



University of Victoria
Faculty of Engineering

ELEC 498 - Honours Thesis

**THE LEAN OPTIMIZATION OF ACOUSTIC DIFFUSERS:
DESIGN BY ARTIFICIAL EVOLUTION, TIME DOMAIN
SIMULATION AND FRACTALS**

Tim Perry
V00213455
Electrical Engineering
timperry@uvic.ca

Prepared for Dr. Peter Driessen

Dec 20, 2011

The latest version of this report can be downloaded for free at
<http://Arqen.com/acoustics/>

CONTENTS

| | |
|--|----|
| LIST OF TABLES AND FIGURES | IV |
| ABSTRACT | VI |
| 1. INTRODUCTION | 1 |
| 1.1. Objective and Scope | 1 |
| 1.2. Structure..... | 1 |
| 1.3. A brief review of diffusers and diffusion..... | 2 |
| 1.3.1. Spatial and temporal Dispersion..... | 3 |
| 2. LITERATURE REVIEW | 4 |
| 2.1. Schroeder Diffusers | 4 |
| 2.1.1. Maximum length sequence diffusers | 5 |
| 2.1.2. Quadratic residue diffusers..... | 6 |
| 2.1.3. Primitive root diffusers | 6 |
| 2.1.4. Other sequences..... | 7 |
| 2.1.5. Modulation schemes and fractal constructions..... | 7 |
| 2.1.6. Two-dimensional (hemispherical) diffusers | 9 |
| 2.1.7. Phase optimized Schroeder and stepped diffusers | 10 |
| 2.2. Geometric Diffusers | 13 |
| 2.2.1. Curved surfaces..... | 13 |
| 2.2.2. Optimized curved surfaces..... | 13 |
| 2.2.3. Fractals | 15 |
| 2.3. Hybrid Surfaces | 16 |
| 2.3.1. Planar hybrid surfaces..... | 16 |
| 2.3.2. Curved hybrid surfaces..... | 18 |
| 3. THE LEAN DIFFUSER OPTIMIZATION PROBLEM..... | 19 |
| 3.1. Design objectives | 19 |
| 3.2. Feasible approaches to optimal diffuser design | 20 |
| 3.3. Practical pitfalls to previous methods of diffuser optimization | 22 |
| 3.4. Design method | 23 |
| 4. PREDICTION OF SCATTERING USING FINITE DIFFERENCE TIME DOMAIN..... | 25 |
| 4.1. A finite difference scheme with a Ricker wavelet source | 26 |
| 4.1.1. Numerical stability..... | 28 |

| | | |
|--------|--|----|
| 4.1.2. | Representing a diffusing surface | 28 |
| 4.1.3. | Excitation | 28 |
| 4.1.4. | Receiver Arc..... | 30 |
| 4.1.5. | Anechoic boundary conditions..... | 32 |
| 4.2. | A simulation environment built with k-Wave for Matlab | 32 |
| 4.2.1. | Adapting k-Wave to simulate scattering..... | 33 |
| 4.2.2. | Flattening the spectral envelope..... | 34 |
| 5. | MEASURES OF DIFFUSION QUALITY | 37 |
| 5.1. | A diffusion parameter based on the standard error | 37 |
| 5.2. | The autocorrelation diffusion coefficient | 38 |
| 6. | THE PRELIMINARY DESIGN OF STEPPED DIFFUSERS | 40 |
| 6.1. | Limitations and considerations for phase grating surfaces..... | 41 |
| 6.2. | Geometric framework for stepped diffusers..... | 42 |
| 7. | OPTIMIZATION | 44 |
| 7.1. | Genetic algorithms..... | 44 |
| 7.2. | Optimization using an integer genetic algorithm..... | 45 |
| 7.2.1. | Tuning and convergence..... | 46 |
| 7.3. | Optimization results | 47 |
| 8. | BANDWIDTH EXTENSION VIA FRACTALIZATION | 51 |
| 9. | SIMULATION AND ANALYSIS OF THE DESIGNS | 54 |
| 9.1. | Feedback from the specular zone..... | 56 |
| 9.2. | Performance results | 57 |
| 9.3. | Analysis of optimized designs at a receiver radius of 5 m..... | 58 |
| 9.3.1. | Applying profiled modulations to the optimized base shape..... | 63 |
| 9.4. | Analysis of fractal designs at a receiver radius of 5 m | 65 |
| 9.5. | Comparing simple optimized diffusers..... | 69 |
| 10. | CONCLUSIONS | 71 |
| 10.1. | Enhancements and Future Research..... | 73 |
| | REFERENCES..... | 75 |
| | APPENDIX A FINITE DIFFERENCE TIME DOMAIN MATLAB CODE..... | 78 |
| | APPENDIX B PREDICTION OF SCATTERING MATLAB CODE..... | 79 |
| | APPENDIX C MATLAB SCRIPTS FOR DIFFUSER OPTIMIZATION | 80 |
| | APPENDIX D SCATTERED POLAR DISTRIBUTION PLOTS..... | 81 |

LIST OF TABLES AND FIGURES

TABLES

| | |
|---|----|
| Table 1 Prospective Design Approaches..... | 21 |
| Table 2 Preliminary Calculations for Phase Grating Diffusers..... | 43 |
| Table 3 Optimization Frameworks..... | 47 |
| Table 4 Optimization Results..... | 48 |
| Table 5 Framework for Designing Fractal Stepped Diffusers | 51 |
| Table 6 Fractal Formations | 51 |
| Table 7 Simulation Frameworks for Full-Scale Analysis | 54 |
| Table 8 Performance of Diffusers at a Receiver Radius of 5m..... | 57 |

FIGURES

| | |
|---|----|
| Figure 1 The “stick room” achieves exceptional diffusion, but invites eye injury. | 2 |
| Figure 2 Spatial and temporal dispersion generated by a Schroeder diffuser..... | 3 |
| Figure 3 Temporal and frequency response for flat (top) and diffusive (bottom) surfaces..... | 3 |
| Figure 4 Scattered pressure from a Schroeder diffuser (left) and plane surface (right). | 4 |
| Figure 5 A one-dimensional Schroder diffuser (After Cox and D’Antonio [6])..... | 5 |
| Figure 6 Cross-section profile through a single period of an $N = 7$ MLS diffuser. | 5 |
| Figure 7 Cross-section profile of an $N = 7$ QRD® (After Cox and D’Antonio [1]). | 6 |
| Figure 8 A basic aperiodic modulated diffuser. | 8 |
| Figure 9 Effects of periodicity and modulation (After Cox and D’Antonio [1]). | 8 |
| Figure 10 The QRD Diffractal® by RPG Diffusor Systems. | 9 |
| Figure 11 Hemidisk scattering pattern for a one-dimensional QRD (left)..... | 10 |
| Figure 12 Optimized stepped diffuser for the rear wall of a performance facility..... | 12 |
| Figure 13 Process to find an optimal well depth sequence for a phase grating diffuser. | 12 |
| Figure 14 Shape optimization process (Source: rpginc.com [3]). | 14 |
| Figure 15 Fourier synthesis generation of a surface that represents Brownian motion..... | 15 |
| Figure 16 Three fractal surfaces generated by Fourier synthesis. | 15 |
| Figure 17 Fractal generation using randomly displaced step functions. | 16 |
| Figure 18 Assembly of the Binary Amplitude Diffusorber, a planar hybrid surface. | 17 |
| Figure 19 Staggered pressure and particle velocity vectors in a FDTD mesh..... | 26 |
| Figure 20 Ricker wavelets and their spectra centered on 250 Hz and 2,000 Hz. | 29 |
| Figure 21 2D FDTD simulation of an $N=17$ stepped diffuser excited by a Ricker wavelet.. | 30 |
| Figure 22 Simulation domain used for preliminary testing with 180 sensors over 180° | 31 |
| Figure 23 Near field temporal response to a Ricker wavelet source at normal incidence. .. | 31 |

| | | |
|-----------|---|----|
| Figure 24 | Near field scattering from a randomly generated stepped diffuser..... | 33 |
| Figure 25 | Temporal response and spectrum of a diffuser measured at the central sensor. | 35 |
| Figure 26 | Left: Reference pulse spectrum and response spectrum. | 36 |
| Figure 27 | A genetic algorithms for integer sequence optimization. | 45 |
| Figure 28 | Minimization progress after 1200 scattering predictions for 5 modules..... | 49 |
| Figure 29 | Near field simulation of the optimized array with base shape A1-LF. | 49 |
| Figure 30 | Standard error $\varepsilon(4.5\text{m}, 2.5\text{m}, f)$ for an array of optimized base shapes A1-LF. ... | 50 |
| Figure 31 | Scattered polar distribution for the optimized array with base shape A1-LF. | 50 |
| Figure 32 | Three arrays of $N = 7$ base shapes with array modulation $\{1\ 0\ 1\ 1\ 0\}$ | 52 |
| Figure 33 | Three arrays of $N = 7$ second order fractals with array modulation $\{1\ 0\ 1\ 1\ 0\}$... | 52 |
| Figure 34 | Temporal response and spectrum of A1-Frac, measured at the central sensor... | 53 |
| Figure 35 | Left: Reference pulse spectrum and response spectrum for A1-Frac. | 53 |
| Figure 36 | Full scale domain used to simulate the scattering from fractal diffusers..... | 55 |
| Figure 37 | Pressure received at each sensor during simulation of the A1-Frac array..... | 55 |
| Figure 38 | Polar response at 5 m from a reflector the same width as the diffuser array..... | 56 |
| Figure 39 | Diffusion parameter $\varepsilon(10\text{m}, 5\text{m}, f)$ for the diffuser array A1-LF. | 59 |
| Figure 40 | Diffusion parameter $\varepsilon(10\text{m}, 5\text{m}, f)$ for the diffuser array B2-LF. | 59 |
| Figure 41 | Scattering from the A1-LF periodic array at a receiver radius of 5 m. | 59 |
| Figure 42 | Scattered polar distribution for the optimized diffuser array A1-LF..... | 60 |
| Figure 43 | Scattered polar distribution for the optimized diffuser array B2-LF..... | 60 |
| Figure 44 | Diffusion coefficient spectra for A1-LF. Diffuse field (left). 0° incidence (right). | 61 |
| Figure 45 | Diffusion coefficient spectra for B2-LF. Diffuse field (left). 0° incidence (right). | 61 |
| Figure 46 | Diffuse field performance report for A1-LF array. Produced in AFMG Reflex.... | 62 |
| Figure 47 | Diffuse field performance report for B2-LF array. Produced in AFMG Reflex.... | 62 |
| Figure 48 | Profiled Modulation 1 coefficient spectra. Diffuse field (left). 0° incidence (right).... | 63 |
| Figure 49 | Profiled Modulation 2 coefficient spectra. Diffuse field (left). 0° incidence (right).... | 64 |
| Figure 50 | Simulation of near field scattering from fractal diffusers A1-Frac. | 65 |
| Figure 51 | Simulation of near field scattering from fractal diffusers B2-Frac. | 65 |
| Figure 52 | Diffusion parameter $\varepsilon(10\text{m}, 5\text{m}, f)$ for fractals based on optimized diffusers. | 66 |
| Figure 53 | Scattering from the A1-Frac array..... | 67 |
| Figure 54 | Scattering from the A1-Frac array (left) and B2-Frac array (right)..... | 68 |
| Figure 55 | Scattered polar distribution for the diffuser array A1-Frac. | 68 |
| Figure 56 | Scattered polar distribution for the diffuser array B2-Frac. | 68 |
| Figure 57 | Diffuse field performance report for A1-LF array. Produced in AFMG Reflex.... | 69 |
| Figure 58 | Diffuse field performance report for 5 periods of optimized diffuser L95 [11]. | 70 |
| Figure 59 | Profiled Modulation 1 coefficient spectra. Diffuse field (left). 0° incidence (right).... | 70 |

ABSTRACT

A simple framework has been developed to optimize acoustic diffusers in reasonable time without the need for boundary element predictions. The approach combines evolutionary optimization and time domain simulation to design shallow, profiled surfaces that create a large amount of diffusion. The “lean” optimization uses an integer genetic algorithm to find candidate designs in a low resolution design space. It compares candidates using a finite difference time domain model to predict diffusion performance. The process has been shown to produce diffusers that offer an excellent trade-off between performance and compact geometry. Fractal forms have been generated from these results to extend the bandwidth over which diffusion occurs. The new optimized and fractal diffusers are compact, modular, and based on the set of integers between zero and 16. This makes them practical to simulate with high accuracy using finite difference time domain and simple to construct using low precision manufacturing.

1. INTRODUCTION

An ideal acoustic diffuser is a surface that causes an incident sound wave from any direction to be evenly scattered in all directions. Until recently the design of diffusers was practiced by a few knowledgeable acousticians—yet many enthusiasts and music industry professionals saw merit in emulating the designs. With the publication of *Acoustic Absorbers and Diffusers* (2004), Cox and D’Antonio [1] have brought diffuser design techniques to a larger audience. The result has been an explosion of Schroeder [2] diffusers in the professional audio marketplace. However, the design of optimized diffusers has remained the domain of experts; notably, the industry’s leading innovator, RPG Diffuser Systems [3]. It appears that diffuser optimization is avoided by designers and acoustical engineers because it requires a sophisticated framework. The heart of this framework is a simulation to predict acoustic scattering. For this a boundary element model is the natural—but not always viable—choice.

1.1. Objective and Scope

This work focusses on solving a practical problem that will be called “the lean diffuser optimization problem”. The design objective is to answer the following question:

What modular sound diffuser provides an ‘optimal’ trade-off between uniform scattering and compact geometry, and how can this surface be discovered without using boundary element predictions?

‘Optimal’ in this context does not mean ideal. It implies that within a given framework, optimization is used to explore the solution space (the set of candidate solutions) and an excellent candidate is picked. The quality of the solution will depend on how thoroughly the optimization process searches the solution space, which will depend on time constraints. In this work a practical optimization run is generally not expected to find the global optimum. The design framework will allow a variety of problems to be solved; therefore, the method is at least as significant as the resulting diffuser designs.

1.2. Structure

Key points from the literature are condensed in Chapter 2. Based on this, the lean optimization problem and the design method used to solve it are presented in Chapter 3.

Chapter 4 covers the implementation of a finite difference time domain model for scattering prediction, and Chapter 5 presents the measures of diffusion quality used for evaluation. After preliminary design considerations in Chapter 6, Chapter 7 demonstrates how the optimization problem was solved using an integer genetic algorithm. The designs are enhanced in Chapter 8 via the self-symmetry properties of fractals. Finally, in Chapter 9 a larger scale time domain simulation is used to further assess the optimized designs, and the results are interpreted to reveal the winning diffuser geometries.

1.3. A brief review of diffusers and diffusion

When designing a room with high quality acoustics, one of the primary goals is to achieve a diffuse reverberant sound field. Diffusers are used in studios and live music spaces to prevent specular reflections that would interfere with critical listening (Figure 3), and to provide a controlled reverberation, or ambience. They are functionally mounted in plain sight, therefore acousticians are interested in developing designs with various visual aesthetics to expand the palette of forms available to the end-user—perhaps a musician with a home studio, or an architect who needs a diffuser to blend in with a new building [4]. Forms with a low profile are desired for the lean optimization problem, and forms that pose obvious hazards to human safety are automatically disqualified. A diffuser shaped like an array of icicles, for example, would pose an eye hazard when mounted on a wall (Figure 1).



Figure 1 The “stick room” achieves exceptional diffusion, but invites eye injury. Most control rooms use significant absorption to create a reflection free zone (RFZ) around the listener. In contrast, this space at Blackbird Studios uses liberal 2D diffusion to achieve a level of clarity that may be acceptable for critical listening (depending on one’s school of thought). One-inch square pegs with lengths varying from 6 to 40 inches are used to realize a design based on the prime number 138,167. (Source: Digizine, 2011) [5]

1.3.1. Spatial and temporal Dispersion

A ‘diffuse reflection’ is one dispersed in both space (spatial dispersion) and time (temporal dispersion), as depicted in Figure 2. Diffusers are often designed to have uniform spatial dispersion, assuming that temporal dispersion will also occur [1]. Temporal variation is less convenient as a design parameter because it does not guarantee uniform spatial dispersion and often introduces colouration to the frequency spectrum (Figure 3). An ideal diffuser will generate both uniform spatial and temporal dispersion over all audible frequencies.

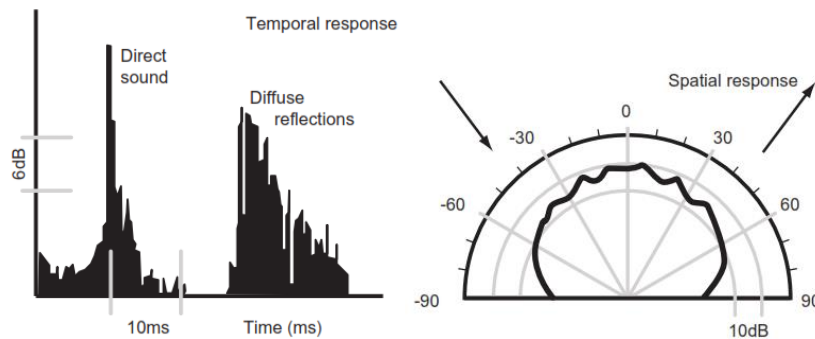


Figure 2 Spatial and temporal dispersion generated by a Schroeder diffuser. Temporal dispersion (left) can be interpreted from an impulse response plot; spatial dispersion (right) can be interpreted from a polar plot (After Cox and D'Antonio [1]).

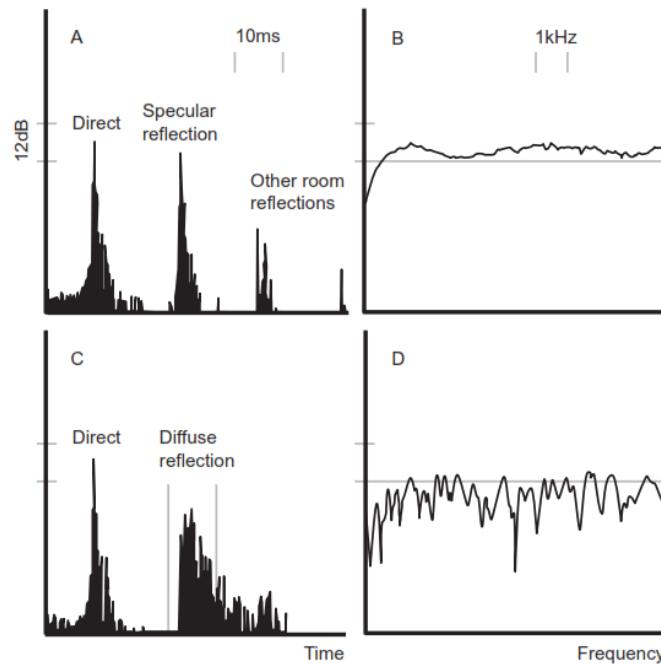


Figure 3 Temporal and frequency response for flat (top) and diffusive (bottom) surfaces. The frequency response is shown for the reflected sound only. The frequency response for the flat surface is characterized by a high pass filter response, and for the diffuser exhibits peaks and nulls spaced irregularly with respect to frequency (After Cox and D'Antonio [1]).

2. LITERATURE REVIEW

This chapter explores existing methods of diffuser design and summarizes their relative merits. Methods that optimize a surface for uniform scattering are emphasized.

2.1. Schroeder Diffusers

Schroeder diffusers are designed using convenient mathematical principles that allow them to be constructed as a series of wells separated by thin fins. The wells have equal width and different depths, with depths determined by a theoretic number sequence. The maximum well depth and well width are commonly used to define the bandwidth for predictable dispersion—but in reality, quality diffusion may extend beyond the predicted upper limit [1].

