3, Sep. 1996.

[10] M. Poletti, "An Investigation of 2D Multizone Surround Sound Systems," in *125th Audio Engineering Society Convention*, vol. 1, pp. 167–175, 2008.

[11] M. Shin, S. Q. Lee, F. M. Fazi, P. A. Nelson, D. Kim, S. Wang, K. H. Park, and J. Seo, "Maximization of Acoustic Energy Difference Between Two Spaces," *The Journal of the Acoustical Society of America*, vol. 128, no. 1, pp. 685–701, 2010.

[12] J. Chang, C. Lee, J. Park, and Y. Kim, "A Realization of Sound Focused Personal Audio System Using Acoustic Contrast Control," *The Journal of the Acoustical Society of America*, vol. 125, no. 4, pp. 2091–2097, 2009.

[13] M. B. Møller, M. Olsen, and F. Jacobsen, "A Hybrid Method Combining Synthesis of a Sound Field and Control of Acoustic Contrast," in *132nd Audio Engineering Society Convention*, 2012.

[14] P. Coleman, P. Jackson, M. Olik, and J. A. Pedersen, "Optimizing the Planarity of Sound Zones," in *Audio Engineering Society: 52nd International Conference: Sound Field Control - Engineering and Perception*, Sep. 2013.

[15] J. Ahrens and S. Spors, "Reproduction of a plane-wave sound field using planar and linear arrays of loudspeakers," in *3rd International Symposium on Communications, Control and Signal Processing*, pp. 1486–1491, 2008.

[16] C. Chung and S. J. Elliott, "Personal Audio Headrest," in *120th Audio Engineering Society Convention*, May 2006.

[17] M. Jones and S. Elliott, "An Active Headrest for Personal Audio," in *International Workshop on Acoustic Echo and Noise Control*, pp. 145–148, 2005.

[18] F. Jacobsen, M. Olsen, M. Møller, and F. Agerkvist, "A Comparison of Two Strategies for Generating Sound Zones in a Room," in *18th International Congress on Sound and Vibration*, 2011.

[19] N. Radmanesh and I. Burnett, "Generation of Isolated Wideband Sound Fields Using a Combined Two-stage Lasso-LS Algorithm," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 21, no. 2, pp. 378–387, Feb 2013.

[20] J. Chang, J. Park, and Y. Kim, "Scattering Effect on the Sound Focused Personal Audio System," *The Journal of the Acoustical Society of America*, vol. 125, no. 5, pp. 3060–3066, 2009.

[21] S. J. Elliott, J. .Cheer, H. Murfet, and K. R. Holland, "Minimally Radiating Sources for Personal Audio," *The Journal of the Acoustical Society of America*, vol. 128, no. 4, pp. 1721–1728, 2010.

[22] M. Jones and S. J. Elliott, "Personal Audio with Multiple Dark Zones," *The Journal of the Acoustical Society of America*, vol. 124, no. 6, pp. 3497–3506, 2008.

[23] J. Cheer, S. J. Elliott, and M. F. S. Gálvez, "Design and Implementation of a Car Cabin Personal Audio System," *Journal of the Audio Engineering Society*, vol. 61, no. 6, pp. 412–424, 2013.

[24] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Pearson, 2010.

[25] P. Coleman, P. J. Jackson, M. Olik, M. Olsen, M. Møller, and J. A. Pedersen, "The Influence of Regularization on Anechoic Performance and Robustness of Sound Zone Methods," *Proceedings of Meetings on Acoustics*, vol. 19, no. 1, pp. –, Jun. 2013.

[26] Y. Wu and T. Abhayapala, "Spatial Multizone Soundfield Reproduction," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2009. ICASSP 2009.*, pp. 93–96, Apr. 2009.

[27] H. Vogel, "A Better Way to Construct the Sunflower Head," *Mathematical Biosciences*, vol. 44, no. 3–4, pp. 179–189, 1979.

[28] D. E. Dudgeon and R. M. Mersereau, *Multidimensional Signal Processing*, ser. Prentice-Hall Signal Processing Series. Prentice-Hall, 1984.

[29] J. H. Rindel, "The Use of Computer Modeling in Room Acoustics," *Journal of Vibroengineering*, vol. 3, no. 4, pp. 219–224, 2000.

[30] F. Mechel, "Improved Mirror Source Method in Room Acoustics," *Journal of Sound and Vibration*, vol. 256, no. 5, pp. 873 – 940, 2002.

[31] J. B. Allen and D. A. Berkley, "Image Method for Efficiently Simulating Small-Room Acoustics," *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, Apr. 1979.

[32] E. A. Lehmann and A. M. Johansson, "Prediction of Energy Decay in Room Impulse Responses Simulated with an Image-Source Model," *The Journal of the Acoustical Society of America*, vol. 124, no. 1, pp. 269–277, Jul. 2008.

[33] P. M. Peterson, "Simulating the Response of Multiple Microphones to a Single Acoustic Source in a Reverberant Room," *The Journal of the Acoustical Society of America*, vol. 80, no. 5, pp. 1527–1529, Nov. 1986.

[34] F. Fahy, *Foundations of Engineering Acoustics.* Academic Press, 2001.

[35] H. Ebert, Ed., *Elektronik Ståbi, 8. Udgave.* Ingeniøren Bøger, 2003.

[36] J. Park, J. Choi, and Y. Kim, "Acoustic contrast sensitivity to transfer function errors in the design of a personal audio system," *The Journal of the Acoustical Society of America*, vol. 134, no. 1, pp. EL112–EL118, 2013.

[37] R. A. Serway and J. W. Jewett, *Physics for Scientists and Engineers.* Thomson Brooks/Cole, 2004.

[38] J. Martinez and R. Heusdens, "On Low-Complexity Simulation of Multichannel Room Impulse Responses," *Signal Processing Letters, IEEE*, vol. 17, no. 7, pp. 667–670, July 2010.

[39] P. Vandewalle, J. Kovačević, and M. Vetterli, "Reproducible Research in Signal Processing," *IEEE Signal Processing Magazine*, vol. 26, no. 3, pp. 37–47, May 2009.

[40] A. Antoniou and W. Lu, *Practical Optimization - Algorithms and Engineering Applications.* Springer, 2007.

[41] I. Kra and S. R. Simanca, "On Circulant Matrices," *Notices of the American Mathematical Society*, vol. 59, no. 3, pp. 368–377, Mar. 2012.

[42] R. M. Gray, "Toeplitz and Circulant Matrices: A review," Jan. 2006.

[43] L. Hogben, *Handbook of Linear Algebra*, 1st ed., ser. Discrete Mathematics and its Applications. Taylor & Francis, 2006.

[44] R. A. Horn and C. R. Johnson, *Matrix Analysis.* Cambridge University Press, 2012.

[45] M. R. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving linear systems," *Journal of research of the National Bureau of Standards*, vol. 49, no. 6, pp. 409–436, 1952.

[46] L. E. Kinsler, A. R. Frey, A. B. Coppens, and J. V. Sanders, *Fundamentals of Acoustics*, 4th ed. Wiley, 2000.

[47] J. Cooley and J. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.

Appendix A

Acoustics

While the present thesis is primarily focused on the signal processing aspects of sound zone technology, it does rely on acoustics to some extent. This appendix shortly introduces the mathematical model used for studying acoustics as well as relevant concepts, quantities and notation. This appendix is not in any way meant as an introduction to acoustics as a whole. For a broad introduction to the topic of acoustics, we refer to textbooks on the subject (e.g. [34, 46]). Section A.1 introduces the mathematical model most commonly used for studying acoustics and discusses the underlying assumptions. Section A.2 introduces the monopole source; an idealised loudspeaker which is heavily employed in the project. Most of the presented material is loosely based on [34].

A.1 The Model Underlying Acoustics

Acoustics is the study of mechanical waves propagating in any medium. Commonly this medium is air. This is also the case of interest in this thesis. Air is a very complicated non-linear medium. The properties of air with respect to waves depend on parameters such as pressure, velocity and temperature which vary as waves are transmitted through the medium and may as well vary across space. Mathematical modelling of acoustics under very general assumptions may therefore be an intractable task. A number of assumptions are therefore most often introduced [34]:

- Air is a homogeneous medium: this indicates that air has the same resting properties (pressure, temperature, speed of sound) everywhere in the considered space,
- Air satisfies the ideal gas law: this provides a linear relationship between pressure and temperature,
- Air is isotropic: air has the same properties in all directions,

- Air is inviscid: it has no "thickness" and flows without resistance.
- Fluctuations are small enough for air to act as a linear elastic medium: changes in pressure and temperature due to waves are too small to cause significant changes in the properties of air.