While commercially successful to date, many acousticians are reluctant to use these designs as they do not visually complement modern architecture. Additionally, the number theoretic designs have performance limitations, most notably:

- They only achieve ‘optimum’ dispersion at discrete frequencies.
- They are designed based on simplified theory: ‘optimum diffusion’ means equal energy in the diffraction lobes [1]. This is not the same as uniform scattered energy in all directions.

Optimization algorithms can be used to improve the design of Schroeder diffusers, with the ultimate goal being uniform broadband dispersion.

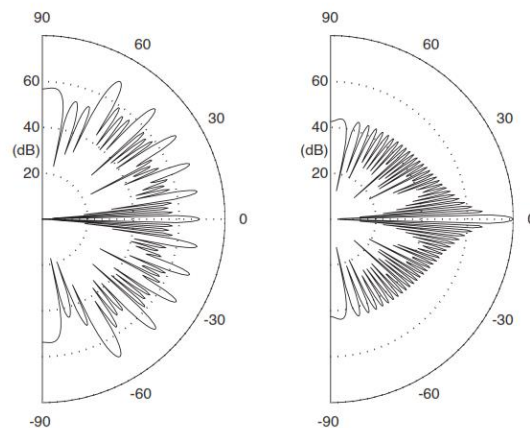


Figure 4 Scattered pressure from a Schroeder diffuser (left) and plane surface (right).
(After Cox and D’Antonio [1])




Figure 5 A one-dimensional Schroeder diffuser (After Cox and D'Antonio [6]).

2.1.1. Maximum length sequence diffusers

In 1975, Schroeder proposed constructing diffusers based on Maximum Length Sequences (MLS) [2]. He justified this using a fact from optics theory: the far field scattering can be approximately predicted by taking the Fourier transform of a 'surface', thus the power spectrum and surface scattering are closely related [1]. The MLS was chosen as it has a flat power spectrum at all frequencies.

One dimensional MLS diffusers consist of strips of material with two different depths, placed according to an MLS. For example, one period of an $N = 7$ MLS surface could be based on the sequence [0, 0, 1, 0, 1, 1, 1]. The diffuser shape is represented by a box with variable admittance on the front surface.

The admittance of the surface is determined from plane wave propagation in the wells, leading to design equations that relate physical dimensions to dispersion performance. This makes it straightforward to design a surface that achieves maximum scattering at a specific frequency.

An octave above the design frequency, MLS diffusers exhibits specular reflection. At this critical frequency the phase-grating fails because the well depth is half a wavelength, causing waves to re-radiate with the same phase. To solve this narrow-bandwidth problem, Schroeder suggested different number sequences, such as the quadratic residue sequence.




Figure 6 Cross-section profile through a single period of an $N = 7$ MLS diffuser. (After Cox and D'Antonio [1])

2.1.2. Quadratic residue diffusers

Quadratic residue diffusers (QRD) are designed to extend the dispersion characteristics of MLS diffusers over a wider bandwidth. The QRD achieves optimum scattering at integer multiples of the design frequency, and generally achieves reasonable dispersion in between these frequencies [1]. Good dispersion over more frequencies can be achieved by using orthogonal modulations, resulting in diffusers with two different design frequencies [7,1].




Figure 7 Cross-section profile of an $N = 7$ QRD® (After Cox and D'Antonio [1]).

2.1.3. Primitive root diffusers

Primitive root diffusers (PRD) are designed to produce a notch in the scattering response, with even energy in the other diffraction lobes. Like the QRD, optimum scattering is only achieved at integer multiples of the design frequency. Unfortunately, the pressure nulls achieved by the PRD are located elsewhere in the spectrum.

The Cox-D'Antonio-modified primitive root diffuser (CDMPRD) is a revised notch diffuser designed to solve this problem. In effect, the technique introduces a frequency shift: the reflection coefficients are appropriately aligned around the unit circle at multiples of the design frequency, resulting in the desired pressure nulls [6,7,1].

It may seem that the PRD is useful in small spaces as a means to minimize the energy reflected in the specular reflection direction; however, as PRDs only work at discrete frequencies, they do not make practical notch filters. The PRD does remain useful as a diffuser, having similar performance to the QRD.

Numerical optimization can be applied to form a broader notch over a wider frequency range [6]—but this is difficult to do. In general, optimization struggles to shape polar responses, but is successful as a method to achieve uniform dispersion.

Triangles or pyramids can be used to achieve a more broadband notch, but they restrict the angle of incidence over which for which diffusion is effective [1].

2.1.4. Other sequences

The MLS, quadratic residue and primitive root sequences are not the only suitable sequences for diffuser design. Other promising sequences include [1]:

- *Index sequences*, which should yield diffusers with similar performance to the PRD, but with extra absorption.
- *Short power residue sequences*, which under certain conditions can be formed by undersampling longer primitive root sequences.
- The *Chu sequence*, which will yield similar performance to a QRD.
- Optimized sequences.

Promising sequences are those with good autocorrelation properties [1]. If the autocorrelation function of the reflection coefficients is a delta function, its Fourier transform will reveal a flat power spectrum. This corresponds to an even scattering distribution—in effect, a good diffuser.

2.1.5. Modulation schemes and fractal constructions

As Schroeder diffusers are periodic, the scattered energy is dominated by grating lobes. A more even scattering distribution can be achieved by making the diffuser aperiodic or by increasing the spacing between periods.

Aperiodicity will result from using a long number sequence with good autocorrelation properties. However, this is rarely a viable design practice because there are few known large aperiodic polyphase sequences [1]. Moreover, periodicity facilitates modular manufacturing and cost effective shipping. Aperiodicity does not.

One practical solution is to use a modulation scheme [1,6,7]. Typically, the best choice is to use a diffuser and its inverse, arranged to achieve aperiodicity (Figure 8).




Figure 8 A basic aperiodic modulated diffuser.
Aperiodic modulation is provided by arranging an optimal binary sequence of the base shape (binary 1) and flipped shape (binary 0) (After Cox and D'Antonio [6]).

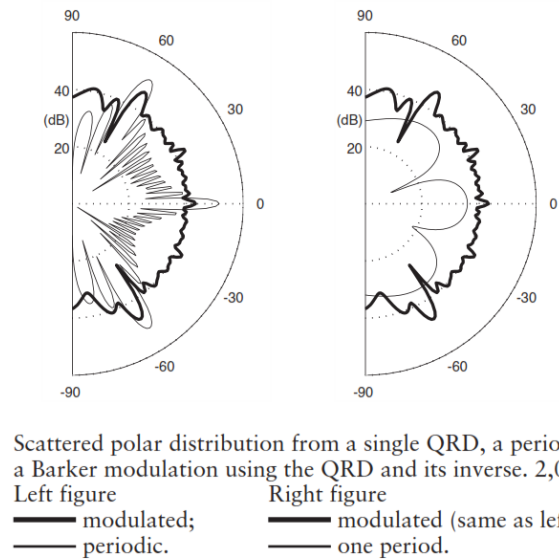


Figure 9 Effects of periodicity and modulation (After Cox and D'Antonio [1]).

Fractal formations are an elegant solution to periodicity, absorption and bandwidth problems. High frequency diffusers can be nested within low frequency diffusers, exploiting the self-symmetry property of fractals to provide full spectrum diffusion within a single device [8,1].




Figure 10 The QRD DiffRACTal® by RPG Diffusor Systems.

Left: Construction of the DiffRACTal consists of first and second generation fractals based on the quadratic residue sequence. Right: Polar scattering response for seven periods of an (a) QRD and (b) QRD DiffRACTal at the high design frequency. Simulated using near-field Kirchhoff diffraction theory (After Cox and D'Antonio [9]).

2.1.6. Two-dimensional (hemispherical) diffusers

Thus far only planar devices have been addressed. Planar diffusers scatter incident sound into a hemi-disc, while two dimensional diffusers disperse sound in a hemispherical pattern (Figure 11). 2D Schroeder diffusers are constructed using two-planes, each designed for optimal scattering. One plane scatters in the x-direction, the other scatters in the z-direction, resulting in even lobes in a hemisphere. The device typically takes the form of a grid, with cavity depths determined by applying the Chinese remainder theorem to fold two 1D sequences into a 2D array [1]. Each 1D sequence should be based on the same prime number.

As 2D diffusers scatter in two planes they deliver less scattered energy to a receiver than 1D devices, making them less efficient. Additionally, 2D diffusers constructed as a grid have higher absorption per unit area than their 1D counterparts.




Figure 11 Hemidisk scattering pattern for a one-dimensional QRD (left). Hemispherical scattering pattern for a two-dimensional QRD (right). Incident plane wave is at 45° with respect to surface normal (After Cox and D'Antonio [9]).

2.1.7. Phase optimized Schroeder and stepped diffusers

Instead of relying on an optimal number sequence with a flat power spectrum, the design of Schroeder diffusers can be improved by optimizing for uniform scattering directly. This approach combines multi-dimensional optimization techniques with boundary element predictions to design *phase optimized diffusers* (POD) [10].

The first step to optimizing the Schroeder is to remove the fins between the wells, yielding a simpler, superior design: the *stepped diffuser*. This simple modification improves dispersion performance [11] and provides these additional benefits:

- Simplified geometry reduces manufacturing costs.
- Removal of the resonant wells results in lower absorption.

The optimization process starts by randomly choosing a well depth sequence, then predicting the scattering and assessing its quality. The goal is to minimize the error between the predicted scattering and desired scattering. This is done by making incremental adjustments to the well depth sequence until a figure of merit is satisfied (in effect, minimizing the error).

To solve the diffuser optimization problem, [1] the following scaffolding must be in place:

1. A model to predict the scattering from a given diffuser design.
2. An error parameter or figure of merit to define the quality of the scattering.
3. An optimization algorithm to change the well depth sequence and search for an appropriate solution.

A boundary element model (BEM) is generally the first choice for scattering prediction in acoustical design, but there are other options including Fraunhofer models, Fourier models, finite difference time domain (FDTD) methods and finite element analysis (FEA) [1,12,13]. While a BEM can be slow to compute, the results are accurate. Older studies used Fraunhofer solutions which offer fast optimization at the expense of solution accuracy [11,13]. Another approach is to use simple models to narrow in on a solution region, and a more accurate model to compute the solution.

A single figure of merit for uniform broadband dispersion can be formed from the mean and standard deviation of the diffusion coefficients across all frequencies. This works as follows [1]:

- The diffusion coefficient at each one-third octave band is computed from the prediction model [11].
- The mean and standard deviation of the diffusion coefficient spectra are calculated.
- The standard deviation is subtracted from the mean.

Thus a penalty is applied to the figure of merit, proportional to the unevenness of the diffusion coefficient spectra. If the predicted diffusion is very uneven across all frequencies, the standard deviation will be large, and a large penalty will be applied to the figure of merit.

If the gradient of the figure of merit is known, using it can greatly speed up the optimization process. In most cases of diffuser optimization the gradient is not available, therefore suitable optimization algorithms are those that depend only on function values.

- Downhill simplex is a natural choice. While slow, it is robust to non-linear constraints and can be applied to wide range of diffuser optimization problems.
- An alternative is to use a genetic algorithm, which requires extra set up as it must be carefully tuned to a given problem.
- Quasi-Newton gradient descent methods are fast, but unreliable when combined with BEM prediction. These methods rely on approximate gradients, calculated with finite differences and data retrieved from the prediction model. As a BEM prediction will exhibit small inaccuracies, there is a high risk of numerical error propagation that will throw off the gradient.

‘Minimizing’ the error between the desired response and the figure of merit involves searching for the global minimum, or else a suitable local minima in the feasible region (the

space of all candidate solutions). As the degrees of freedom increase, the region to search becomes more complex and the global minimum becomes virtually impossible to find—but at the same time it becomes less important. Searching for a suitable solution will typically involve evaluating the scattering a thousand times [1].




Figure 12 Optimized stepped diffuser for the rear wall of a performance facility. (After Cox and D'Antonio [12])

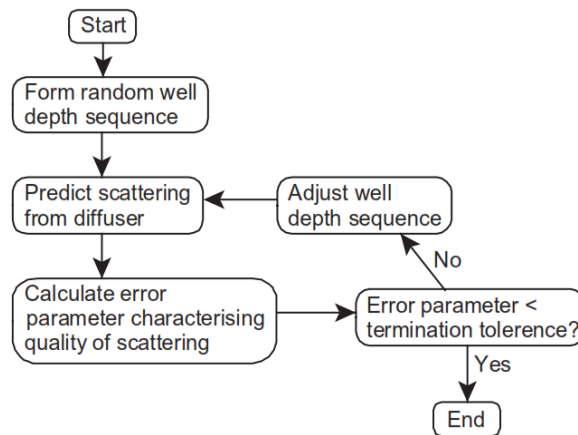


Figure 13 Process to find an optimal well depth sequence for a phase grating diffuser. (After Cox and D'Antonio [1])

2.2. Geometric Diffusers

2.2.1. Curved surfaces

Compared with other diffuser designs, curved surfaces have the potential benefit of lower construction costs and lower absorption. In theory, the cylinder appears to be the perfect diffuser design [1]; but in practice, a single-cylinder design would be too deep for most architectural applications.

Convex reflectors based on part of a circle—such as a semicylinder (ellipse)—only disperse sound well for normal incidence. Better response at oblique incidence can be achieved by forming an array of semicylinders, or by creating a more complex shape using optimization.

When placed in an array the response from a single cylinder becomes secondary to the response of the array [1]. For such an array to be effective cylinders must be spaced far apart, as randomly as possible; otherwise, modulation schemes are required to reduce periodicity.

2.2.2. Optimized curved surfaces

Curved surfaces can be optimized to meet performance and aesthetic requirements for most architectural applications. For an optimization algorithm to change the shape of the surface, the shape is first described as a set of numbers represented by a Fourier series. In theory, an infinite Fourier series can represent any shape—but for optimization to be possible, the series must be truncated. 4-6 harmonics are typically used to avoid unnecessary complexity which would increase manufacturing costs [1].

To achieve an acoustically optimized shape that satisfies physical design specifications, non-acoustical constraints are typically needed. Three constraint methods are mentioned by D'Antonio and Cox [1]:

- Fuzzy constraints check to see if a surface is sufficiently close to constraint points during optimization. If not, a penalty is applied to the error parameter. Fuzzy constraints add complexity to the optimization problem, as the error parameter depends on both scattering quality and shape quality. While this system can be used to avoid simple physical obstacles, it is inelegant as a means to satisfy a desired visual aesthetic.

- A spline construction with linear constraints can be used to simplify the above problem.
- A base shape can be designed from shape variables, and distorted to change the acoustical performance. The same compression, modulation and warping techniques used in image processing can be applied to distort the surface while preserving desired visual characteristics.

The shape optimization process is given in Figure 14. *Downhill simplex* is typically used as it is robust to non-linear constraints—but the resulting search process is slow, and may require many trials with different starting locations [14].

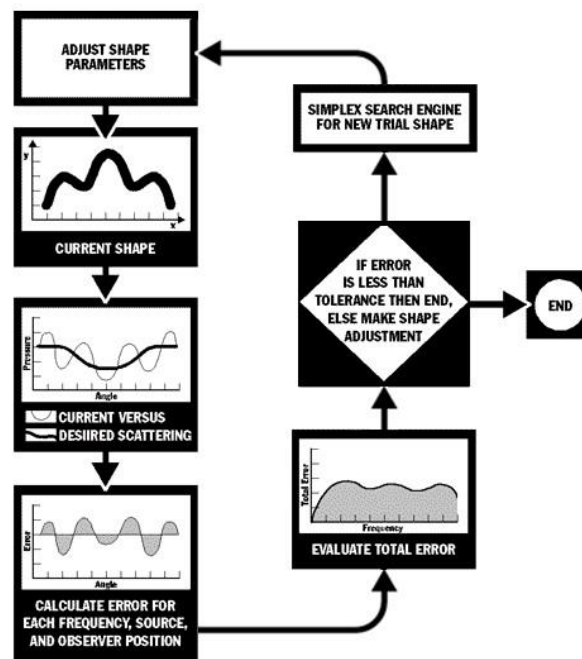


Figure 14 Shape optimization process (Source: rpginc.com [3]).

All optimized curved diffusers tested by Cox and D’Antonio [1] performed at least as good as tan arcs of a circle. In general, these surfaces have the best dispersion performance of all diffuser designs—provided that periodicity can be avoided. Periodicity may be dealt with by using a modulation scheme if diffusers are arranged in arrays, or by constructing a single large surface.

2.2.3. Fractals

High quality dispersion can be achieved by constructing fractal surfaces that simulate fractional Brownian motion. Fractional Brownian diffusers (FBD) constructed using Fourier synthesis result in complex surfaces governed by many shape parameters [8]. Such surfaces may be manufactured using extrusion, but due to the many shape variables it is not practical to optimize them for best diffusing performance.




Figure 15 Fourier synthesis generation of a surface that represents Brownian motion. By extrusion, the surface can be manufactured into a fractional Brownian diffuser, or FBD (After Cox and D'Antonio [1]).




Figure 16 Three fractal surfaces generated by Fourier synthesis. Input noise: (a) Brown noise; (b) Pink noise; (c) White noise (After Cox and D'Antonio [1]).

Fractal generation using step function addition

A series of randomly displaced step functions can be used to generate a fractal surface that simulates Brownian motion. This enables the number of parameters defining the surface to be reduced, and therefore the use of optimization techniques. True Brownian motion would require an infinite number of superimposed step functions, each having random amplitude and a random displacement along the width of the diffuser (the x-axis of Figure 17).




Figure 17 Fractal generation using randomly displaced step functions.
 (a) 20 step functions. (b) 10 step functions. (c) 1 step function (After Cox and D'Antonio [1]).

Tanh Addition Diffusers

Step function addition may generate surfaces that exhibit specular reflection. To avoid this problem, functions with smoother transitions are typically used. Tanh Addition Diffusers, also called Random Addition Diffusers (RAD), use the sum of many hyperbolic tangent functions to form surfaces that can be optimized—with excellent results [1]. The optimized fractal surfaces can achieve comparable performance to optimized curved surfaces, and better performance than the arc of a circle for random incident sound. Surprisingly, under certain conditions this fractal generation technique will yield a semicircular surface [8].

2.3. Hybrid Surfaces

Hybrid surfaces achieve variable impedance by using patches of absorptive and reflective material. These diffusers can be designed with excellent dispersion characteristics, but their applications are limited as they cannot be designed for minimum absorption. In environments that require simultaneous control of absorption and diffusion, hybrid surfaces offer attractive benefits:

- Manufacturing is simple and cheap.
- Hybrid surfaces have a shallow profile compared with other diffusers.
- Treatment can be easily hidden to blend in with any environment.
- Optimization gives complete control over the reflectivity [15].

2.3.1. Planar hybrid surfaces

The Binary Amplitude Diffsorber (or BAD™ panel) by RPG is a flat hybrid surface constructed from a porous absorber faced with a perforated mask (Figure 18). Such panels

are often highly absorptive up to about 2 kHz, thus only the dispersion performance over mid-to-high frequencies needs to be considered [16,1].

One-dimensional sequences

To maximize dispersion in a one-dimensional hybrid diffuser, the surface binary elements are distributed based on an optimal binary sequence with good autocorrelation properties [1,17]. Potential sequences include the *MLS*, *1D optical sequences*, *1D ternary* and *quadriphase sequences* [18,19].

Unfortunately, existing number sequences offer few choices for absorption coverage. For example, the MLS gives a panel open area of about 50%, which is generally more absorption coverage than desired. To solve this problem, numerical optimization is typically used to find a family of discrete sequences with low mutual cross-correlation. These are then concatenated together to form a longer sequence. Searching for this family of sequences is a discrete optimization problem, best solved using a genetic algorithm on sequences of length $N < 48$, or an exhaustive search when $N < 20$ [1,15].

Two-dimensional sequences

RPG developed a method to design hemispherical scattering hybrid surfaces by using the *Chinese Remainder Theorem* to fold 1D sequences into 2D arrays, while preserving the good autocorrelation and Fourier properties [16]. It is also possible to construct optimal binary sequences on a hexagonal array pattern [1].




Figure 18 Assembly of the Binary Amplitude Diffuser, a planar hybrid surface. Porous absorber (left), reflective mask (middle), and optional fabric covering (right). (After Cox and D'Antonio [1])

Improving the Binary Amplitude Diffusor

Payne-Johnson, Gehring and Angus suggest that there is an opportunity to improve the BAD design, proposing an easily manufactured surface using small variable-size square panels [20]. The size of these panels would be determined using an M-sequence, and they would be grouped in such a way that diffusion is preserved despite periodic effects. This

may be done with modulation techniques [20], similar to those used for arrays of Schroeder and curved diffusers.

2.3.2. Curved hybrid surfaces

Specular reflections are attenuated, but not eliminated when using flat hybrid surfaces. Curvature can be incorporated into the design to break up specular reflections, resulting in a surface with similar absorption to the BAD panel, but greatly improved dispersion.

Curved hybrid surfaces are welcome guests in listening rooms and recording studio control rooms, as they enable the sweet spot to be spatially expanded [1]. In applications where moderate absorption is acceptable, they may compete with other diffuser designs: for example, a curved hybrid surface can be designed with one quarter the depth of a rigid optimized curved surface, and achieve dispersion of almost the same quality [1].

3. THE LEAN DIFFUSER OPTIMIZATION PROBLEM

This work focusses on solving an established problem—diffuser optimization—in a new practical context: “The lean diffuser optimization problem” places material and time constraints on the researcher and requires that the diffuser design method is simple. The bulk of the work was to implement a lean optimization framework, and with it solve the problem:

What modular diffuser provides an ‘optimal’ trade-off between uniform scattering and compact geometry, and how can such a surface be designed without using boundary element predictions?

3.1. Design objectives

Specifically, a successful diffuser design must satisfy the following criteria:

1. It must be shallower than existing profiled diffusers, with comparable or better broadband diffusion performance.
2. It should be modular, and designed to disperse sound optimally when modules are arranged an array.
3. It must be economical to manufacture and distribute.
4. The geometry must not pose obvious hazards to human safety.
5. The design method should utilize high quality scattering predictions, but it should not require a custom boundary element implementation¹.
6. Optimization must produce a suitable design within a reasonable solution time.
7. The design should be suitable for application on the back wall of a small-to-medium sized recording studio control room, where the purpose is to disperse first reflections from a hypothetical sound wave arriving at zero degrees incidence².
8. The design method should also be applicable to solving a larger problem: optimizing a surface for uniform scattering of an incident wave from any direction.

¹ There was a tight time frame on this project, and limited tools were available.

² This is a simplifying assumption.

3.2. Feasible approaches to optimal diffuser design

Numerous design methods were covered in the literature review, but not all are viable or worthwhile. Table 1 summarizes the various methods of designing a diffuser using numerical optimization. For any of the non-hybrid diffuser design methods, the minimum requirements to design a diffuser optimized for uniform scattering include:

1. A validated model for scattering prediction, such as a BEM or FDTD simulation.
2. A figure of merit or error parameter.
3. An implementation of the optimization problem, typically using downhill simplex to perform the minimization.

For numeric theory Schroeder diffusers and hybrid diffusers, a discrete optimization problem can be formed to search for a theoretically optimal number sequence with good autocorrelation properties. While this is a logical approach for the design of hybrid diffusers with customized absorption, it may have little merit for Schroeder diffuser design because optimal number sequences such as the quadratic residue sequence are already available. Moreover, the above approach is based on an approximate theory: it assumes that a number sequence with a flat power spectrum will yield optimal scattering. The current work does not make that assumption; and will instead use a design-by-simulation approach.

Rigid optimized curved surfaces generally have the best performance potential, and hybrid surfaces are blessed with simplified construction and a neutral aesthetic. However, the diffusers that most naturally facilitate efficient scattering with minimal absorption are

- Fractal Schroeder and stepped diffusers.
- Phase optimized stepped diffusers.
- Tanh addition diffusers.

The above designs approaches have a common objective: to achieve optimal diffusion by simulating a rough surface based on many instances of a simple, repeatable unit (e.g., a step function). Because they are related, these design approaches might use a similar framework for simulation and optimization.

Table 1 Prospective Design Approaches

| Diffuser Design Method | Optimization Approach | Analysis Methods | Advantages | Disadvantages |
|---|---|---|--|---|
| Numeric Theory Schroeder Diffuser | Search for a sequence with minimum side lobe energy in its autocorrelation. | <i>BEM, Thin panel BEM, Kirchhoff, FDTD, FEA.</i> | Elegant design equations. Optimization can be used to improve the theoretical designs, particularly those with few wells per period. Fractal construction can increase bandwidth [21]. | Visual appearance may conflict with room aesthetics. |
| Phase Optimized Schroeder or Stepped Diffuser | Manipulate the well depth series to minimize the standard deviation of the polar response. Check progress with boundary element predictions. Gradient is typically not available. Minimization with N degrees of freedom (e.g. $N = 7, 36$). Typical techniques: <i>downhill simplex, genetic algorithm, quasi-Newton.</i> | <i>BEM, Kirchhoff, FDTD, FEA.</i> | Uniform scattering is optimized directly, resulting in designs with better dispersion than numeric theory diffusers. Potentially straightforward to simulate using FDTD. | Computational cost. Requires boundary element or FDTD predictions during optimization [10], [3]. |
| Rigid Optimized Curved Surface | Manipulate the surface design using either a) Fuzzy constraints (inelegant). b) Distortion applied to shape variables (superior) [1]. Use a standard minimization technique like <i>downhill simplex</i> , with many different starting locations [14]. | <i>BEM, Kirchhoff, FDTD, FEA.</i> | Simple geometry affords lower construction costs and lower absorption. Best diffusion performance potential (requires a modulation scheme, unless the surface is very large). Customizable aesthetics. | Much deeper than comparable performing curved hybrid surfaces. Requires boundary element predictions during optimization [10], [3]. |
| Fractal: Random Addition Diffuser (RAD) | Enable optimization by using a sum of <i>tanh</i> functions to reduce the number of parameters defining the surface. Alter shape coefficients using a standard routine like <i>downhill simplex</i> [8]. | <i>BEM, Kirchhoff, FDTD, FEA.</i> | For random incident sound, the fractal performs better than the arc of a circle [1]. Performance of an optimized fractal is comparable to an optimized curve [1]. | No improvement over optimized curved surfaces—fractals often have inferior dispersion. Requires boundary element predictions during optimization [10], [3]. |
| Hybrid Diffusers | | | | |
| Planar Hybrid Surface: 1-D or 2-D Binary Amplitude Diffuser (BAD) | Discrete optimization of an N length sequence. For $N < 20$, an exhaustive search is possible. For $20 < N < 48$, a <i>genetic algorithm</i> can be used. To create longer sequences, concatenate several optimized sequences with low mutual cross-correlation [15,18,1]. | <i>BEM</i> (2-D analysis). For 3-D polar balloons, must use <i>thin panel periodic BEM</i> [1]. <i>FDTD, FEA.</i> | Low cost. Shallow profile. Treatment is hidden. Optimization gives complete control over the reflectivity. Often highly absorptive up to about 2kHz, thus only mid-to-high frequency dispersion needs consideration [1]. | Optimization is slow, and not possible for long sequences. Optimal design is not straight forward: requires deriving a family of number sequences, concatenating sequences into arrays and modulating to minimize periodicity. Patented design process. |
| Improved 2-D BAD using variable-size | Further improve the variable square size panel proposed in [20]. For an N-length number sequence with M different square sizes to | <i>BEM, FDTD, FEA. Fourier analysis</i> was | Potentially better performance than existing BADs. | Optimization is slow, and not possible for long sequences. If N is fixed, adding more |

| | | | |
|-----------------------|---|--|---|
| square panels | choose from, there are N^M degrees of freedom. Using modulation techniques on a family of optimized number sequences, there is the potential to position many small variable sized panels. | used for 3-D polar plots in [20], but it is only valid for far field scattering. | square size options results in exponentially more degrees of freedom. If the number of square sizes is fixed while N is increased, the degrees of freedom grow geometrically. |
| Curved Hybrid Surface | Potential approaches: a) Optimize the curvature shape for low frequency diffusion; use a pre-optimized number sequence for the hybrid surface coverage. b) Use a simple curvature with excellent diffusion (e.g., semi-cylinder) to break up specular reflection; use custom-optimized number sequences to control mid-to-high frequency dispersion. c) Custom-optimize both shape and number sequence [15]. | <i>BEM</i> , <i>FDTD</i> , <i>FEA</i> . | Curve breaks up the specular reflection found in flat hybrid surfaces. Can be designed with one quarter the depth of a rigid optimized curved surface, and achieve dispersion of almost the same quality [1]. The use of custom optimizations greatly complicates the design process. Advanced design technique with little literature publically available. |

3.3. Practical pitfalls to previous methods of diffuser optimization

Previous approaches to optimal diffuser design do not address the need for high quality design using simple to implement tools. Designs based on theoretical optimal sequences are simple, but less transferable to the real world than designs found via physical modeling. A superior design approach should optimize for uniform scattering directly using the desired performance as a figure of merit. This way, form directly follows function.

Phase-optimized diffusers are an example of form following function. However, they can only be designed when one has access to (or resources to implement) an optimization framework that includes accurate scattering prediction. A BEM is the standard prediction method, but it is not simple or quick to implement. Finite difference time domain is attractive from an implementation perspective, but at the time of this writing it is computationally unfeasible and will result in enormous solution times when applied to the standard phase-optimization problem.

The standard diffuser optimization method, as utilized by Cox [11], is to perform a high-resolution search of the solution space to find an optimal set of well depths, typically using downhill simplex. Cox defined solutions using 200 possible well depths with a resolution of 1 mm. If this problem is set up on a simple FDTD domain, with no special considerations—it is computationally ridiculous. This provided partial motivation for the lean optimization problem.

3.4. Design method

Agile prototyping was the strategy: iterative design, frequent testing. Rapid feedback was needed, thus the lean optimization framework arose out of practical necessity. It started as an FDTD simulation and evolved into a design system. It uses high resolution scattering predictions (FDTD) to evaluate candidates, but searches for solutions over a low resolution domain using an intrinsically parallel search method (the integer genetic algorithm). The resulting solutions are convenient, low-natural number sequences (e.g., zero to 16).

The design process utilizes trial and error feedback where appropriate. Most of this feedback is handled by numerical optimization; however, the designer must ensure that the system is working by interpreting visualizations of the available data. Additionally, the finite difference time domain simulation provides intuitive, real-time visual feedback.

The stepped diffuser: a simple geometric template

Stepped diffusers are simpler in form, and perform better than traditional Schroeder diffusers [1,11]. While traditional Schroeder diffusers are visually controversial, finless, optimized stepped diffusers may have more universal appeal because of a well-known principle: *beauty in design stems from purity of function*. If the *form follows function* principle holds true, a wall that intuitively looks like it is designed to scatter sound will have greater public acceptance than a strange looking wall without obvious function. Lastly, the design principle *Occam's razor* can be used to establish bias: of two diffusers with equal performance, the simpler design is generally the better choice—and easier to simulate.

A simple geometric template like the stepped diffuser can be placed in a modulated array, then optimized to design large diffusive surfaces based on smaller units. Additionally, design based on a common shape is potentially iterative, for example:

- Based on test results, optimization constraints can be quickly implemented by the designer, making the optimization process interactive.
- A stepped diffuser can be designed for a specific bandwidth, and then potentially updated to a fractal formation with expanded bandwidth. This technique combines optimization with Schroeder diffuser design theory. Optimization can be carried out on a single level of the fractal—or, if the scattering simulation is efficient, the entire fractal geometry can be simulated while the optimization algorithm manipulates the base shape of the fractal.
- A 1D stepped diffuser can be optimized, then updated to a 2D stepped diffuser using the Chinese remainder theorem.

A finite difference time domain prediction model: practical to implement

For stepped surfaces, a finite difference time domain (FDTD) simulation is practical to code from scratch and facilitates high quality predictions during the optimization process.

However, FDTD is computationally intensive compared with a BEM. The choice of FDTD is justified in Chapter 4.

An evolutionary algorithm: intrinsically parallel

Rather than searching for a single ‘optimal diffusing surface’, this work is concerned with generating designs that offer an excellent trade-off between uniform scattering and compact form factor. The design approach benefits from an intrinsically parallel optimization algorithm that enables the solution space to be thoroughly explored. It follows that an evolutionary algorithm is preferred over downhill simplex. Specifically, an integer genetic algorithm is the logical choice (justified in Chapter 7) due to practical constraints that arise when using FDTD simulation.

Iterative design-by-simulation: interactive and flexible

The design approach can be described as iterative design-by-simulation applied to an array of geometric templates. The optimization is constrained such that solutions will be stepped diffusers (or fractals based on stepped diffusers). However, the overall design approach is interactive and relies on feedback from a human designer. Design proceeds as follows:

1. Develop and test a finite difference time domain simulation framework.
2. Implement a parameter that measures the quality of diffusion achieved by a simulated surface. This is the objective function to be optimized.
3. Define preliminary constraints for the diffuser geometry based on a) the simulation framework and b) the design equations for Schroeder diffusers.
4. Develop, test and tune the optimization framework.
5. Perform optimization on modules in an array.
6. ←Feedback. Revise the design framework using an iterative process of cyclic prototyping. E.g., change the number of variables, number of modules, array modulation, well width, constraints on the well depths, etc.
7. Generate fractal forms based on the stepped diffusers with the best performance. The main goal here is to extend the bandwidth over which diffusion occurs.
8. Simulate and assess the results.
9. ←Feedback and iterative design.

4. PREDICTION OF SCATTERING USING FINITE DIFFERENCE TIME DOMAIN

For optimization based on uniform scattering to work, a model that predicts the scattering had to be developed and refined. Because the model needed to communicate with the optimization routine, it was not viable to rely on standalone acoustical simulation software for predictions. Instead, the prediction model was implemented as a Matlab function that can be called during optimization. The two processes are linked to form a design-by-simulation framework.

The typical modern approach to scattering prediction is to employ a boundary element model (BEM); however, given the time constraints on this work it was not feasible to develop a BEM from scratch. A finite difference time domain (FDTD) simulation has been implemented instead, offering practical benefits:

- FDTD is straightforward to code from scratch.
- FDTD gives accurate results [22], particularly near the specular reflection angle.
- Time domain simulation gives real-time visual feedback. At a glance the animated FDTD mesh can reveal problems such as numerical dispersion, and acoustical patterns such as grating lobes. This makes both designing surfaces and debugging code more intuitive.
- While frequency domain BEMs compute scattering at discrete frequencies, FDTD allows a wide frequency range to be excited by a single pulse, resulting in more data per prediction. Additionally, the temporal response is obtained directly by recording the scattered pressure at a desired location on the mesh. FDTD may be the most efficient way to gain reflected impulse responses, making it natural to evaluate the temporal dispersion³.
- Having only recently become a feasible method of evaluating acoustic scattering, FDTD has seen limited use compared with BEMs. This trend is expected to change because FDTD can simulate a largely unexplored class of problems: time variant systems.

³ The interpretation of temporal dispersion data is not well established; therefore, this work evaluates diffuser performance in terms of spatial dispersion. The study of temporal dispersion using FDTD is a future research opportunity.

The main weakness of FDTD is that the entire simulation domain must be meshed, and the space between mesh nodes limits the geometry of the diffuser and maximum frequency that can be represented. As the size or resolution of the domain increases, the size of the mesh increases, and the solution time increases by orders of magnitude. For a square domain of containing dim^2 mesh nodes, it was found that doubling the number of grid points increased the solution time by a factor of about 20.

If the goal is to simulate the reflected pressure in the far field, the near field to far field transformation (NFFFT) combined with absorbing boundary conditions can be used to reduce the computation time of FDTD. This work does not consider far field scattering therefore the NFFFT was not employed; however, an absorbing boundary condition—in this case perfectly matched layers (PMLs)—was crucial. Without PML boundaries a huge FDTD grid would have been required, and simulation times would inflate by orders of magnitude.

For direct comparison with the BEM predictions of Cox and D’Antonio [11,14,1], diffusers would need to be optimized on a FDTD mesh slightly larger than 5 m x 10 m, with a mesh spacing of 1 mm. This was not computationally feasible. Instead, the domain size and resolution were varied depending on the application (see Table 2, Table 3 and Table 7). The resulting FDTD model can be used to accurately assess scattering—but only for surfaces that can be represented as a rigid wall on the finite difference mesh.

4.1. A finite difference scheme with a Ricker wavelet source

This section outlines the two-dimensional FDTD scheme used to simulate acoustic scattering. The Matlab implementation is given in Appendix A.




Figure 19 Staggered pressure and particle velocity vectors in a FDTD mesh. (After Cox and D’Antonio [1])

FDTD is a numerical method that sees widespread use in computational electrodynamics. The method has been adapted to acoustics based on the conservation of momentum and continuity equations. FDTD models these using central-difference approximations to the space and time partial derivatives.

Central differences minimize the significance of higher order terms in the discretized partial derivatives, making the computation less prone to numerical error and therefore more robust. The sound pressure p and particle velocity u are formulated on staggered spatial grids in a Yee lattice (also known as a leapfrog scheme), resulting in an explicit time stepping scheme that avoids the need to solve simultaneous equations.

Figure 19 illustrates how the pressure vector component in the Yee lattice is located midway between a pair of velocity vector components. The mesh for the particle velocity in the x plane is shifted by $\Delta x/2$ with respect to the pressure mesh, and the mesh for the y component of the velocity is shifted by $\Delta y/2$. Additionally, the particle velocity time meshes are shifted by $\Delta t/2$ with respect to the pressure mesh.

This work utilizes a 2D FDTD scheme as the groundwork for scattering prediction. The central finite difference approximations to the pressure and particle velocity are implemented using the following update formulations [1]:

$$p_{i,j}^{n+\frac{1}{2}} = p_{i,j}^{n-\frac{1}{2}} - \frac{1}{K_e} \Delta t \left(\frac{ux_{i+\frac{1}{2},j}^n - ux_{i-\frac{1}{2},j}^n}{\Delta x} + \frac{uy_{i,j+\frac{1}{2}}^n - uy_{i,j-\frac{1}{2}}^n}{\Delta y} \right) \quad (4.1)$$

$$ux_{i+\frac{1}{2},j}^{n+1} = ux_{i+\frac{1}{2},j}^n - \frac{\Delta t}{\rho_0} \left(\frac{p_{i+1,j}^{n+\frac{1}{2}} - p_{i,j}^{n+\frac{1}{2}}}{\Delta x} \right) \quad (4.2)$$

$$uy_{i,j+\frac{1}{2}}^{n+1} = uy_{i,j+\frac{1}{2}}^n - \frac{\Delta t}{\rho_0} \left(\frac{p_{i,j+1}^{n+\frac{1}{2}} - p_{i,j}^{n+\frac{1}{2}}}{\Delta y} \right) \quad (4.3)$$

where superscripts denote the time index and subscripts denote the spatial indices, Δx and Δy are the mesh spacing in the x and y directions, Δt is the time step between computations, and $K_e = 1/\rho_0 c^2$ is the bulk modulus of the medium having density ρ_0 .