These assumptions are generally accepted and hold to within reasonable accuracy for acoustics in small rooms as considered in this thesis [34]. Under these assumptions, sound waves (in the absence of boundaries or sources) obey the wave equation [34]:

$$\nabla^2 p = \frac{1}{c^2}\frac{\partial^2 p}{\partial t^2}, \tag{A.1}$$

where

$p = p(\mathbf{x}, t)$ is the sound pressure deviation from resting pressure [Pa],
$\mathbf{x} \in \mathbb{R}^{3\times 1}$ is spatial displacement [m],
t is time [s],
c is the speed of sound [m/s].

The speed of sound is given by [34]:

$$c = \sqrt{\frac{\gamma P_0}{\rho_0}} \tag{A.2}$$

where

ρ_0 is the density of the air [kg / m^3],

P_0 is the air pressure at rest [Pa],

γ is a material constant ($\gamma_{\text{air}} \approx 1.4$).

One important solution to the wave equation, (A.1), is the plane wave [34]:

$$p(\mathbf{x}, t) = f(ct - \mathbf{x} \cdot \mathbf{n}) \tag{A.3}$$

where:

f is any function of one real variable,

$\mathbf{n} \in \mathbb{R}^{3\times 1}$ is a unit vector pointing in the direction of propagation.

Another one is the spherical wave, expressed in spherical coordinates [34]:

$$p(r, t) = \frac{1}{r}[f(ct - r) + g(ct + r)] \tag{A.4}$$

where:

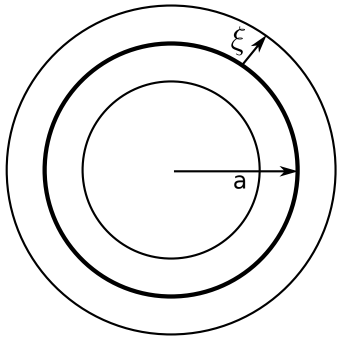

Figure A.1: A sphere of radius a, vibrating with amplitude ξ. The figure is an almost exact replica of Fig. 6.6 in [34].

r is the distance from the origin of the coordinate system.

f is any function of one real variable, representing a wave travelling outwards from the centre of the coordinate system,

g is any function of one real variable, representing a wave travelling inwards to the centre of the coordinate system.

Since the model treats air as a linear medium, any linear combination of the above solutions is as well a solution.

A.2 The Monopole Source

The wave equation, (A.1), describes the behaviour of acoustical waves under the assumption that no external forces are present. The wave equation does therefore not explain where acoustical waves originate from. It is possible to extend the wave equation to take external forces into account. We do not go into any details with the specifics of doing so here (this is covered in [34]). This makes it possible to take into account the effects of acoustical waves emitted from vibrating arrangements of solid matter. A simple example of such a vibrating arrangement of material is the pulsating sphere as shown in Figure A.1. The radius of the sphere oscillates with a displacement given by:

$$\xi = \xi_0 \exp(j2\pi f t), \tag{A.5}$$

where

ξ_0 is the amplitude of the oscilation [m],

f is the frequency of the oscilation [Hz].

It is assumed that the amplitude of the oscilation, ξ_0, is much smaller than the radius of the sphere, a. When the sphere oscilates, it emits acoustical waves which are spherically symmetric around the center of the sphere [34]. The volume velocity of the source is defined as the time rate of change of the amount of air that the source displaces. The volume velocity oscillates synchronously with the sphere with amplitude, $\tilde{Q}_0$ [34]:

$$Q = \tilde{Q}_0 \exp{(j\omega t)}, \tag{A.6}$$

where the amplitude Q_0 is given by [34]:

$$\tilde{Q}_0 = j8\pi^2 f \tilde{\xi}_0 a^2. \tag{A.7}$$

The pressure outside of the sphere is given by [34]:

$$p(r) = -\left[\frac{j2\pi f \rho_0 \tilde{Q}_0}{4\pi r}\right]\left[\frac{1}{1 + j\frac{2\pi f}{c}a}\right] \exp{(-j\frac{2\pi f}{c}(r-a))}. \tag{A.8}$$