4.1.1. Numerical stability

Numerical stability requires that the time step be small enough to describe acoustic wave propagation. This is enforced using the Courant criteria, where the Courant number, s , is defined in 2D as

$$s = c\Delta t \sqrt{\left(\frac{1}{\Delta x}\right)^2 + \left(\frac{1}{\Delta y}\right)^2} \leq 1 \quad (4.4)$$

Additionally, on a mesh with discretization step size Δx , the maximum frequency that can be accurately simulated is restricted by $f \leq c/10\Delta x$ [1]. The highest frequency of interest for predicting room acoustics is typically about 5,500 Hz (the upper end of the 5 kHz one-third octave band), calling for a mesh spacing of 6mm. In this work, mesh spacing was chosen to be 10 mm for low frequency design and 2.85 mm for high frequency design⁴, therefore the maximum frequencies that can be simulated are 3440 Hz and 12,070 Hz, respectively. For these two discretization sizes, the Courant condition (Eq. (4.4)) required that the sampling frequency be no less than 48.7 kHz and 171 kHz, respectively.

4.1.2. Representing a diffusing surface

The diffusing surface was modeled as a perfectly rigid wall occupying desired coordinates on the mesh. At those points the particle velocity was set to zero for the duration of the simulation. Another approach is to use an impedance boundary condition that relates the pressure to the particle velocity normal to the surface. An acoustic impedance function can simulate geometries that do not snap directly to the mesh nodes; however, this may be problematic for complex surfaces as the time-domain acoustic impedance may not be a causal function [23].

4.1.3. Excitation

Three methods of pulsed excitation were tested on the FDTD domain: Gaussian pulses, Ricker wavelets and disk displacement initial conditions. In preliminary testing, the Ricker wavelet (also known as a Mexican hat wavelet) was found to offer the widest bandwidth

⁴ Mesh spacing was chosen based on diffuser geometry, discussed in Chapter 6.

without causing visible numerical dispersion artefacts on the mesh. Ricker wavelets are particularly useful because they do not introduce significant frequency components near DC. Gaussian pulses are acceptable; however, they contain the greatest energy at DC which can potentially introduce non-physical artefacts. For example, large-amplitude frequency components near DC might be resonant with the mesh.

The Ricker wavelet is equivalent to the second derivative of the Gaussian function and has spectral content fixed by a single parameter, the central frequency f_{cent} . The basic form of a Ricker wavelet is

$$W(t) = -\sqrt{2} \pi f_{cent} \left[\left(\sqrt{2} \pi f_{cent} t \right)^2 - 1 \right] \exp \left[-\frac{1}{2} \left(\sqrt{2} \pi f_{cent} t \right)^2 \right] \quad (4.5)$$

This wavelet is not zero for $t < 0$, therefore a transient can be expected at $t = 0$. To reduce the transient caused by activating the wavelet a temporal delay, t_0 , was applied. As t_0 can be any desired amount it was set using trial and error, but for convenience it can be expressed as an integer multiple of $1/f_{cent}$. This results in a time shifted wavelet like those shown in Figure 20. The wavelet was implemented in Matlab by applying the following update to the pressure mesh at a point (x_{cor}, y_{cor}) , while incrementing the time step n :

```
p(xcor,ycor) = -sqrt(2)*pi*fcent*((sqrt(2)*pi*fcent*(n-t0)).^2 - 1) ...
               *exp(-0.5*(sqrt(2)*pi*fcent*(n-t0)).^2);
```




Figure 20 Ricker wavelets and their spectra centered on 250 Hz and 2,000 Hz. The wavelets have been time shifted to reduce transients at $t = 0$ (after Redondo *et al.* [22]).




Figure 21 2D FDTD simulation of an $N=17$ stepped diffuser excited by a Ricker wavelet.

4.1.4. Receiver Arc

To obtain a polar representation of the scattered pressure, the response over the pressure mesh is recorded at many points on a semicircle. For a desired receiver radius r , the arc is approximated so that lattice points on the semicircle correspond to nodes on the FDTD mesh. This allows the array of sensors to be efficiently implemented, but the approximation must be done carefully so that sensors are evenly distributed. Figure 22 illustrates that with an arc radius of 250 grid steps, 180 sensors can be placed with nearly equal spacing. Higher accuracy is achieved when the radius is increased to 500 grid steps. The highest resolution simulations were those performed on fractalized diffusers, in which sensors were arranged on an arc with a radius of 1938 grid steps (corresponding to 5 m). The Matlab function *arclattice.m* was created to arrange sensors on an arc and is included in Appendix A.

The pressure is read at each sensor for the duration of the simulation, yielding a collection of temporal response vectors (Figure 23). This data is then processed using the methods in Chapter 5 to reveal the quality of diffusion.




Figure 22 Simulation domain used for preliminary testing with 180 sensors over 180° . Left: a single period of an $N=7$ diffuser. Right: An array of $N=7$ asymmetric diffusers in the aperiodic modulation $[1\ 0\ 1\ 1\ 0]$. The excitation source is depicted as a red point.




Figure 23 Near field temporal response to a Ricker wavelet source at normal incidence. Sensors are distributed on a $\pm 90^\circ$ receiver arc, such that the incident pulse is first received by the central sensor (at 0° from the normal), and the reflected response is first received by the outside sensors (located at $\pm 90^\circ$). The surface is an $N = 17$ stepped diffuser.

4.1.5. Anechoic boundary conditions

An anechoic environment was needed to enable scattering predictions with reasonable solution times. Without absorbing boundary conditions, propagating waves reflect off the boundaries of the mesh and interfere with ongoing data collection. Elimination of these reflected waves called for the use of perfectly matched layers (PMLs), which are regions of lossy medium that pad the interior of the boundaries.

A PML applies attenuation gradually to minimize the significance of discretization errors. The attenuation factor is set to zero inside the integration area and increases near the boundaries according to the expression

$$\gamma_x = \gamma_{x \max} \left| \frac{x - x_0}{x_{\max} - x_0} \right|^n \quad (4.6)$$

When PMLs are necessary the elegance of FDTD is compromised. Consequently it was not viable to continue building a robust FDTD prediction model from scratch, because the focus of this work was to design diffusers using lean optimization.

4.2. A simulation environment built with k-Wave for Matlab

The complications of PML prompted the search for an existing Matlab implementation, and an excellent solution was found: k-Wave, a free Matlab toolbox for the time-domain simulation of acoustic wave fields [24]. k-Wave places a level of abstraction between the numerical implementation (a k-space pseudospectral method with a PML) and the simulated physical environment. Using k-Wave as the scaffolding for simulation, an environment has been built to evaluate acoustic scattering and visualize the process. This prediction model can be run as a function during optimization (in which case it will display an animation and return a single-valued broadband diffusion parameter), or as a script when detailed analysis needs to be performed. Appendix B contains the Matlab code.

4.2.1. Adapting k-Wave to simulate scattering

Adapting k-Wave to the scattering problem required special considerations. To define a diffusing surface, a mask was created over part of the simulation domain with different material properties than air. k-Wave has a function to set the properties of a propagation medium, designed to be used for tissue realistic acoustics, ultrasound and photoacoustics. Realistically, the density of a solid object like a diffuser is much higher than that of air, and because the particles are closer together, sound propagates through the diffuser faster than it does through air. If these properties are set literally, the Courant criteria (Eq. (4.4)) requires that the mesh spacing be tiny, and given the size of the domain (meters) the resulting simulation becomes computationally unfeasible. Fortunately, propagation of sound inside the diffuser is of little concern for this work. The density was simply specified to be that of air throughout the domain, and to make the diffuser mask reflective the speed of sound at those points was set arbitrarily close to zero (1×10^{-9} m/s). This is not the same as setting the particle velocity to zero, as was done in Section 4.1; however, on a macro scale the resulting scattering was indistinguishable.

The Ricker wavelet pulse was implemented in k-Wave by defining a custom time-varying source. Without full access to the k-Wave time series, specifying a suitable amplitude pulse required scaling the Ricker expression. It proved more reliable to produce a short duration pulsed excitation by setting a point displacement initial condition on the mesh. This was the excitation method of choice when optimizing diffusers.



Figure 24 Near field scattering from a randomly generated stepped diffuser. The excitation pulse was created using a point displacement initial condition.

4.2.2. Flattening the spectral envelope

At a given sensor, the sound pressure level (SPL) spectrum $L_p(f)$ can be obtained from the fast Fourier transform (FFT) of the pressure $p(t)$:

$$P(f) = FFT\{p(t)\}$$

$$L_p(f) = 20 \log_{10} \left[\frac{|P(f)|}{\sqrt{2}P_0} \right] \quad \text{dB SPL}$$

where $P_0 = 20 \cdot 10^{-6}$ Pa (RMS) is the threshold of human hearing. $p(t)$ is a time-varying quantity with positive and negative values, thus the pressure spectrum $|P(f)|$ is divided by $\sqrt{2}$ to get the quadratic mean (RMS value).

Regardless of the excitation method, the bandwidth of the pulse is limited by the mesh spacing and the spectrum is not flat. Consequently the temporal response (Figure 25) is not identical to the impulse response, and the frequency spectrum of the scattered pressure does not equal the frequency response. While optimization can function without knowing the true frequency response⁵, it is more intuitive to check that the model works when the visual representation of an even response matches the human expectation of ‘even looking’.

A rough correction of the response spectrum was applied to assist visualization. First, an inverse envelope was created from the SPL spectrum of a reference incident pulse, $L_{p,ref}(f)$. To isolate the spectral envelope from undesired local fluctuations, pseudo-Gaussian smoothing was applied to the reference pulse spectrum. This consisted of three passes of a moving average filter with a width of 10 FFT bins. The reference was also normalized to have a maximum magnitude of unity. The inverse envelope, $H(f)$, is simply the inverse of the smoothed normalized pulsed spectrum:

$$H(f) = \left(\frac{\text{smooth}\{L_{p,ref}(f)\}}{\max[\text{smooth}\{L_{p,ref}(f)\}]} \right)^{-1}$$

⁵ Diffusers are evaluated using the statistical criteria in Chapter 5.

For any sensor θ on the receiver arc, the SPL spectrum of the scattered response $L_p(q, f)$ can be approximately flattened through multiplication with the inverse envelope:

$$\bar{L}_p(q, f) = H(f) \times L_p(q, f)$$

This deconvolution yields an approximate frequency response suitable for visualization (Figure 26), but not for critical analysis. The deconvolution assumes that propagation losses are equal at all frequencies, which is false. Sound propagating in air experiences atmospheric effects that causes high frequencies to be absorbed more than low. The amount of absorption depends on the temperature and humidity, which the numerical scheme can model by adjusting the sound speed and density in the medium. Consequently the high frequencies are undercompensated by the inverse envelope. A more suitable reference spectrum might be created from a pulse that has propagated the same distance as the reflected wave front, so that each will have experienced similar attenuation upon reaching the sensor. Another concern is that the smoothed reference pulse does not perfectly align with its original spectral envelope—and when the original is multiplied with the inverse envelope, a small amount of amplitude distortion appears in the flattened spectrum. A less noise-prone method of flattening the spectrum is Wiener deconvolution, which finds an optimal compromise between inverse filtering and noise smoothing.




Figure 25 Temporal response and spectrum of a diffuser measured at the central sensor.

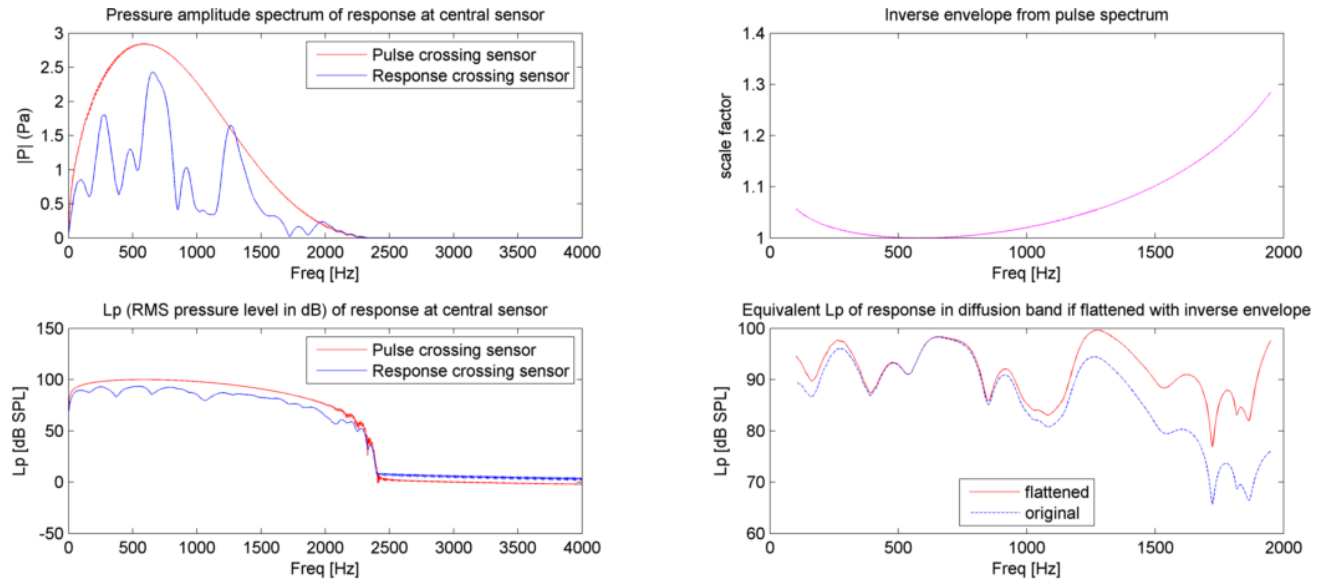


Figure 26 Left: Reference pulse spectrum and response spectrum.
Top right: inverse envelope. Bottom right: response spectrum flattened by inverse envelope.

5. MEASURES OF DIFFUSION QUALITY

5.1. A diffusion parameter based on the standard error

For optimization to be viable the scattering pressure distribution needed to be reduced into a single merit of diffusion. Cox developed and tested several diffusion parameters based on the standard errors for all frequencies, source locations and receiver arcs [11,14]. Minor anomalies arose when the standard error was calculated via intensities, therefore a revised diffusion parameter evaluates the standard error using sound pressure levels: [14]

$$\varepsilon(\mathbf{r}_s, r, f) = 2\sqrt{\frac{\sum_{\theta=-85}^{85} [L_p(r, \theta) - \bar{L}_p(r, \theta)]^2}{n(n-1)}} \quad \text{dB} \quad (5.1)$$
$$\bar{L}_p(r, \theta) = \log \left[\frac{1}{n} \sum_{\theta=-85}^{85} 10^{L_p(r, \theta)} \right]$$

where n is the number of receiver angles in 180° and $\bar{L}_p(r, \theta)$ is the mean value calculated via intensities. $\varepsilon(\mathbf{r}_s, r, f)$ is the standard error for each receiver arc radius r , source position \mathbf{r}_s and frequency f . This formula assumes that the sound pressure level is viewed and heard on a roughly linear perceptual scale. According to equal loudness contours, the assumption is most valid for high sound pressure levels (>80 dB) between 300 Hz and 1000 Hz.

To ensure that the standard error is dominated by the general envelope rather than local fluctuations, spatial averaging was applied to the response prior to computing Eq. (5.1). This was accomplished using a sliding average filter with a width equal to the number of sensors⁶ in 10° . An alternative approach is to smooth the frequency spectrum (e.g., by averaging in one-third-octave bands). When using a time domain model the FFT can be readily computed from the temporal response, therefore smoothing the spectrum is a viable alternative to spatial averaging.

⁶ For a receiver radius of 250 grid steps, a dense sensor mask was used with 39 sensors per 10° . For a receiver radius of 500 grid steps it was possible to create a sparse sensor mask with sensors spaced by roughly 1° .

The mean diffusion $\bar{\varepsilon}$ can be obtained by averaging the standard error for all frequencies, source locations and receiver arcs. However, there may be cases where poor diffusion at a particular frequency or angle is compensated for by excellent diffusion at another frequency or angle. To penalize these cases, a standard error of the standard errors is added to the mean. The result is a single-valued broad-band diffusion parameter, ε' :

$$\varepsilon' = \bar{\varepsilon} + \frac{2}{\sqrt{m(m-1)}} \sqrt{\sum_{j=1}^m [w_j(\mathbf{r}_s, r, f) \varepsilon_j(\mathbf{r}_s, r, f) - \bar{\varepsilon}]^2} \quad (5.2)$$

$$\bar{\varepsilon} = \frac{1}{m} \sum_{j=1}^m w_j(\mathbf{r}_s, r, f) \varepsilon_j(\mathbf{r}_s, r, f)$$

where m is the total number of frequencies, incident angle and receiver arcs, each having a standard error $\varepsilon_j(\mathbf{r}_s, r, f)$ and a weighting function $w_j(\mathbf{r}_s, r, f)$ [14]. The weighting functions may be used to assign relative importance to certain source positions, frequency ranges or receiver arcs.

In this investigation, optimization was performed for a single receiver radius and source position, with the weighting functions set to unity. Therefore, Eq. (5.2) can be simplified and represented in terms of n frequencies [11]:

$$\varepsilon' = \bar{\varepsilon} + \frac{2}{\sqrt{n(n-1)}} \sqrt{\sum_{i=1}^n [\varepsilon_i(f) - \bar{\varepsilon}]^2} \quad (5.3)$$

$$\bar{\varepsilon} = \frac{1}{n} \sum_{j=1}^n \varepsilon_i(f)$$

5.2. The autocorrelation diffusion coefficient

There are various statistical approaches to calculating a diffusion coefficient; however, because these methods involve data reduction they are prone to inaccuracies [1]. The autocorrelation diffuser coefficient has proven to be the most reliable method to compare diffusion between devices that were tested under different conditions.

While autocorrelation is typically used to reveal self-similarity between a signal and a delayed version of itself, it is used here to evaluate the spatial similarity of the scattered energy. High values in the spatial autocorrelation function indicate that a surface scatters sound uniformly to all receivers; low values indicate that a surface is concentrating reflected energy in one direction [1]. The autocorrelation diffusion coefficient for a fixed source position was implemented as

$$d_{\psi} = \frac{\left(\sum_{\phi=-90}^{90} 10^{L_p(\phi)/10} \right)^2 - \sum_{\phi=-90}^{90} (10^{L_p(\phi)/10})^2}{(n-1) \sum_{\phi=-90}^{90} (10^{L_p(\phi)/10})^2} \quad (5.4)$$

where L_i are the set of sound pressure levels arriving at each receiver (sensor) in decibels, n is the number of receivers arranged on a semicircle, and ψ is the angle of incidence. The above equation requires an equal angular spacing between each receiver. This was approximated by using a radius of 250+ grid steps, with each receiver location rounded to a point on the grid that approximately intersects a semicircle.

Eq. (5.4) was evaluated in one-third octave bandwidths which results in a smoothed coefficient characterized by the spectral envelope⁷.

At low frequencies the surface acts like a point source due to the increased prominence of edge scattering, resulting in a diffusion coefficient that increase as the frequency approaches zero. As this behavior can cause confusion, a normalized diffusion coefficient can be used to remove it. The normalized diffusion coefficient was calculated by first applying Eq. (5.4) to the test surface, yielding d_{ψ} , and to a reference flat surface of equal size, yielding $d_{\psi,r}$. From d_{ψ} and $d_{\psi,r}$ the normalized diffusion coefficient was computed as

$$d_{y,n} = \frac{d_y - d_{y,r}}{1 - d_{y,r}} \quad (5.5)$$

⁷ In Chapter 9 the diffusion coefficient spectra is shown for several surfaces.