We consider the limit for $a \to 0$:

$$\tilde{p}(r) = \lim_{a\to 0} p(r) = -\left[\frac{j2\pi f \rho_0 \tilde{Q}_0}{4\pi r}\right] \exp{(-j\frac{2\pi f}{c}r)}. \tag{A.9}$$

A source with $a \to 0$ is known as a point monopole [34]. The term $j2\pi f \rho_0 \tilde{Q}_0$ is the harmonic monopole source strength. The point monopole represents an ideal acoustic source which emits spherical waves without having physical extent itself. The relationship between the monopole source strength and sound pressure at a point is known as the *harmonic free-space Green's function* of the source:

$$G(\mathbf{x}|\mathbf{x}_0, f) = \frac{\exp{(-j\frac{2\pi f}{c}r)}}{4\pi r} \tag{A.10}$$

where $\mathbf{x}_0 \in \mathbb{R}^{3\times 1}$ is the location of the monopole, $\mathbf{x} \in \mathbb{R}^{3\times 1}$ is the location at which the pressure is computed and $r = ||\mathbf{x} - \mathbf{x}_0||$.
The point monopole represents a very simple loudspeaker model and we use it extensively thoughout the thesis in place of more complicated loudspeaker models.

Appendix B

The Discrete Fourier Transform and Convolution

Throughout the thesis we make heavy use of the discrete Fourier transformation (DFT) and convolution operators. This appendix clarifies the used notation and introduces some results of importance. It is, however, not intended to be an introduction to either the DFT or convolution as such (for an introduction to these matters, we refer to [24] or similar works). Section B.1 presents the

B.1 The discrete Fourier Transform

Consider the sequence $\mathbf{x} \in \mathbb{C}^{K \times 1}$:

$$\mathbf{x} = [\mathbf{x}[0], \mathbf{x}[1], \ldots, \mathbf{x}[K-1]]. \tag{B.1}$$

We denote the DFT of $\mathbf{x}$ as follows:

$$\hat{\mathbf{x}} = \boldsymbol{\mathcal{F}}\mathbf{x} = [\hat{\mathbf{x}}[0], \hat{\mathbf{x}}[1], \ldots, \hat{\mathbf{x}}[K-1]], \tag{B.2}$$

where $\boldsymbol{\mathcal{F}} \in \mathbb{C}^{K \times K}$ is the DFT. There are multiple possible definitions of the DFT. In this thesis we use the following definition:

$$\hat{\mathbf{x}}[k] = (\boldsymbol{\mathcal{F}}\mathbf{x})[k] = \frac{1}{\sqrt{K}} \sum_{l=0}^{K-1} \mathbf{x}[l] \exp \frac{-j2\pi kl}{K}. \tag{B.3}$$

The DFT can be represented as a matrix with elements:

$$(\boldsymbol{\mathcal{F}})[l, k] = \frac{1}{\sqrt{K}} \exp\left(\frac{-j2\pi kl}{K}\right). \tag{B.4}$$

It can be seen that the Herimitian conjugate of $\boldsymbol{\mathcal{F}}$ is given by:

$$(\boldsymbol{\mathcal{F}}^{\mathrm{H}}\mathbf{x})[k] = \frac{1}{\sqrt{K}} \sum_{l=0}^{K-1} \mathbf{x}[l] \exp\left(\frac{j2\pi kl}{K}\right). \tag{B.5}$$

This leads us to the following convenient property of the chosen definition of the DFT.

Lemma B.1. The DFT as defined by (B.3) is unitary:

$$\boldsymbol{\mathcal{F}}^{\mathrm{H}} = \boldsymbol{\mathcal{F}}^{-1} \tag{B.6}$$

Proof. To prove that $\boldsymbol{\mathcal{F}}$ is unitary, we must show that:

$$\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{F}}^{\mathrm{H}} = \boldsymbol{\mathcal{F}}^{\mathrm{H}}\boldsymbol{\mathcal{F}} = \mathbf{I}, \tag{B.7}$$

where $\mathbf{I} \in \mathbb{C}^{K\times K}$. Define $\mathbf{c}$, $\mathbf{d} \in \mathbb{C}^{K\times 1}$:

$$\mathbf{c} = \boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{F}}^{\mathrm{H}}\mathbf{x} \tag{B.8}$$

$$\mathbf{d} = \boldsymbol{\mathcal{F}}^{\mathrm{H}}\boldsymbol{\mathcal{F}}\mathbf{x}. \tag{B.9}$$