6. THE PRELIMINARY DESIGN OF STEPPED DIFFUSERS

Meaningful constraints were needed to help formulate the optimization problem. Given a desired specification—in this case the operational bandwidth and the number of wells—the design equations for Schroeder diffusers were used to determine the viable range of well depths and widths.

Application of this design theory requires that plane wave propagation (propagation in one dimension) dominates in the wells. Based on this assumption, the upper frequency of diffusion is conveniently related to the well width:

$$w = \frac{\lambda_{\min}}{2} = \frac{c}{2f_{\max}} \quad (6.1)$$

where w is the well width, λ_{\min} is the shortest wavelength to avoid cross-modes in the wells, and f_{\max} is the upper limit of dispersion based on this basic design theory. This limit is used only as a guideline for preliminary design, because in practice diffusion will occur above f_{\max} [1]. This is especially true for stepped diffusers, as the plane wave propagation theory is greatly compromised by removing the fins between wells.

The quadratic residue sequence was used to estimate a practical range of depths based on the desired design frequency. First, the sequence numbers for a QRD were calculated using

$$s_n = n^2 \text{ modulo } N \quad (6.2)$$

where s_n is the sequence number for the n^{th} well and N is a prime number. For an $N = 7$ QRD, $s_n = \{0, 1, 4, 2, 2, 4, 1\}$. The objective was to design an optimized stepped diffuser, not a QRD; therefore, only the maximum and minimum sequence numbers were used. Well depths for a Schroeder diffuser are determined from the sequence numbers using

$$d_n = \frac{s_n \lambda_0}{2N} \quad (6.3)$$

where λ_0 is the design wavelength. It follows that a Schroeder diffuser has well depth that vary between 0 and $\lambda_0/2$. These limits were exploited as constraints for the optimization problem. While the maximum depth d_{\max} is flexible, the depth of the diffuser was limited to

an accepted standard for the sake of aesthetics and production feasibility. Additionally, this depth constraint speeds up the optimization.

The design frequency, f_0 , was taken to be the lower limit of controlled diffusion for a Schroeder diffuser:

$$f_0 = \frac{s_{\max}}{N} \frac{c}{2d_{\max}} \quad (6.4)$$

Quadratic residue diffusers are designed to have optimum diffusion at integer multiples of f_0 ; however, the bandwidth is limited because the QRD behaves like a plane reflector at frequencies mNf_0 where $m = 1, 2, 3 \dots$ [1]. At these *critical frequencies*, all well depths are integer multiples of half the wavelength, thus all wells re-radiate in phase. It follows that to design a QRD, w and N must be chosen so that f_{\max} is less than the first critical frequency, f_0 . This condition can be written as a constraint on N :

$$N > \frac{c}{2wf_0} \quad (6.5)$$

6.1. Limitations and considerations for phase grating surfaces

A practical implication of Eq. (6.5) is that for a given design frequency, the well width can only be reduced if N is increased—but doing so increases the complexity of the optimization problem and the geometry. Eq. (6.5) also implies that for a given design frequency there is a minimum repeat width for one period of the diffuser (Nw). If made too narrow, f_0 will be governed by the period width instead of by d_{\max} .

The design theory relies on grating lobes generated by periodicity in a surface, thus a Schroeder diffuser must be periodic to exhibit ‘optimal’ dispersion [1]. Conversely, too many periods cause narrow grating lobes with large nulls in between, resulting in uneven scattering. The best design based purely on Schroeder’s theory will consist of an array of about five periods [1]. An alternative is to use modulation schemes to form an array, which is a general technique that is not restricted to Schroeder diffusers.

If the diffuser wells are too narrow the viscous boundary layer becomes significant, causing the absorption of the device to increase. The minimum practical well width to avoid this is about 2.5 cm, with 5 cm being a more typical choice [1]. For narrow-welled designs, it is expected that absorption will be less severe for the typical stepped diffuser. Stepped

diffusers lack dividers between wells, and thus present less bounding surface to a plane wave at normal incidence.

The low frequency efficiency depends on the ratio of s_{\max}/N . For Schroeder diffusers, the choice of prime number is limited by low frequency performance, critical frequencies and manufacturing considerations. For stepped diffuser optimization, the choice of N is also governed by computational constraints. In this case N corresponds to the number of wells and is not required to be a prime number.

6.2. Geometric framework for stepped diffusers

Preliminary geometries to optimize were formulated using Eq. (6.1) - (6.5). Geometries are named according to their scale:

- *Module A-LF* refers to a single period of a structure to optimize with a low design frequency (for these purposes, $f_0 < 800$ Hz is considered LF).
- *Array A-LF* refers to multiple periods of Module A-LF in a modulated array.
- *Module A-HF* refers to the second stage of a fractal based on Module A-LF.
- A first order structure is any module or array with a LF design frequency.
- A fractal refers to a nested structure with both a LF and HF design frequency.

The choices were narrowed as additional practical constraints were introduced:

1. Each discrete unit length or width must be an integer multiple of the simulation grid step size; otherwise, the geometry will be rounded to the nearest point on the grid. Physical dimensions were chosen carefully so that the first generation module and a second generation fractal may be placed on a common grid and simulated within a reasonable time frame.
2. The number of wells N must be sufficiently small to allow fractal formations in a small form factor. This will enable a simple, compact diffuser with a wide bandwidth.

Table 2 Preliminary Calculations for Phase Grating Diffusers⁸

| Identifier for single period of geometry | No. wells | Well width (mm) | Max depth (mm) | Round to grid step (mm) | Design freq. (Hz) | Max freq. (Hz) | Module width (mm) | Total depth (mm) | Total bandwidth (Hz) |
|--|-----------|-----------------|----------------|-------------------------|-------------------|----------------|-------------------|------------------|----------------------|
| Module A-LF <i>Base Struct.</i> | 7 | 60 | 160 | 10 | 614 | 2867 | 420 | 205.6 | ~19450 |
| Module A-HF <i>2nd Stg. of Fractal</i> | 7 | 8.57 | 45.6 | 2.85 | 2155 | 20067 | 60 | | |
| Module B-LF <i>Base Struct.</i> | 7 | 60 | 140 | 10 | 702 | 2867 | 420 | 180 | ~19370 |
| Module B-HF <i>2nd Stg. of Fractal</i> | 7 | 8.57 | 40 | 2.85 | 2457 | 20067 | 60 | | |
| Module C-LF <i>Base Struct.</i> | 7 | 70 | 200 | 10 | 491 | 2457 | 490 | 240 | ~16700 |
| Module C-HF <i>2nd Stg. of Fractal</i> | 7 | 10 | 40 | 2 | 2457 | 17200 | 70 | | |
| Module B-LF <i>Base Struct.</i> | 7 | 140 | 140 | 10 | 702 | 1229 | 980 | 220 | ~7850 |
| Module B-HF <i>2nd Stg. of Fractal</i> | 7 | 20 | 70 | 5 | 1229 | 8600 | 140 | | |

⁸ This work focussed on the framework shown in bold. The other frameworks were used for preliminary testing.

7. OPTIMIZATION

Optimization was performed using five modules in an array with aperiodic modulation. As there is a trade-off between the pitfalls of periodicity (nulls in the polar response) and the benefits (symmetrical scattering), the optimization problem has been set up to determine whether the best array is periodic with a symmetrical base shape or aperiodic with an asymmetrical base shape. The aperiodic modulation $\{1\ 0\ 1\ 1\ 0\}$ was used where ‘1’ denotes the base shape and ‘0’ denotes a flipped version. For symmetrical modules the array naturally becomes periodic with modulation $\{1\ 1\ 1\ 1\ 1\}$. If the best base shape is known to be symmetrical, the optimization can be greatly sped up by exploiting symmetry to reduce the number of variables.

With only 16 discrete well depths to choose from the optimization problem called for mixed integer programming, therefore downhill simplex and gradient descent methods were not suitable. The mixed integer programming problem is non-deterministic polynomial-time hard (NP-hard), which implies that there is no efficient algorithm known to solve it. In this work the objective function is the (nonlinear) scattering prediction model, while the solution space is the set of all possible well depth combinations. $N = 7$ yields a solution space of $16^7 = 268,435,456$ candidate solutions; consequently, converging to a solution within reasonable time required a minimization algorithm that could be tuned to this particular nonlinear, mixed integer problem. The integer genetic algorithm (GA) was the natural choice.

7.1. Genetic algorithms

Genetic algorithms are based on natural selection. They modify a population of ‘individuals’ (solutions) through multiple generations. To create each new generation the algorithm randomly selects individuals from the current population to breed, producing the children for the next generation. Given multiple generations the population evolves toward an optimal solution. Figure 27 describes the optimization processes for a genetic algorithm adapted to mixed integer programming.

Three operators are used to create the next generation from the current population at each step:

- *Selection* chooses which parents will survive and contribute to the next generation.
- *Crossover* represents mating; it derives children from parents.
- *Mutation* applies random modifications to individual parents to create children.

Genetic algorithms are intrinsically parallel: they have multiple offspring that can search the feasible region (design space) in multiple directions simultaneously, such that a feasible region can be explored with a reasonable chance of finding the global optimum. If a dead end is found, genetic algorithms can eliminate it and continue exploring. In contrast, most minimization algorithms are serial and only traverse the design space in one direction at a time, stopping when they find an apparent solution or hit a dead end.

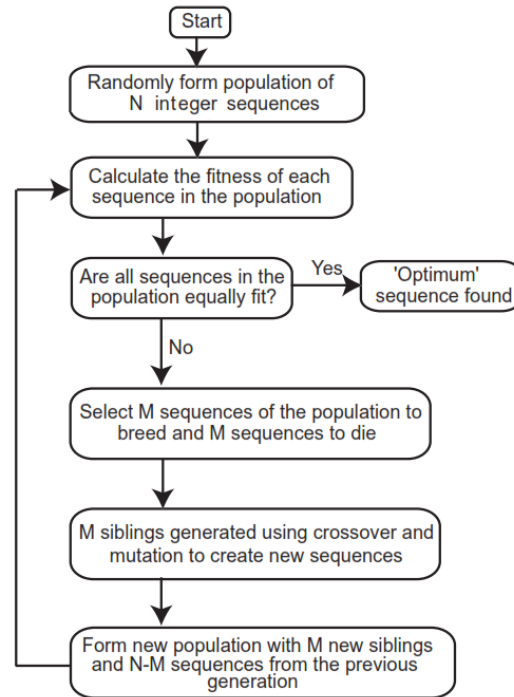


Figure 27 A genetic algorithms for integer sequence optimization.
(Modified after Cox and D'Antonio [1])

7.2. Optimization using an integer genetic algorithm

Optimization was performed across a range of frequencies that roughly span the 10dB-bandwidth of the excitation pulse. This will be called the *diffusion band*. 12 discrete frequencies were chosen for optimization, distributed quasi-randomly in the diffusion band. The response spectrum was characterized by a 4096 point FFT, and the optimization frequencies were the FFT bins centered on {102, 250, 449, 574, 700, 824, 949, 1102, 1250, 1450, 1700 and 1950 Hz}. This sparse set of FFT bins was chosen in an effort to clearly define the optimization objective, aid in visualizing progress and perhaps speed up

convergence⁹. When the goal is to optimize for twelve specific frequencies (rather than all frequencies), the disparity between good solutions and average solutions is higher and results can be compared more meaningfully. Moreover, Cox found that optimization at seven frequencies [11] was sufficient to achieve good dispersion over the entire bandwidth.

To emphasize the importance of the lower-mid frequencies, the range 400-1250 Hz received the highest density of optimization frequencies. Another option is to optimize across all FFT bins in the diffusion band, assigning relative weights to each bin when solving Eq. (5.2).

7.2.1. Tuning and convergence

Tuning genetic algorithms to find an appropriate population size and mutation rate was time consuming. On a 512 x 512 simulation mesh, each scattering prediction took about 2 minutes on an i7 2600K desktop processor. A typical optimization run took about 40 hours during which the scattering was evaluated 1200 times with the goal of minimizing the single-valued broadband diffusion parameter ε' . Given that the solution space contained over 260 million candidates, the algorithm was considered to be well tuned if it made progress toward an apparent local minimum while maintaining a diverse population.

When the population size was too small (e.g., 20) the genetic pool would quickly become dominated by a few individuals and the algorithm would converge prematurely to a local minima—typically resulting in a poor quality ‘solution’. With a large population size the solution space was explored more thoroughly; however, progress toward a minima was slower. A suitable balance was found by setting the population size to 40, resulting in 30 generations per 1200 function evaluations (Figure 28). For the integer GA, a thorough search of the solution space would require a larger population size and number of generations, say, 100 and 300 respectively. This is possible, but the optimization would take over a month with the given set up.

Test runs provided insight on what additional constraints might assist convergence. In early testing all well depths were constrained to integer multiples of the mesh spacing, where s_{max} denotes the maximum integer value (typically 16). The two outside well depths were later constrained to integers between $s_{max}/2$ and s_{max} . This prevents the algorithm from searching undesired forms that were expected to exhibit particularly poor scattering

⁹ It was hypothesized that convergence might be fastest for a problem with narrow criteria; however, this was not confirmed.

and a deep profile. One intention was to avoid concave forms which have a focussing rather than a dispersing effect.

Table 3 Optimization Frameworks

| Framework name | No. wells | Well width (mm) | No. depths | Minimization algorithm parameters | Simulation parameters | Sensors | Frequencies optimized |
|--------------------------|-------------------------|-----------------|------------|---|---|--|----------------------------------|
| Preliminary Tuning Setup | 7+ (solo module) | 60 | 20 | Meta-heuristic GA, DE, PSO ¹⁰ Population: 20+ | Grid xDim: 256 Grid zDim: 256 Grid step: 10 mm Source: 2.2 m, 0° | Radius: 1.25 m Angle: $\pm 85^\circ$ Quantity: 335 | All FFT bins between 100-1950 Hz |
| Design Setup Op250M | 7 (solo module) | 60 | 14 | Integer GA Population: 40 Tol Function: 1e-3 | Grid xDim: 512 Grid zDim: 512 Grid step: 10 mm Source: 4.5 m, 0° | Radius: 2.5 m Angle: $\pm 85^\circ$ Quantity: 667 | 12 FFT bins between 100-1950 Hz |
| Design Setup Op250 | 7 (x5 in array [10110]) | 60 | 16 | Integer GA Population: 40 Tol Function: 1e-4 | Grid xDim: 512 Grid zDim: 512 Grid step: 10 mm Source: 4.5 m, 0° | Radius: 2.5 m Angle: $\pm 85^\circ$ Quantity: 667 | 12 FFT bins between 100-1950 Hz |

7.3. Optimization results

Optimization results were collected over three weeks, which involved numerous runs of the integer GA using design frameworks Op250M and Op250 (frameworks are outlined in Table 3). The default optimization run relied on about 1200 scattering predictions; however, early runs using framework Op250M were limited to 800 scattering predictions. Diffusers were optimized at a receiver radius of 2.5 m using a pulsed excitation at 4.5 m and normal incidence. The best performing shapes were also tested using a full-scale scattering simulation (see Chapter 9) to ensure excellent scattering at a radius of 5 m with a source at 10 m ¹¹. Due to the large number of runs, only the best results are presented in Table 4.

¹⁰ Before utilizing the integer GA, tuning was performed on a *meta-heuristic global optimizer* that combines the GA, differential evolution (DE) and particle swarm optimization (PSO) to increase the likelihood of finding the global optimum [34]. This optimizer, when modified for use with integers, had comparable performance to the integer genetic algorithm. While the meta-heuristic optimizer is expected to outperform the GA for problems with larger populations [34], the integer GA was preferred for this work due to its relative simplicity.

¹¹ In the literature, diffusers performance is typically assessed at a receiver radius of 5 m with a source at 10 m [1].

The best result found using array-based optimization is the symmetrical base shape named A1-LF, having well depths of {13 9 8 10 8 9 13} cm. Figure 28 shows the progress toward this optima, where the goal was to minimize the penalty value ϵ' . The mean penalty for the population did not converge to the best penalty value, indicating that the integer GA maintained a diverse population throughout the 30 generations. While true convergence was not achieved, Figure 28 shows significant minimization progress followed by a plateau in both the best value and mean value of ϵ' . This was the preferred method of finding a solution using only 1200 scattering evaluations. A better solution might be found by letting the optimization process run much longer until it converges.

Table 4 Optimization Results^{12,13}

| Result name | Framework used to obtain result | No. wells | Well width (cm) | Effective array | Depths and resulting heights ¹⁴ (cm) | Diffusion param. ϵ' | Significance of this result |
|-------------|---|-----------|-----------------|-----------------|---|------------------------------|--|
| A1-LF | Design Setup <i>Op250</i> (integer GA) | 7 | 6 | 1 1 1 1 1 | 13 9 8 10 8 9 13 → stepheights → 0 4 5 3 5 4 0 | 9.7553 (in array) | Best symmetrical base shape found via optimization in the array [1 0 1 1 0]. |
| A2-LF | Design Setup <i>Op250</i> (integer GA) | 7 | 6 | 1 0 1 1 0 | 13 12 7 4 6 11 12 → stepheights → 0 1 6 9 7 2 1 | 9.7707 (in array) | Best asymmetrical base shape found via optimization in the array [1 0 1 1 0]. |
| B1-LF | Design Setup <i>Op250M</i> (integer GA) | 7 | 6 | 1 (solo) | 10 6 3 4 3 5 10 → stepheights → 0 4 7 6 7 4 0 | 9.4861 (solo) | Of the top ten base shapes found via solo optimization (no array), this shape maintained the best performance when later placed in the array 1 0 1 1 0. |
| B2-LF | Analysis Setup ¹⁵ <i>Sim250</i> | 7 | 6 | 1 0 1 1 0 | 10 6 3 4 2 5 9 → stepheights → 0 4 7 6 8 5 1 | N/A ¹⁶ | The above shape was mutated into several asymmetric shapes via input from a human designer. Each candidate was tested in the array 1 0 1 1 0, and the winner was selected. |

¹² Depth sequences shown in bold represent the best designs, particularly when converted to fractal formations. These designs are fractalized in Chapter 8 and further analyzed in Chapter 9.

¹³ The diffusion parameter ϵ' depends on the scattering prediction model set up, and can only be used to compare results that have an identical simulation set up. Side-by-side test results and a universal comparison using the autocorrelation diffusion coefficient are given in Chapter 9.

¹⁴ To minimize the depth, heights are offset such that lowest step is flush with the diffuser base.

¹⁵ See Table 7 for details on the FDTD simulation framework used for full-scale analysis.

¹⁶ This result was obtained using the full scale FDTD simulation framework *Sim1* from Table 7, therefore the diffusion parameter is not directly comparable to the others in Table 4.

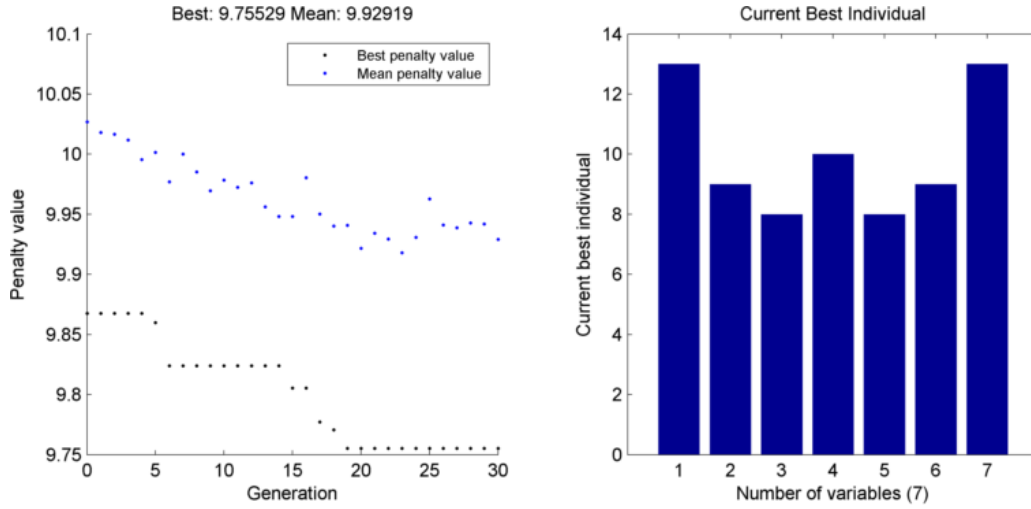


Figure 28 Minimization progress after 1200 scattering predictions for 5 modules in an array with modulation [1 0 1 1 0]. The goal was to minimize the penalty value, ε' .




Figure 29 Near field simulation of the optimized array with base shape A1-LF.

The solution A1-LF is of special significance because it has a much shallower form factor than previously optimized stepped diffusers. An $N=7$ stepped diffuser optimized by Cox had well depths spanning a 15 cm range [11]; A1-LF has an operational depth of 5 cm. The low value of ε' , computed using Eq. (5.3), can be visually understood by examining Figure 30. The standard error at a 2.5 m receiver radius for a source at 4.5 m $\angle 0^\circ$ is denoted as $\varepsilon(4.5\text{m}, 2.5\text{m}, f)$, or shorthand as $\varepsilon(f)$. Figure 30 shows that the standard error, when scaled with the inverse envelope, has a relatively flat spectrum. Additionally, the local fluctuations in $\varepsilon(f)$ are not severe, therefore the variance in $\varepsilon(f)$ is acceptably low. Moreover the mean value of $\varepsilon(f)$ is small compared with other diffusers. Since the mean and variance of $\varepsilon(f)$ are both small, the single-value broadband diffusion parameter ε' is small.

It follows that because ε' is small, the scattered response is relatively uniform throughout the diffusion band. Figure 31 confirms this; however, because the A1-LF array is periodic grating lobes contribute peaks and nulls to the scattered polar distribution. This effect is further analyzed in Chapter 9.




Figure 30 Standard error $\varepsilon(4.5\text{m}, 2.5\text{m}, f)$ for an array of optimized base shapes A1-LF. ε has been scaled by the inverse envelope, and is shown at each FFT bin (red line). Optimization frequencies are indicated with an 'o'.

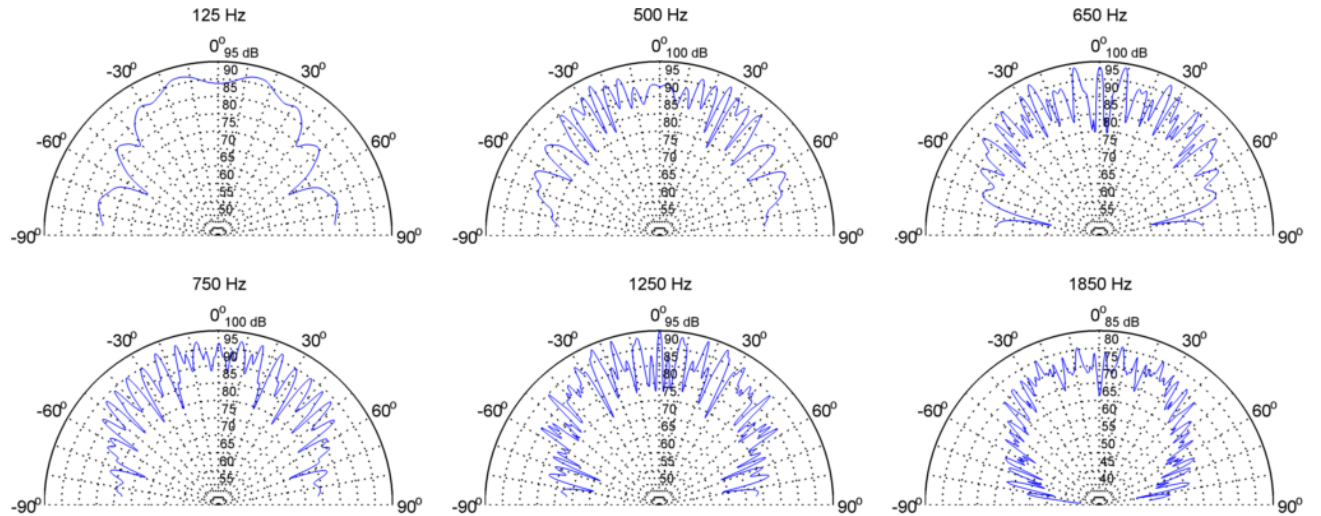


Figure 31 Scattered polar distribution for the optimized array with base shape A1-LF. The scattering is shown at one optimization frequency (1250 Hz) and five others. Results were obtained via k-Wave prediction with a 2.5 m receiver radius and a source at 4.5 m, 0°.

8. BANDWIDTH EXTENSION VIA FRACTALIZATION

Fractal formations have been used as an elegant way to extend the bandwidth of the optimized diffusers. This required a higher resolution simulation mesh, which in turn expanded the bandwidth of the excitation pulse (Figure 35). Based on the pulse spectrum, the diffusion band spans 100-5100 Hz. Diffusion throughout the band is assessed in Chapter 9.

The top performing base shapes have been converted to fractal form according to the specifications in Table 5. In addition, a fractal formation was created from a stepped diffuser that was optimized by Cox [11]. This diffuser from the literature will be named L95. L95 was not precisely reproduced on the limited-resolution mesh; therefore, quantitative performance comparisons with L95 are avoided. L95 has been useful as an intuitive reference when comparing animated time domain simulations. Here, it is simply a geometric reference.

Table 5 Framework for Designing Fractal Stepped Diffusers

| Type of geometry | No. wells | Well width (mm) | Round to grid step (mm) | Max freq. (Hz) | Module width (mm) |
|----------------------------------|-----------|-----------------|-------------------------|----------------|-------------------|
| Base Shape | 7 | 60 | 10 | 2867 | 420 |
| 2 nd Stage of Fractal | 7 | 8.57 | 2.85 | 20067 | 60 |

Table 6 Fractal Formations

| Fractal name | Base shape | Base shape depths and heights ¹⁷ (mm) | 2 nd stg step heights (mm) | HF dsgn freq. (Hz) ¹⁸ | Base shape depth (mm) ¹⁹ | Tot. depth (mm) |
|--------------|------------|---|---------------------------------------|----------------------------------|-------------------------------------|-----------------|
| A1-Frac | A1-LF | 130 90 80 100 80 90 130 → stepheights → 0 40 50 30 50 40 0 | 0 11.4 14.25 8.55 14.25 11.4 0 | 6897 | 60 | 74.25 |
| B2-Frac | B2-LF | 100 60 30 40 20 50 90 → stepheights → 0 40 70 60 80 50 10 | 0 11.4 19.95 17.1 22.8 14.25 2.85 | 4311 | 90 | 112.80 |
| L95-Frac | L95 [11] | 168 20 55 48 55 20 168 → stepheights → 0 148 113 120 113 148 0 | 0 42.75 31.35 34.2 31.35 42.75 0 | 2299 | 158 | 200.75 |

¹⁷ To minimize the depth, heights are offset such that lowest step is flush with the diffuser base.

¹⁸ The resulting ‘design frequency’ for the second stage of the fractal has been calculated using Eq. (6.4). This quantity was not used in design; it is simply a predictor of gaps in the diffusion spectra.

¹⁹ The base shape depth includes a 10 mm deep base to enable mounting on a wall or ceiling.



Figure 32 Three arrays of $N = 7$ base shapes with array modulation $\{1\ 0\ 1\ 1\ 0\}$. Base shapes A1-LF (top), B2-LD (middle) and L95 (bottom). L95 is a close approximation of a stepped diffuser optimized by Cox [11].




Figure 33 Three arrays of $N = 7$ second order fractals with array modulation $\{1\ 0\ 1\ 1\ 0\}$. Fractal formations A1-Frac (top), B2-Frac (middle) and L95-Frac (bottom).

The geometry of each fractal reveals whether or not the design is appropriate for practical use. L95 is clearly not suitable as a base shape for a fractal. It looks particularly jagged because the fractal framework with a 2.85 mm grid step (designed in Chapter 6) creates an elongated second stage. The framework has been designed to facilitate time domain simulation—not to preserve the height-to-width ratio of the base shape. The resulting formation, L95-Frac, would protrude 200 mm out of the wall, exposing long barbs that could

seriously injure a person upon collision. Additionally, the deep, narrow wells increase the significance of the viscous boundary layer, making the fractal subject to higher absorption at high frequencies. Conversely, it is the depth of these wells that gives L95-Frac a suitable HF design frequency (2299 Hz), “theoretically” suggesting a smooth transition between the LF and HF diffusion bands (Chapter 9 refutes this). The other fractal diffusers are safer, lower profile and easier to manufacture; however, as they have a high HF design frequency (Table 6) they may introduce gaps between the LF and HF diffusion bands.




Figure 34 Temporal response and spectrum of A1-Frac, measured at the central sensor.

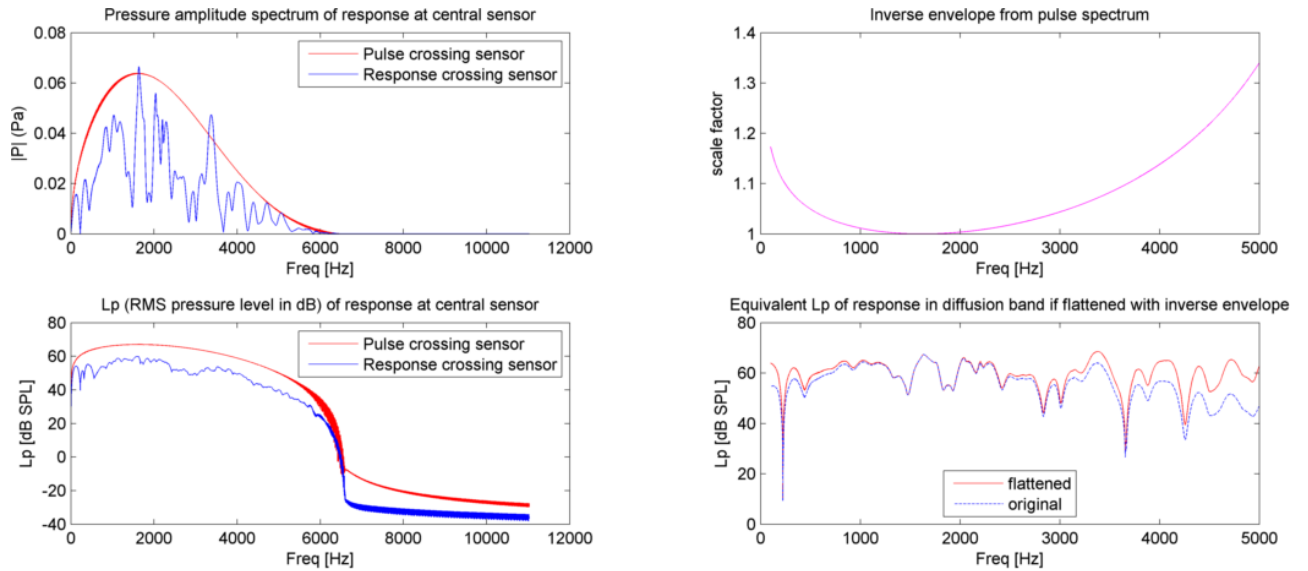


Figure 35 Left: Reference pulse spectrum and response spectrum for A1-Frac. Top right: inverse envelope. Bottom right: response spectrum flattened by inverse envelope.

9. SIMULATION AND ANALYSIS OF THE DESIGNS

A compact simulation domain with a receiver radius of 2.5 m was used during optimization. This was based on the assumption that the diffuser would typically be used in a small-to-medium sized recording studio control room. However, Eq. (5.1) (used to calculate γ) requires that the receiver radius is significantly larger than the diffuser width [14].

In the literature diffusion is typically measured at a receiver radius of 5 m with an excitation source at 10 m. This is the setup called for by the diffusion coefficient standard, AES-4id-2001 [25]. For real-world measurement, AES-4id recommends 36 microphones spaced 5° apart over a $\pm 90^\circ$ arc. For predictions using a BEM, Cox [14] used 54 sensors per receiver arc, distributed over $\pm 85^\circ$.

Table 7 Simulation Frameworks for Full-Scale Analysis

| Framework Name | FDTD mesh parameters | Excitation Pulse | Sensors ²⁰ | Frequencies tested | Solution time using i7 2600K processor |
|---|---|---|--|----------------------------------|--|
| <i>Sim500</i> based on AES-4id-2001 [25] | xDim: 1100 zDim: 1024 dx,dz: 10mm | Distance: 10 m Angle: 0° Magnitude: 1 | Radius: 5 m AngleA: $\pm 85^\circ$ NumSensorsA: 1337 AngleB: $\pm 90^\circ$ NumSensorsB: 180 | All FFT bins between 100-1950 Hz | 20 min |
| <i>SimF250</i> Fractal forms near field testing | xDim: 1796 zDim: 1796 dx,dz: 2.85mm | Distance: 4.5 m Angle: 0° Magnitude: 1 | Radius: 2.5 m AngleA: $\pm 85^\circ$ NumSensorsA: 2343 AngleB: $\pm 90^\circ$ NumSensorsB: 180 | All FFT bins between 100-5100 Hz | 2 hr 45 min |
| <i>SimF500</i> Fractal forms analysis based on AES-4id [25] | xDim: 4098 zDim: 4098 dx,dz: 2.85mm | Distance: 10 m Angle: 0° Magnitude: 1 | Radius: 5 m AngleA: $\pm 85^\circ$ NumSensorsA: 4685 AngleB: $\pm 90^\circ$ NumSensorsB: 180 | All FFT bins between 100-5100 Hz | 22 hr |

In this work a 10 m x 11 m domain has been used to analyze the performance of the most promising designs. The new simulation frameworks, Sim500 and SimF500 (Table 7), are based on AES-4id-2001 [25]. Sensors on the simulation mesh cost next to nothing, therefore at least 180 receivers were used for all simulations. However, a side effect of fitting the receiver arc to a rectangular grid is that as more sensors are added, they become less evenly spaced. To ensure uniform spacing, only 180 sensors were used when measuring the

²⁰ For generating polar plots, the sensor arc consisted of *NumSensorsA* receivers over the angular range *AngleA*. For computing the autocorrelation diffusion coefficient, *NumSensorsB* evenly spaced receivers were used over the angular range *AngleB*.

autocorrelation diffusion coefficient. In contrast, the polar plots were produced using the detailed scattered response captured from a dense array of sensors.




Figure 36 Full scale domain used to simulate the scattering from fractal diffusers.




Figure 37 Pressure received at each sensor during simulation of the A1-Frac array.

9.1. Feedback from the specular zone

Figure 38 shows the scattering from a plane reflector the same width as the diffuser array. When the receiver radius is 5 m, geometric reflection is received between $\pm 15^\circ$. This is the *specular zone*. As the receiver arc shrinks, the specular zone widens, and Eq. (5.1) becomes less valid as a means to measure the standard error, $\varepsilon(r_s, r, f)$. For large surfaces (e.g., a wall) the specular zone can be as much as $\pm 90^\circ$, and a different diffusion parameter is necessary [14]. AES-4id-2001 recommends that 80 percent of receivers are outside the specular zone [25,1].

A good polar response can be recognized by understanding that a diffuser should disperse energy from the specular zone to other positions. When compared to diffusers, the plane reflector has a poor polar response on the Sim500 domain. This suggests that Eq. (5.1) is valid at the 5 m receiver radius: its use can be justified because the $\pm 15^\circ$ specular zone is small compared to the $\pm 85^\circ$ receiver arc. Specifically, 82.4 percent of receivers are outside the specular zone, which complies with AES-4id-2001. Finally, this knowledge can be used as feedback to assess the design method: If diffusers tested at a 2.5 m receiver radius retain their relative performance ranking at the 5 m radius, then it supports the case that

1. Eq. (5.1) was valid during optimization.
2. Diffusion at a 2.5 m receiver radius may function as a rough predictor of performance at the 5 m radius.

The results in Table 4 and Table 8 are consistent with these arguments. Likewise, all other diffusers that were compared retained their relative performance ranking at the 2.5 m and 5 m radius. While the results support the above arguments, the sample size was too small to confirm a strong trend.




Figure 38 Polar response at 5 m from a reflector the same width as the diffuser array.

9.2. Performance results

Table 8 compares the performance of selected diffusers that were tested using a 5 m receiver radius. The best diffuser should have the smallest single-value broadband diffusion parameter, ϵ' . This is the same parameter that optimization sought to minimize at a 2.5 m receiver radius. The mean autocorrelation diffusion coefficient for a 0° incident source is denoted \overline{d}_0 . A good diffuser should have a larger value for \overline{d}_0 ; however, this is not a reliable metric. Unlike ϵ' , \overline{d}_0 pays no attention to the evenness of the diffusion spectrum (this will be shown shortly).

Table 8 Performance of Diffusers at a Receiver Radius of 5m²¹

| Diffuser | Simulation Framework | N | Array | Well depths based on the sequence | Diffusion param at 5m, ϵ' | Mean diffusion coeff, \overline{d}_0 | Significance of this diffuser |
|-------------|----------------------|----------------|------------------------|-----------------------------------|------------------------------------|--|--|
| A1-LF | Sim500 | 7 | 1 1 1 1 1 | 13 9 8 10 8 9 13 | <u>4.9786</u> | 0.4716 | Best symmetrical base shape found via optimization in the array [1 0 1 1 0]. |
| B2-LF | Sim500 | 7 | 1 0 1 1 0 | 10 6 3 4 2 5 9 | 4.9873 | <u>0.5107</u> | Best asymmetrical base shape found via optimization and input from a human designer. |
| L95 | Sim500 | 7 | 1 1 1 1 1 | 17 2 6 5 6 2 17 | 5.1029 | 0.4039 | Loosely based on a stepped diffuser optimized by Cox [11]. ²² |
| Stepped QRD | Sim500 | 7 | 1 1 1 1 1 1 0 1 1 0 | 0 4 16 8 8 16 4 | 5.0970 5.2207 | 0.4810 0.4466 | QRD diffusers typically have fins between the wells. This one is finless. |
| A1-Frac | SimF500 | 7 ² | 1 1 1 1 1 | 13 9 8 10 8 9 13 | 2.3238 | 0.5728 | Fractal form of A1-LF. |
| B2-Frac | SimF500 | 7 ² | 1 0 1 1 0 | 10 6 3 4 2 5 9 | <u>2.3038</u> | <u>0.6331</u> | Fractal form of B2-LF. |
| L95-Frac | SimF500 | 7 ² | 1 1 1 1 1 | 16.8 2 5.5 4.8 5.5 2 16.8 | 2.4171 | 0.5016 | 1 st stage of this fractal closely approximates a stepped diffuser optimized by Cox [11]. |

²¹ Good diffusion is expected when ϵ' is minimized and \overline{d}_0 is maximized.

²² L95 is not identical to the stepped diffuser optimized by Cox [10], thus it should be viewed as a generic diffusive surface, not as an optimized diffuser. Likewise, the stepped QRD results do not represent the performance of a typical QRD (QRDs are typically designed with fins).

Included in Table 8 is a stepped diffuser based on the $N = 7$ quadratic residue sequence. There are no dividers between the wells of the stepped QRD, thus it does not recreate the performance of the equivalent finned QRD. The scattered polar distribution for this diffuser in both periodic and aperiodic arrays is featured in Appendix D.