The elements of $\mathbf{c}$ are given by:

$$\begin{aligned}
\mathbf{c}[k] &= \sum_{l=0}^{K-1} \left[(\boldsymbol{\mathcal{F}}^{\mathrm{H}}\mathbf{x})[l] \exp\left(\frac{-j2\pi kl}{K}\right)\right] \\
&= \frac{1}{\sqrt{K}} \sum_{l=0}^{K-1} \left[\frac{1}{\sqrt{K}} \sum_{i=0}^{K-1} \left[x[i] \exp\left(\frac{j2\pi il}{K}\right)\right] \exp\left(\frac{-j2\pi kl}{K}\right)\right] \\
&= \frac{1}{K} \sum_{i=0}^{K-1} \left[\mathbf{x}[i] \sum_{l=0}^{K-1} \left[\exp\left(\frac{j2\pi(i-k)l}{K}\right)\right]\right] \\
&= \frac{1}{K} \sum_{i=0}^{K-1} \mathbf{x}[i] K\delta[i-k] \\
&= \mathbf{x}[k].
\end{aligned} \tag{B.10}$$

This implies that $\mathbf{c} = \mathbf{x}$. A similar argument shows that $\mathbf{d} = \mathbf{x}$. We therefore have:

$$\mathbf{c} = \mathbf{d} = \mathbf{x} \tag{B.11}$$

which implies that:

$$\boldsymbol{\mathcal{F}}\boldsymbol{\mathcal{F}}^{\mathrm{H}}\mathbf{x} = \boldsymbol{\mathcal{F}}^{\mathrm{H}}\boldsymbol{\mathcal{F}}\mathbf{x} = \mathbf{I}\mathbf{x}. \tag{B.12}$$

Since this holds for any $\mathbf{x} \in \mathbb{C}^{K\times 1}$, we have shown that $\boldsymbol{\mathcal{F}}$ is unitary. □

The fact that the DFT operator is unitary implies that:

$$||\mathbf{x}||^2 = ||\boldsymbol{\mathcal{F}}\mathbf{x}||^2. \tag{B.13}$$

That is, the energy of the DFT of a signal is the same as the energy of the signal itself.

It is worth noting that while the DFT can be represented as a matrix, it is rarely practical to apply it to a vector by use of conventional matrix multiplication. Typically the application of $\boldsymbol{\mathcal{F}}$ is carried out by use of Fast Fourier Transformation (FFT) [47].

B.2 Discrete Fourier transformation with Zero Padding

The frequency grid of the DFT is determined by the length of the transformed signal. In some cases this can be inconvenient. By adding zeros to the end of the signal, we can obtain a frequency representation with a more dense frequency grid. We may represent zero padding with the following operator:

$$\mathbf{I}_{\{K,L\}} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}, \tag{B.14}$$

where $\mathbf{I} \in \mathbb{C}^{L \times L}$ is the indentity operator and $\mathbf{0} \in \mathbb{C}^{M-L \times L}$ is the zero matrix. We use this to define the zero-padded DFT:

$$\boldsymbol{\mathcal{F}}_{\{M,L\}} = \boldsymbol{\mathcal{F}}\,\mathbf{I}_{\{M,L\}}. \tag{B.15}$$

The operator $\boldsymbol{\mathcal{F}}_{\{M,L\}}$, with $L \leq M$, takes a vector in $\mathbb{C}^{L \times 1}$, zero-pads to obtain a vector in $\mathbb{C}^{M \times 1}$ and takes the DFT of this vector.

Lemma B.1. The left inverse of $\boldsymbol{\mathcal{F}}_{\{K,L\}}$ is given by $\boldsymbol{\mathcal{F}}^{\mathrm{H}}_{\{K,L\}}$.