9.3. Analysis of optimized designs at a receiver radius of 5 m

The diffusion parameter for the A1-LF array is reasonably uniform throughout the diffusion band (Figure 39), particularly above 1000 Hz. The plot of $\varepsilon(10\text{m}, 5\text{m}, f)$ dips between 200-1100 Hz. This is where the best diffusion is achieved. This range also contains the highest density of optimization frequencies, thus the achievements on the Sim500 domain appear to be correlated with optimization on the Op250 domain. Regardless, the equivalent plot for Op250 (Figure 30) showed a more uniform diffusion parameter, demonstrating that optimal design at one radius does not equal optimal design at another radius.

Based on the broadband diffusion parameter ε' , the A1-LF array was the top performer at both the 2.5 m and 5 m radius (Table 8). A1-LF is also the most compact design, having a total depth of just 6 cm. One pitfall of the design is that that periodicity lobes cause an uneven polar response at certain frequencies. For example, Figure 41 depicts significant peaks and nulls in the scattered polar distribution at 750 Hz. This was an expected side effect of placing modules in a periodic array—and to circumvent this, B2-LF was designed using aperiodic modulation. Another shortcoming of the shallow A1-LF design is that it is expected to exhibit less temporal dispersion than a deeper diffuser. This is not a major concern because uniform spatial dispersion produces temporal variation as a side effect [1].

The B2-LF array has slightly inferior diffusion when compared with A1-LF. Figure 40 reveals that $\varepsilon(f)$ has significant variance, therefore ε' will accumulate a significant penalty²³. Spectral variance in $\varepsilon(f)$ exists because the standard error of the polar response varies with frequency. At each frequency, the polar response is effected by geometry, and it follows that the asymmetry in B2-LF has contributed to a less-than-uniform scattered distribution. This is depicted in Figure 43 and the supplemental plots in Appendix D. Even though the aperiodic modulation thwarts the effects of periodicity, the resulting design is inferior to A1-LF. B2-LF has worse performance, more complicated geometry and a deeper form factor. However, at 9 cm deep B2-LF remains shallower than most commercial diffusers.

²³ Recall that ε' depends on both the mean value and standard error of $\varepsilon(f)$.




Figure 39 Diffusion parameter $\varepsilon(10\text{m}, 5\text{m}, f)$ for the diffuser array A1-LF.




Figure 40 Diffusion parameter $\varepsilon(10\text{m}, 5\text{m}, f)$ for the diffuser array B2-LF.




Figure 41 Scattering from the A1-LF periodic array at a receiver radius of 5 m. The local peaks and nulls at 750 Hz are caused by periodicity.

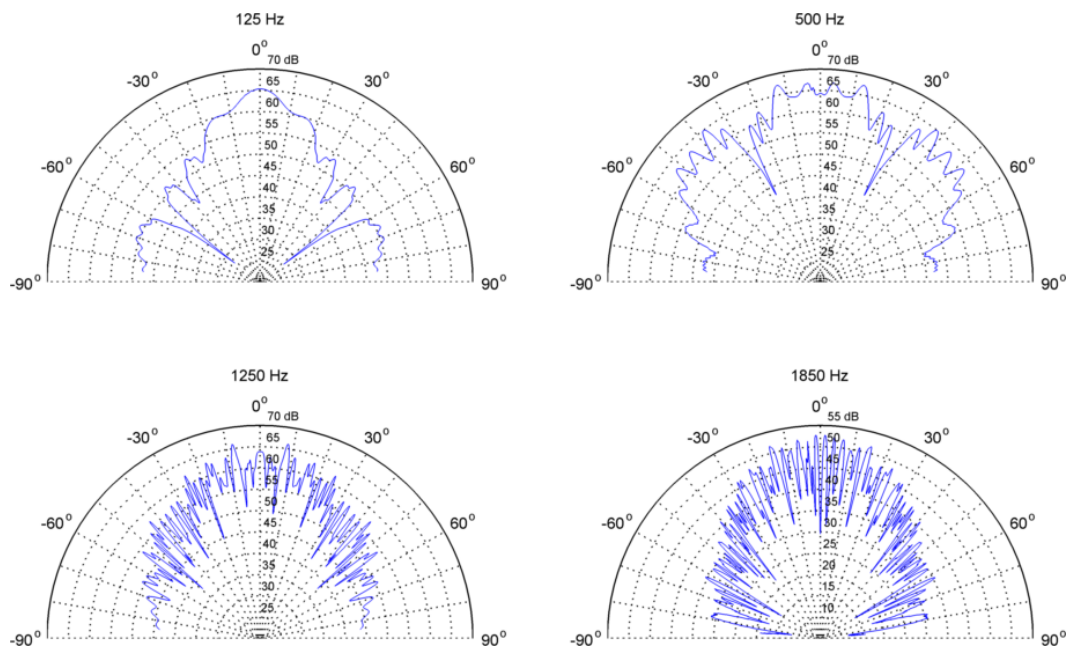


Figure 42 Scattered polar distribution for the optimized diffuser array A1-LF.

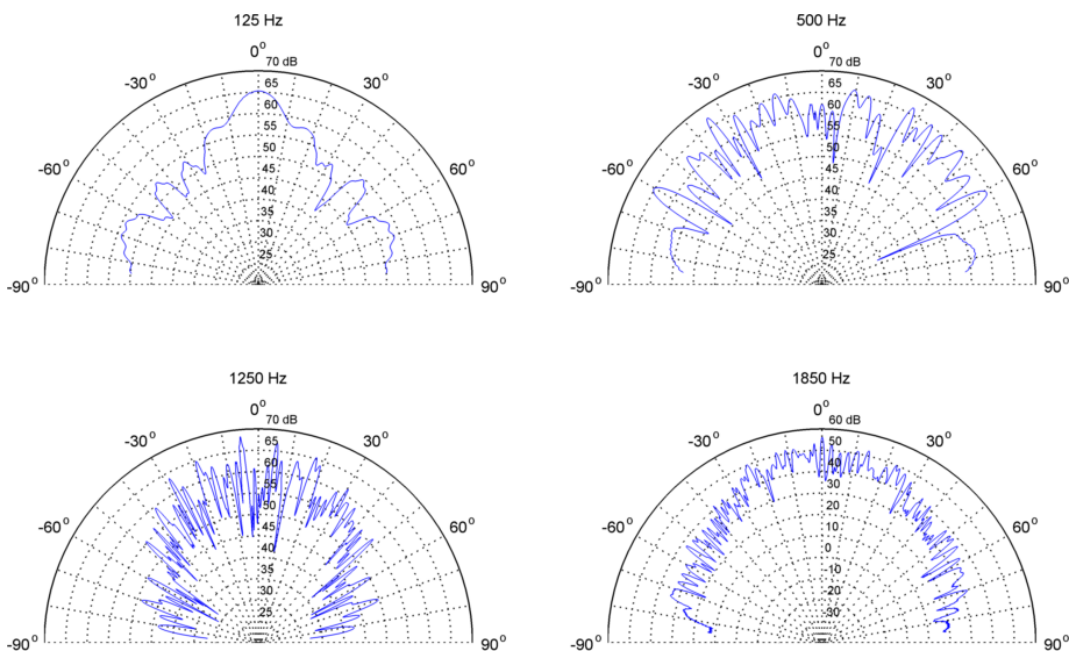


Figure 43 Scattered polar distribution for the optimized diffuser array B2-LF.

Figure 44 and Figure 45 show the autocorrelation diffusion coefficient spectra for the two optimized diffusers. The diffusion coefficient for a 0° incident source, d_0 , was calculated using Eq. (5.4). Eq. (5.5) was used to obtain the normalized diffusion coefficient, d_{n0} . A reference for the normalization was created by simulating a plane reflector the same width as the diffuser.

Research Update: The following coefficient plots were updated in 2012

In the images below (produced using *AFMG Reflex Software*), the diffusion coefficient has been averaged in one-third octave bands and plotted as a function of frequency. High values of the diffusion coefficient indicate that a surface scatters sound uniformly to all receivers; low values tell us that a surface is concentrating reflected energy in one direction. Detailed performance reports for these diffusers — and other optimized diffusers — are available at <http://argen.com/sound-diffusers/>.




Figure 44 Diffusion coefficient spectra for A1-LF. Diffuse field (left). 0° incidence (right).




Figure 45 Diffusion coefficient spectra for B2-LF. Diffuse field (left). 0° incidence (right).

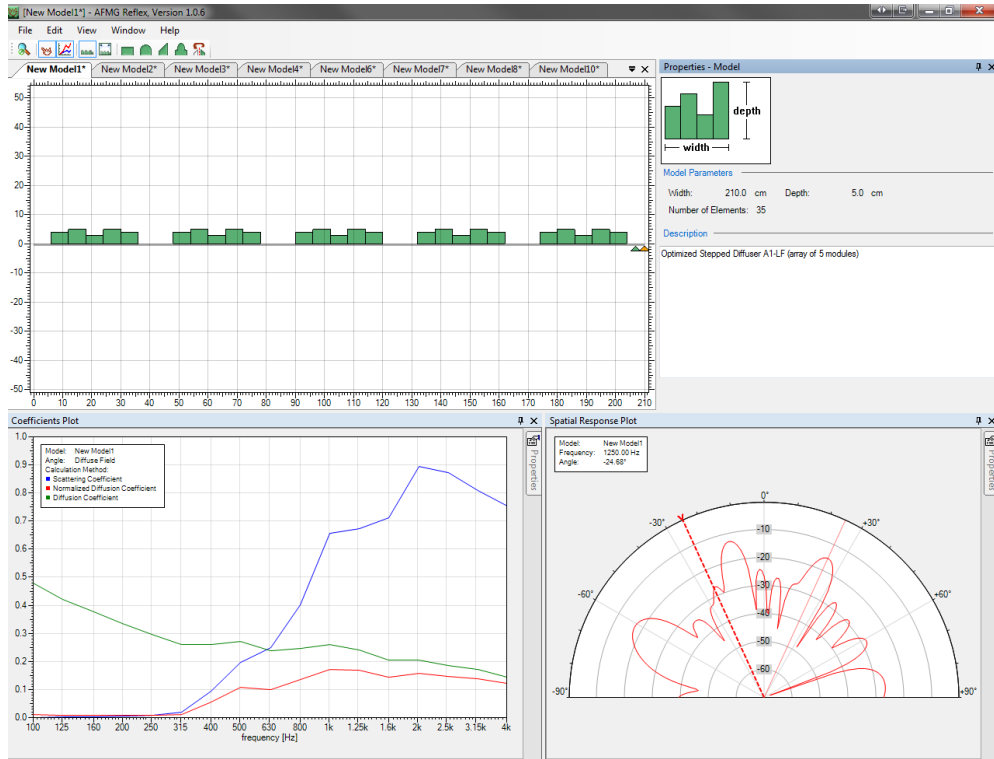


Figure 46 Diffuse field performance report for A1-LF array. Produced in AFMG Reflex.

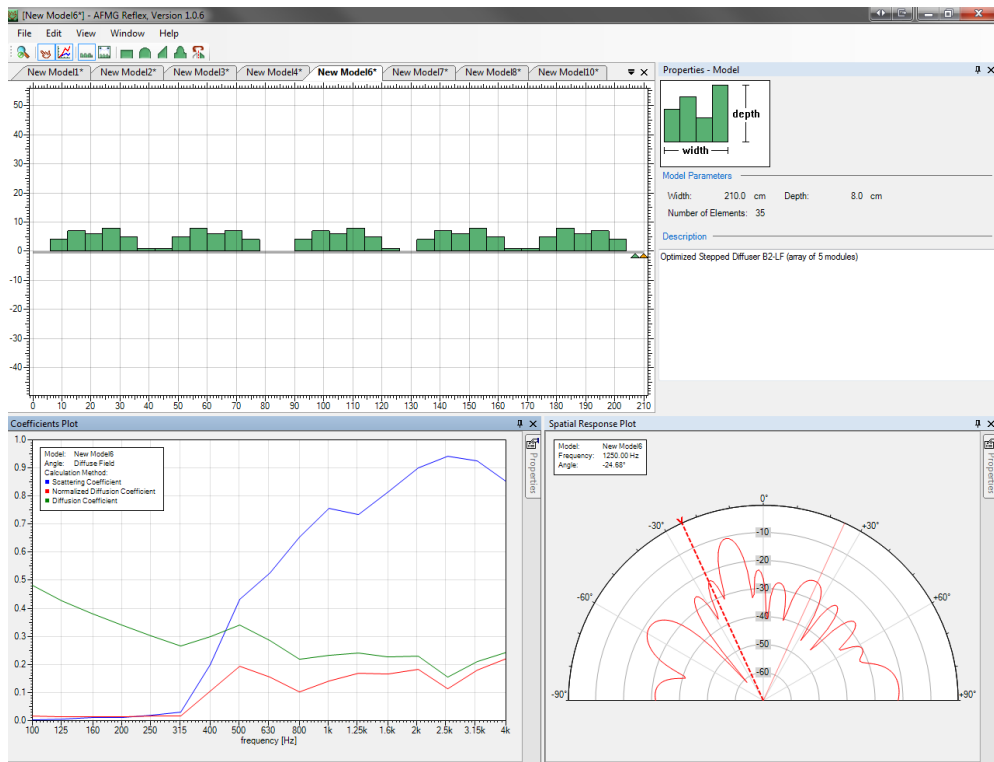


Figure 47 Diffuse field performance report for B2-LF array. Produced in AFMG Reflex.

9.3.1. Applying profiled modulations to the optimized base shape

Because the A1-LF module is symmetrical, 5 modules placed side-by-side form a periodic array. A great way to increase the diffusion performance is to reduce this periodicity by mounting the modules at different depths. This will be called a *profiled modulation*.

Profiled Modulation 1: Mounting 5 Modules to Prevent Periodicity

The image below shows an example where 5 A1-LF modules have been mounted at specific depths. Using BEM simulations in AFMG Reflex, these profile depths were determined by trial and error to be [0 cm, 5 cm, 6 cm, 5 cm, 0 cm].




Figure 48 Profiled Modulation 1 coefficient spectra. Diffuse field (left). 0° incidence (right).

Profiled Modulation 2: Mounting 7 Modules to Prevent Periodicity (Fractal Modulation)

Profiled Modulation 2 uses seven A1-LF (Leanfuser™) modules mounted at different depths based on fractal self-symmetry. The seven mounting depths are derived from the depth sequence of the A1-LF module. The resulting diffuser has basic fractal geometry; therefore, this will be called a *fractal modulation*. This modulation could also be applied to the A1-frac module (the fractal version of A1-LF), resulting in a 3rd order nested fractal.

The mounting sequence [0, 8, 10, 6, 10, 8, 0] cm is based on the depth sequence of the A1-LF module multiplied by two: $2[0, 4, 5, 3, 5, 4, 0]$ cm = [0, 8, 10, 6, 10, 8, 0] cm.

The scaling factor was determined experimentally by simulating various profiled modulations using AMFG Reflex. A scaling factor of two was chosen because it creates a shallow profiled diffuser that offers good performance with minimal depth.




Figure 49 Profiled Modulation 2 coefficient spectra. Diffuse field (left). 0° incidence (right).

Deep Fractal Modulation:

If a thick diffuser was desired, one could take the fractal self-symmetry concept more literally. For example, a low frequency fractal stage can be designed with the proportions equal to the stepped diffuser proportions (I.e., keeping the same width:depth ratio). For the A1-LF module, the width:depth ratio = 42 cm / 5 cm = 8.4. Therefore, for an array of 7 diffuser modules (294 cm wide), the deepest step of the profiled modulation would be 294 cm / 8.4 = 35 cm. Since the deepest step of the A1-LF module is 5 cm, one would scale all the depths by 35 cm / 5 cm = 7.

This would give a profiled modulation of 7x[0, 4, 5, 3, 5, 4, 0] cm = [0, 28, 35, 21, 35, 28, 0] cm. The resulting diffuser would have an operational depth of 40 cm, which is much deeper than the other designs mentioned here. The expected benefit of this deeper diffuser would be a lower cut-in frequency (extending the diffusion bandwidth into the low end).

9.4. Analysis of fractal designs at a receiver radius of 5 m

Fractalization has resulted in high frequency performance gains. The results show that while A1-LF has superior performance as a base shape, B2-Frac is the better fractal sound diffuser. B2-Frac dominates Table 8 with a broadband diffusion parameter of 2.3038 and a mean autocorrelation diffusion coefficient of 0.6331. The diffusion coefficient can only be compared between designs that were simulated on the same simulation domain, in this case SimF500. Given that a different simulation framework was used in Section 9.3, verifying the performance gains will require simulating all diffusers on the SimF500 framework. This is not a quick task, as it takes about 22 hours to compute a single SimF500 prediction on an i7 2600K processor.




Figure 50 Simulation of near field scattering from fractal diffusers A1-Frac.




Figure 51 Simulation of near field scattering from fractal diffusers B2-Frac.

Figure 52 shows that for A1-Frac and B2-Frac, the diffusion parameter spectral envelope is reasonably²⁴ uniform between 200-5000 Hz (ignoring local fluctuations). The diffusers achieve an ε' of 2.3238 and 2.3038 respectively, which is substantially lower than the ε' of 2.4171 achieved by L95-Frac.




Figure 52 Diffusion parameter $\varepsilon(10m, 5m, f)$ for fractals based on optimized diffusers. Top: A1-Frac periodic array. Bottom: B2-Frac aperiodic array with modulation $\{1\ 0\ 1\ 1\ 0\}$.

²⁴ Qualitative observations such as “reasonably uniform” are based on comparisons to other surfaces. In this chapter, it means noticeably better than the equivalent result of L95-Frac.

The scattered polar distributions give a more intuitive indication of the scattering quality; however, this is not a reliable way to critically assess diffusers. Here is why: the scattered pressure level from A1-Frac looks highly uniform for all three frequencies shown in Figure 53. In Figure 55 the frequency 1250 Hz was swapped for the 750 Hz and the dB range of the polar plot was changed. The resulting plot tells a completely different story about this diffuser. The response at 750 Hz much less uniform than at 1250 Hz, and the rescaling of the plot reveals two narrow nulls in the polar distribution at 500 Hz.

Figure 55 shows that A1-Frac exhibits the same periodicity lobes as A1-LF at 750 Hz. This confirms that the mid-frequency characteristics of the original optimized form, A1-LF, have been preserved in the fractal form. B2-Frac, however, has undergone a more notable change during fractalization—and it shows in the polar response. Figure 55 clearly illustrates the superior uniform dispersion achieved by B2-Frac below 800 Hz. Figure 56 shows that B2-Frac produces an asymmetrical response that tends to approximate scattering. The asymmetry is fundamentally due to the asymmetrical base shape. Fortunately, the aperiodic modulation {1 0 1 1 0} has a secondary purpose: it creates symmetrical forms with four of the modules in the B2-Frac array (Figure 51). Thanks to this sequence, asymmetry in the polar response is mild, and periodicity is avoided.

When they become available, the autocorrelation diffusion coefficients for these fractal diffusers will be posted to <http://argen.com/sound-diffusers/>.