Proof. For any $\mathbf{x} \in \mathbb{C}^{L \times 1}$:

$$\begin{aligned} \boldsymbol{\mathcal{F}}^{\mathrm{H}}_{\{K,L\}}\boldsymbol{\mathcal{F}}_{\{K,L\}}\mathbf{x} &= \mathbf{I}^{\mathrm{H}}_{\{K,L\}}\boldsymbol{\mathcal{F}}^{\mathrm{H}}\boldsymbol{\mathcal{F}}\,\mathbf{I}_{\{K,L\}}\mathbf{x} \\ &= \mathbf{I}^{\mathrm{H}}_{\{K,L\}}\mathbf{I}_{\{K,L\}}\mathbf{x} \\ &= \mathbf{x}. \end{aligned} \tag{B.16}$$

This completes the proof. □

The converse is not true in general:

$$\boldsymbol{\mathcal{F}}^{\mathrm{H}}_{\{K,L\}}\boldsymbol{\mathcal{F}}_{\{K,L\}}\mathbf{x} \neq \mathbf{x}, \tag{B.17}$$

for some $\mathbf{x} \in \mathbb{C}^{K \times 1}$.

B.3 Convolution Matrices

We denote discrete linear convolution of $\mathbf{x} \in \mathbb{C}^{K\times 1}$ with $\mathbf{h} \in \mathbb{C}^{L\times 1}$ as follows:

$$\mathbf{y} = \mathbf{h} * \mathbf{x}, \tag{B.18}$$

where $\mathbf{y} \in \mathbb{C}^{(K+L-1)\times 1}$. The entries of $\mathbf{y}$ are given by:

$$\mathbf{y}[n] = \sum_i \mathbf{x}[n-i]\mathbf{h}[i], \tag{B.19}$$

where the sum runs over all values of i for which the indices are within bounds. We may express the convolution as a matrix multiplication:

$$\mathbf{y} = \mathbf{H}\mathbf{x}, \tag{B.20}$$

where $\mathbf{H} \in \mathbb{C}^{(K+L-1)\times K}$ is the Toeplitz matrix given by:

$$\mathbf{H} = \begin{bmatrix} \mathbf{h}[0] & 0 & \dots & 0 & 0 \\ \mathbf{h}[1] & \mathbf{h}[0] & & \vdots & 0 \\ \vdots & \mathbf{h}[1] & \ddots & 0 & \vdots \\ \mathbf{h}[L-2] & \vdots & \ddots & \mathbf{h}[0] & 0 \\ \mathbf{h}[L-1] & \mathbf{h}[L-2] & & \mathbf{h}[1] & \mathbf{h}[0] \\ 0 & \mathbf{h}[L-1] & \ddots & \vdots & \mathbf{h}[1] \\ 0 & 0 & \ddots & \mathbf{h}[L-2] & \vdots \\ \vdots & \vdots & & \mathbf{h}[L-1] & \mathbf{h}[L-2] \\ 0 & 0 & \dots & 0 & \mathbf{h}[L-1] \end{bmatrix}. \tag{B.21}$$

B.4 Circulant Matrices

A type of matrix similar to the convolution matrices discussed above is the circulant matrices. A circulant matrix is a matrix of the form:

$$\mathbf{C} = \begin{bmatrix} \mathbf{h}[0] & \mathbf{h}[L-1] & \dots & \mathbf{h}[2] & \mathbf{h}[1] \\ \mathbf{h}[1] & \mathbf{h}[0] & \dots & \mathbf{h}[3] & \mathbf{h}[2] \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{h}[L-2] & \mathbf{h}[L-3] & \dots & \mathbf{h}[0] & \mathbf{h}[L-1] \\ \mathbf{h}[L-1] & \mathbf{h}[L-2] & \dots & \mathbf{h}[1] & \mathbf{h}[0] \end{bmatrix} \in \mathbb{C}^{L\times L}. \tag{B.22}$$

While convolution matrices carry out linear convolution, a circulant matrix carries out circular convolution. Circulant matrices have the property of being diagonalised by the DFT [41]:

$$\boldsymbol{\mathcal{F}}^{\mathrm{H}}\mathbf{C}\boldsymbol{\mathcal{F}} = \Lambda, \tag{B.23}$$

where Λ is a diagonal matrix. The diagonal entries are the eigenvalues of $\mathbf{C}$ which coincide with the DFT coefficients of the first column of the matrix.
A simple property of convolution and circulant matrices is that:

$$\mathbf{H}^{\mathrm{H}}\mathbf{H} = \mathbf{C}, \tag{B.24}$$

where $\mathbf{H}$ is a convolution matrix and $\mathbf{C}$ is a circulant matrix. This fact is rather easily seen by writing out the matrix multiplication.

B.5 Block Matrices

We assume the reader to be more or less familiar with the concept of block matrices. We do, however, show one simple result which is necessary for the developments of the thesis. The result concerns the eigenvalues of self-adjoint block diagonal matrices. Such a matric takes the form:

$$\widetilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_0 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_1 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{A}_{M-1} \end{bmatrix} \in \mathbb{C}^{MN \times MN}, \tag{B.25}$$

where $\mathbf{A}_m \in \mathbb{C}^{N \times N}$ is self-adjoint. Let the set eigenvalues of $\mathbf{A}_m$ be denoted by $\sigma(\mathbf{A}_m)$. The eigenvalues and eigenvectors of $\widetilde{\mathbf{A}}$ are connected to those of the individual block by the following result:

Lemma B.1. The spectrum of $\widetilde{\mathbf{A}}$ is given by:

$$\sigma(\widetilde{\mathbf{A}}) = \sigma(\mathbf{A}_0) \cup \sigma(\mathbf{A}_1) \cup \ldots \cup \sigma(\mathbf{A}_{M-1}). \tag{B.26}$$

Consider eigenvalue $\lambda \in \sigma(\mathbf{A}_m)$ and the corresponding eigenvector of $\mathbf{A}_m$, $\mathbf{v} \in \mathbb{C}^{N\times 1}$. By the above, λ is also an eigenvalue of $\widetilde{\mathbf{A}}$, but the corresponding eigenvector is obviously not the same (due to dimensionality). The corresponding eigenvector is given by:

$$\tilde{\mathbf{v}} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{v} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \in \mathbb{C}^{NM \times 1}, \tag{B.27}$$

where the vector $\mathbf{v}$ is placed in entires Nm to $N(m+1)-1$.

Proof. The above is easily seen by considering that:

$$\widetilde{\mathbf{A}}\tilde{\mathbf{v}} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \mathbf{A}_m\mathbf{v} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \lambda\mathbf{v} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} = \lambda\tilde{\mathbf{v}}. \tag{B.28}$$

Thus, λ is an eigenvalue of $\widetilde{\mathbf{A}}$ and $\tilde{\mathbf{v}}$ is the corresponding eigenvector. The matrix $\widetilde{\mathbf{A}}$ can have at most NM eigenvalues/eigenvectors. Each of the M blocks can contribute N eigenvalues/eigenvectors, giving a total of NM. Thus, $\widetilde{\mathbf{A}}$ cannot have other forms of eigenvectors. □

Appendix C

Codebase and Other Material on the CD

The thesis is supplied with a CD which contains various supplements to the material presented in the thesis. The CD contains:

- A digital copy of the thesis.
- A scan of the signed study plan which was handed in at the beginning of the semester.
- A copy of the source code for generating all plots in the report as well as the additional material on the CD.
- The results generated by the supplied source code, including:
  - All the plots shown in the thesis.
  - A selection of audio demonstrations of the discussed algorithms.
  - A `.gif` animation similar to Figure 3.2 on page 34 but for changing virtual source location.

This appendix serves to provide an overview of the material on the CD as well as the system used to generate the material. A short introduction to the hardware and third party software used for the simulations are given in Section C.1. A short overview of the supplied source code is given in Section C.2. An overview of the provided audio demonstrations are given in Section C.3

C.1 Platform and Python Environment Set-up

The important hardware specifications of the used workstation is shown in Table C.1. The used software components are shown in Table C.2.

Workstation	Antec (AAU no. 78464)
CPU	Intel Core-i7-920 @ 2.67GHz (4 cores / 8 threads)
RAM	3×2 GiB DDR3 @ (3×)400 Mhz

Table C.1: General specifications for the workstation used for testing.

C.2 Software

The supplied source code is listed in Figure C.1. Instructions on how to run the software is found in the `README`-file in the source code directory.

```
<CD>/Software/Source/
├── room_impulse/                  # Source code for room simulation module.
│   ├── _room_impulse.c            # C-functions for room simulation module.
│   ├── _room_impulse.h            # Header for _room_impulse.c.
│   ├── room_impulse.pyx           # Cython interface to C-functions.
│   └── room_impulse-setup.py      # Compile script for room simulation module.
├── run_all.sh                     # Bash script for producing all results.
├── parameters.py                  # Global parameters and other settings.
├── room.py                        # Class for room simulation.
├── vogel.py                       # Module for Vogel distribution (and Figure 2.3, right).
├── cgls.py                        # Module for conjugate gradient algorithm.
├── setup_layout.py                # Figure 2.3, left.
├── room_demo.c                    # Figures 2.7, 2.8.
├── room_resp_compare.py           # Figure 2.6.
├── methods_single_frequency.py    # Figures 3.1, 3.2 3.3.
├── time_domain_aliasing.py        # Figures 4.3, 4.4.
├── methods_wideband.py            # Figures 5.1, 5.2, 5.3, 5.4, 5.5.
├── filter_length.py               # Figure 5.6.
├── vsource_position.py            # Figure 5.7.
├── gif_vsource_position.py        # Animation of changing virtual source position.
├── audio_comparison.py            # Audio demonstrations.
└── frontpage.py                   # Generates the front page to the report.
```

Figure C.1: The source code as found on the CD.

Linux Mint 16 64-bit	2.0.14	Distribution of the Linux operating system `http://www.linuxmint.com/`
Kernel	3.11.0-12-generic	
Ananconda	1.9.1 (64bit)	Python distribution by Continuum Analytics `http://continuum.io/`
Python	2.7.6	Interpreter for the Python programming language `http://www.python.org/`
`numpy`	1.8.1	Fundamental package for scientific computing with Python `http://www.numpy.org/`
`scipy`	0.13.3	Package of Python-based software tools for science and engineering `http://www.scipy.org/`
`matplotlib`	1.3.1	A Python 2D plotting library `http://matplotlib.org/`
`cython`	0.20.1	An optimising static compiler for the Cython language. `http://cython.org/`
Python Packages not in Anaconda		
`scikits.audiolab`	0.11.0	Scikit for reading and writing audio files `http://cournape.github.io/audiolab/`
`scikits.samplerate`	0.3.3	Scikit for sample rate coversion `http://www.ar.media.kyoto-u.ac.jp/members/david/softwares/samplerate/`
ImageMagick	6.7.7	Command line tool for image editing (used for making `.gif` animations) `http://www.imagemagick.org/`
-	-	-

Table C.2: Important software tools.

C.3 Audio Demonstrations

A selection of audio demonstrations of selected sound zone methods are provided on the CD (in `<CD>/Software/Output/audio_comparison/`). A number of audio clips from the Sound Quality Assessment Material (SQAM) collection has been used [1]. The clips are

[1]SQAM is a collection of audio clips for quality assessment of audio quality. The entire collection can be downloaded on `https://tech.ebu.ch/publications/sqamcd`. A manual for the collection is provided

listed in Table C.3. The audio demonstrations have been carried out for the system

SQAM no.	Description
39	Grand piano
49	Female speech, English
50	Male speech, English
69	ABBA
70	Eddie Rabbitt

Table C.3: The clips used for performing audio comparison of sound zone methods.

specified in Table 2.1 on page 14. The SQAM clips have been re-sampled to match the sampling rate of this system. Each demonstration clip is composed of two short segments. First, an audio clip is played, as heard in the middle of the bright zone. Secondly, the same clip is repeated, as heard in the silent zone. The "recordings" are made with ideal microphones and is recorded in two channels spaced by 20 cm and centered on the middle of the given zone. It is advised that headphones are used for listening to the demonstrations. The files are names are given as follows:

```
<SQAM no.>_<sound zone method>(_<filter length>)(_changed).wav
```

The first section specifies the audio clip used (see Table C.3). The second section specifies the method used to create sound zones. The third section specifies the used filter length but is only applicable to the CG-PM method (as this is the only algorithm for which this is a parameter). Lastly, "`_changed`" is added when the recording has been made after slightly changing the room conditions (in the same manner as discussed in the thesis). E.g. `49_cgls_200_changed.wav` is a demonstration of the CG-PM method with filter length 200 under changed room conditions using the female speech audio clip.

on `https://tech.ebu.ch/docs/tech/tech3253.pdf`.