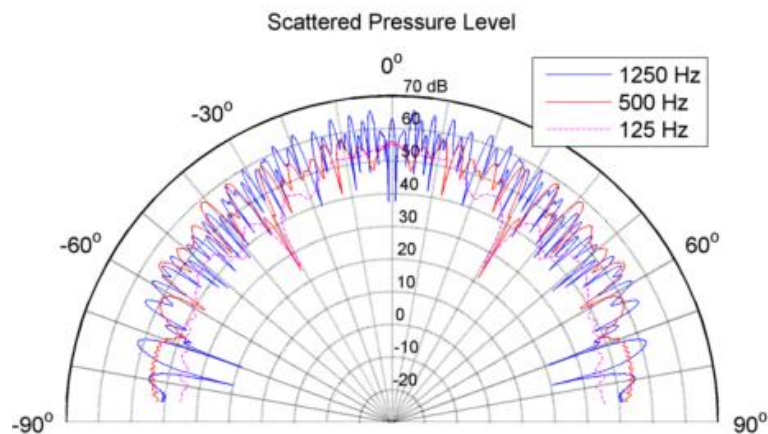


Figure 53 Scattering from the A1-Frac array. Fluctuations are under-exaggerated because the 0 dB contour is not at the center of the plot.

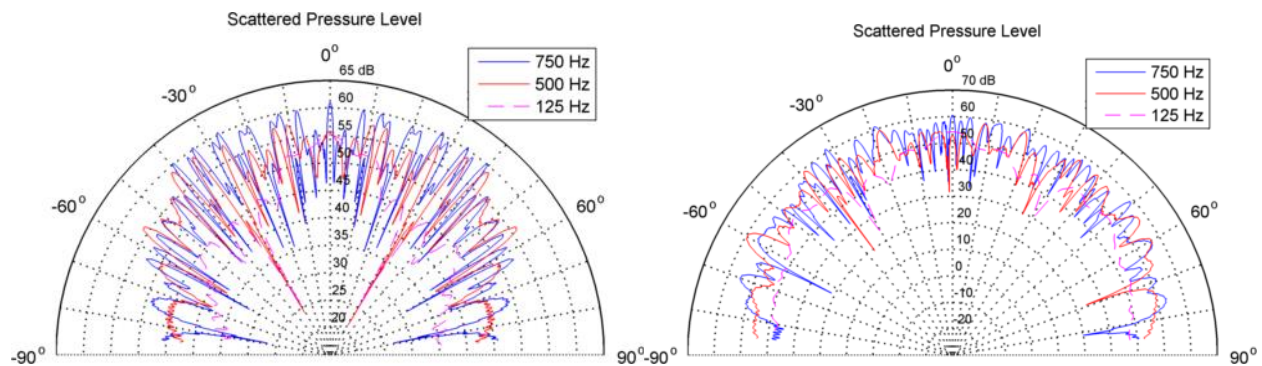


Figure 54 Scattering from the A1-Frac array (left) and B2-Frac array (right)

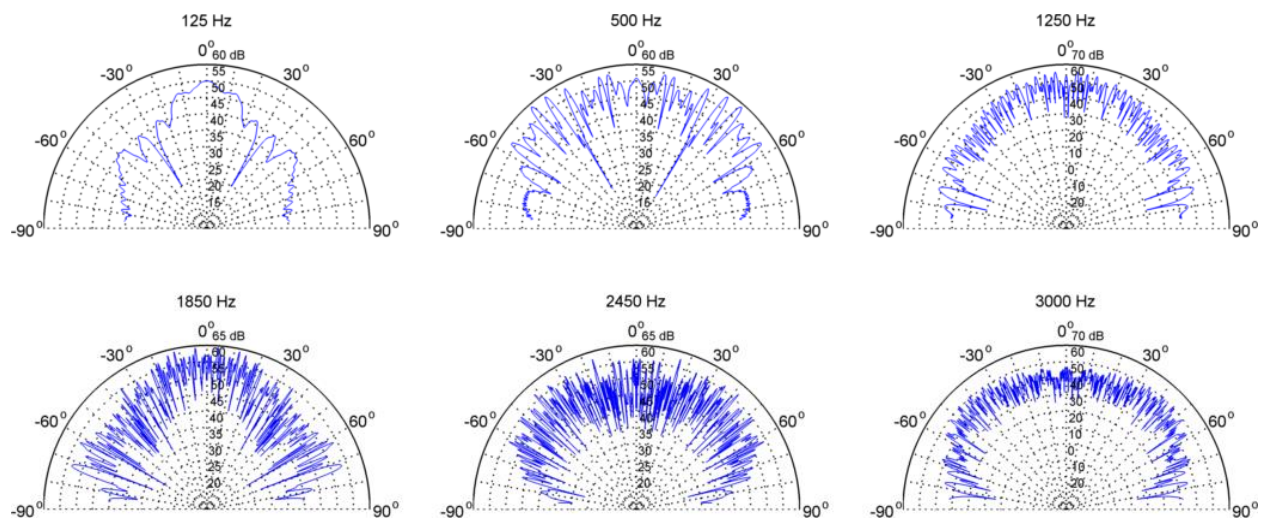


Figure 55 Scattered polar distribution for the diffuser array A1-Frac.

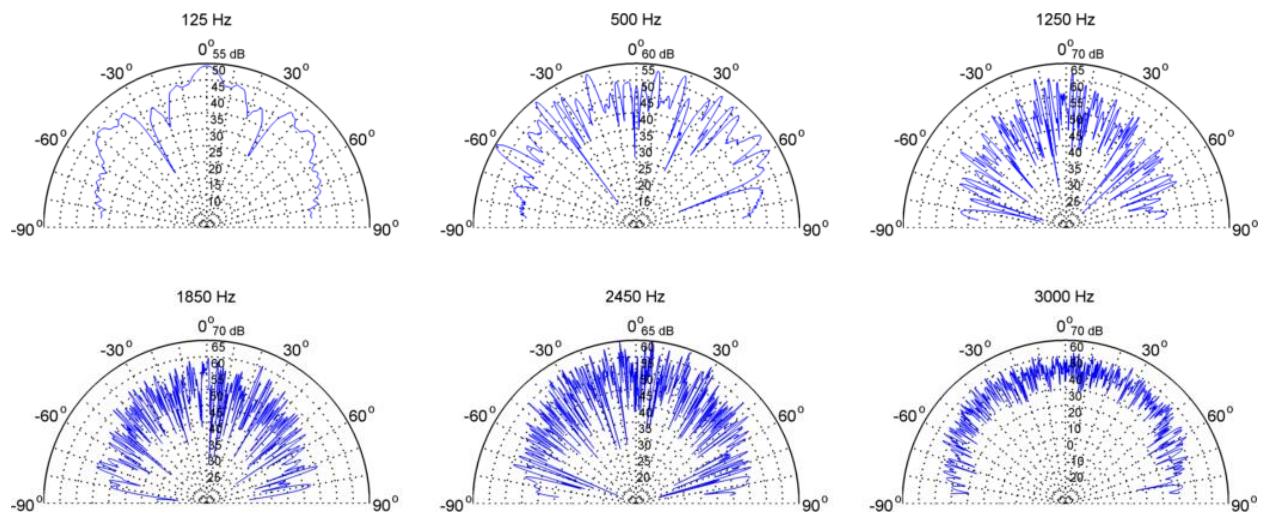


Figure 56 Scattered polar distribution for the diffuser array B2-Frac.

9.5. Comparing simple optimized diffusers

Figure 58 shows the autocorrelation diffusion coefficient for 5 periods of an N=7 stepped diffuser optimized by Cox [11]. This design will be called L95. It has higher performance below 630 Hz than the periodic array of A1-LF modules (Figure 57), which is expected, because A1-LF is much shallower. However, above 500 Hz the periodic array of L95 diffusers cannot compete with the performance of Profiled Modulation 1 (Figure 59).

When the A1-LF modules are arranged into Profiled Modulation 1, they achieve superior performance to L95, with similar build complexity and less operational depth (11 cm for Profiled Modulation 1 vs. 14.8 cm for L96).

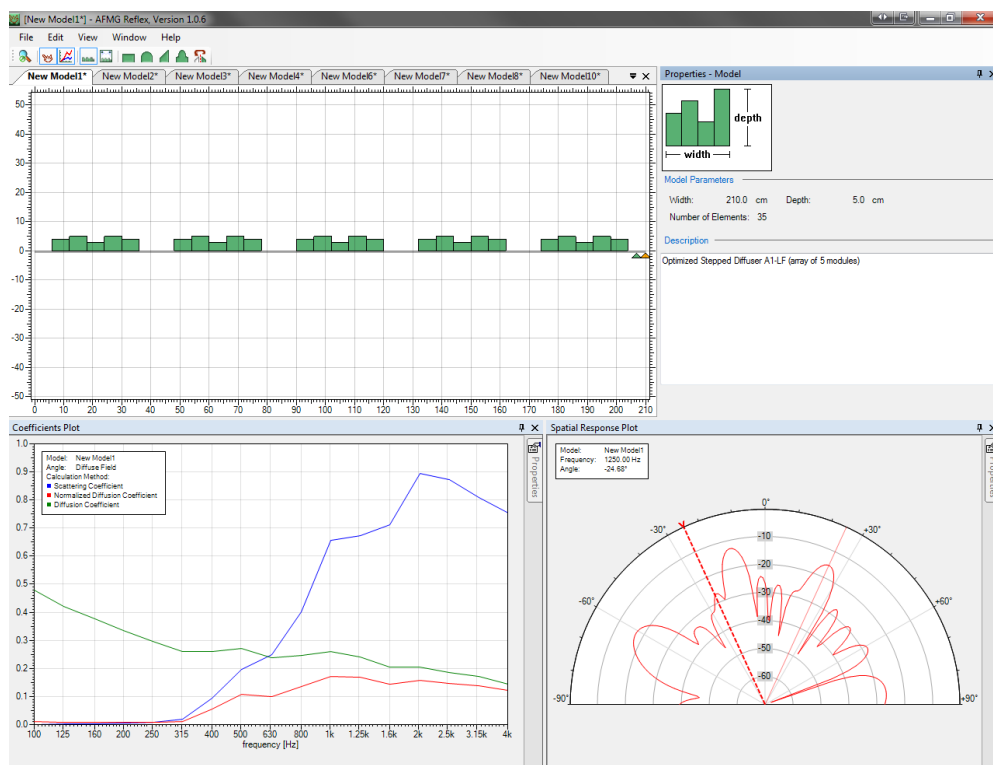


Figure 57 Diffuse field performance report for A1-LF array. Produced in AFMG Reflex.

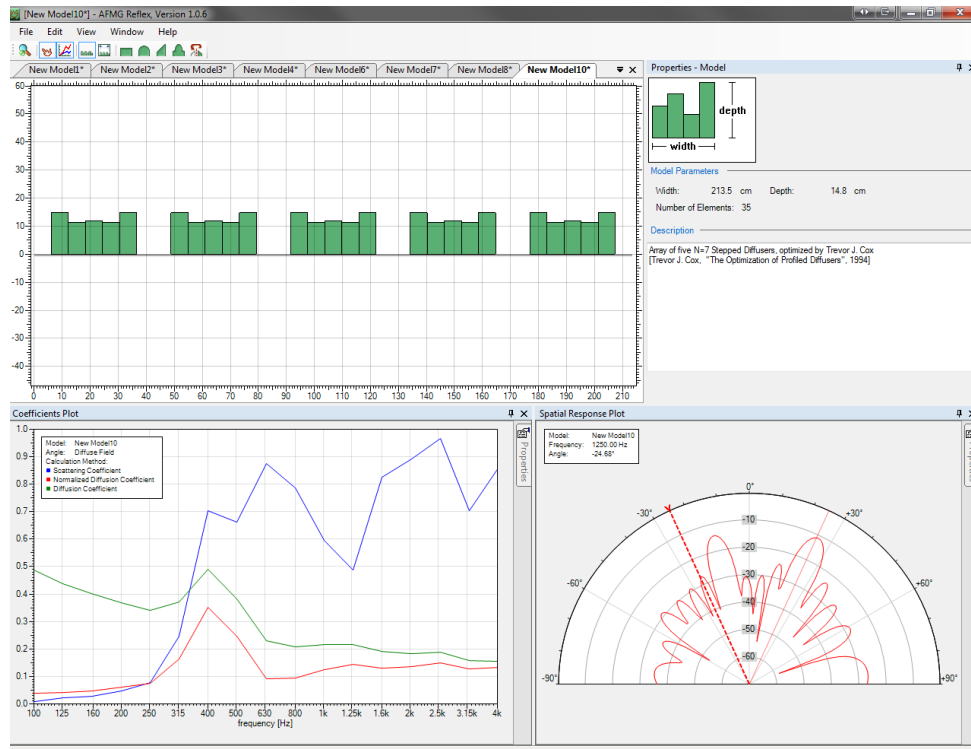


Figure 58 Diffuse field performance report for 5 periods of optimized diffuser L95 [11].

Depth Profile of Diffuser Array with Profiled Modulation 1 (Simulation View in AFMG Reflex)



Performance Coefficients for Profiled Modulation 1




Figure 59 Profiled Modulation 1 coefficient spectra. Diffuse field (left). 0° incidence (right).

10. CONCLUSIONS

The objective was to design a modular diffuser that provides an optimal trade-off between uniform scattering and compact geometry, and to design it without using boundary element predictions. The approach involved leveraging the strengths of finite difference time domain simulation, while accommodating the weaknesses. Success relied on cyclic prototyping and the intrinsically parallel property of evolutionary algorithms.

The following conclusions were drawn from this work:

1. A suitable solution to the design problem in Chapter 3 has been found. As shown in Sections 7.3 and 9.3, constrained optimization using seven variables produced a low profile modular stepped diffuser, A1-LF. A1-LF consists of five periods of a module with step heights of {0 4 5 3 5 4 0} cm. Of all candidate solutions evaluated, this ‘optimum’ design had the most uniform dispersion, which it accomplished with a structural depth of just 6 cm—one third the depth of typical stepped diffusers [11,1].
2. Diffusers with more uniform broadband scattering (better performance) were found by simulating fractal formations based on optimized diffusers.
3. While computationally intensive, time domain methods can be used as a practical prediction model for diffuser optimization. Special accommodations were necessary to make optimization feasible on the discrete time domain mesh. Namely, mixed integer programming was necessary because the solution space was constrained to discrete points. This was justified as an avenue worth exploring because it allowed a simple prediction model to be used, with no need for a well-depth impedance function. The resulting “lean” optimization framework chooses designs using low-resolution parallel search optimization, and simulates them with accuracy comparable to a BEM. For 1200 scattering evaluations the solution time was reasonable (about 40 hours). Additionally, FDTD and other time domain methods are worth exploring because they gives direct access to the temporal response, which may be used in future work to evaluate the temporal dispersion.
4. Optimization using an intrinsically parallel method that includes an element of randomness, like the integer genetic algorithm, is fruitful for finding designs with comparable performance but very different geometries.

5. Seven-well diffusers based purely on the set of natural numbers up to 16 can achieve very good performance. Traditional diffuser design has relied on a few known optimal number sequences (e.g., quadratic residue and primitive root sequences). In contrast, Cox [11] defined optimized stepped diffusers using high precision (e.g., the set of natural numbers up to 200). The current work executed a broader search of the solution space using low precision, such that solutions are based on simple number sequences. The results confirm that when designing a stepped diffuser with uniform broadband dispersion, there are numerous low-integer sequences that are superior to the quadratic residue sequence.
6. The designs are simple to implement. The optimization framework produces sequences of low natural numbers, useful for designers who need to work in low precision. The resulting designs can be constructed by carpenters or hobbyists, on-site, out of any suitable material that has a low absorption coefficient. Additionally, these low-natural number sequences reduce the cost of manufacturing and simulating fractal formations because they enable high accuracy with low precision.
7. Results support the hypothesis that diffusers optimized for uniform scattering over a 2.5 m receiver radius tend to exhibit as good or better diffusion at a 5 m receiver radius. Conversely, diffuser design standards focus on the 5 m receiver radius, which does not guarantee good performance at a smaller radius. In home studios and compact control rooms with suboptimal dimensions the 2.5 m receiver radius may be more realistic²⁵.

²⁵ This assumes that the diffuser array is placed with the intention of dispersing first reflections from either the back wall or a lofted ceiling.

10.1. Enhancements and Future Research

The optimization is designed to be functional, but not efficient. Future investigations might address the following:

1. If the near field to far field transformation (NFFFT) is utilized, FDTD scattering predictions will be computationally feasible at a much larger radius.
2. It is expected that the computation can be sped up by orders of magnitude if it is run on graphics processing units instead of a CPU.
3. To search the solution space more efficiently, multiple evolutionary algorithms can be run in parallel with the Matlab parallel computing toolbox. An alternative that was explored during testing is to use a meta-heuristic global optimizer that combines multiple optimization methods to increase the likelihood of finding the global optimum [26].
4. The computation can be greatly sped up if a boundary element model is used. The solution space will no longer be restricted to integers, but it may still be desirable to perform a low resolution global search of the design space to reveal unexpected regions that meet the design criteria. When a desired region is found, a continuous optimization method can be used to fine tune the well depths to achieve optimal diffusion.
5. An impedance model can be used to define diffusers instead of setting the particle velocity to zero at selected mesh nodes. This will enable optimization over a higher resolution solution space without the requirement for a tiny mesh spacing. With a high resolution design space, continuous optimization techniques will be able to establish clear trends and the computational complexity of finding the global optimum will be reduced.
6. Fractal forms may be optimized directly if one of the above approaches is used to either a) increase the efficiency of the scattering prediction or b) increase the resolution of the design space. The optimization algorithm could search for a solution over N degrees of freedom and by doing so optimize a 2nd order fractal module with N^2 wells (i.e., steps). Alternatively, N low frequency wells and N high frequency wells might be optimized independently and nested to form a quasi-fractal with N^2 wells.

7. A two-dimensional diffuser can be formed by using the Chinese remainder theorem to fold an optimized 1D design sequence into a 2D array [1].
8. Currently, optimization does not utilize all available data because it is performed at discrete frequencies. Instead, optimization can be applied across all FFT bins, using relative weighting to place greater importance on selected frequencies.
9. The prediction model needs to be validated by comparing simulation results to real world measured results. The simplest approach is to precisely implement and simulate a design that has already been validated, such as those outlined by Cox [11,1] and D'Antonio [3,1].
10. Data to analyze temporal dispersion is readily available from the time domain simulation. Future research might focus on interpreting this data [22,1].

REFERENCES

- [1] T.J. Cox and P. D'Antonio, *Acoustic Absorbers and Diffusers*, 2nd ed. Abington, Oxfordshire: Taylor & Francis, 2009.
- [2] M.R. Schroeder, "Diffuse sound reflection by maximum-length sequences," *J. Acoust. Soc. Am.*, vol. 65, no. 4, pp. 958-63, 1975.
- [3] RPG Diffusor Systems, Inc. (2000) CHAOS: The Collaborative Holistic Acoustical Optimization System. [Online]. <http://www.rpginc.com/products/chaos/index.htm>
- [4] T.J. Cox, "Acoustic diffusers: the good, the bad and the ugly," in *Proc. Inst. Acoust. Reproduced Sound 20th Conf.*, Oxford, 2004.
- [5] Digizine. (2011, Oct.) George Massenburg builds a blackbird room. Image. [Online]. http://www2.digidesign.com/digizine/dz_main.cfm?edition_id=101&navid=907
- [6] T.J. Cox and P. D'Antonio, "Acoustic phase gratings for reduced specular reflection ," *Applied Acoustics*, vol. 60, no. 2, pp. 167-86, 2000.
- [7] J.A.S. Angus, "Using grating modulation to achieve wideband large area diffusers ," *Applied Acoustics*, vol. 60, no. 2, pp. 143-65, 2000.
- [8] T.J. Cox and P. D'Antonio, "Fractal Sound Diffusers," in *Proc. Audio Eng. Soc. 103rd Conv.*, New York, 1997, p. 4578.
- [9] P. D'Antonio and J. Konnert, "The QRD Diffractal: A New One- or Two-Dimensional Fractal Sound Diffusor*," *J. Audio Eng. Soc.*, vol. 40, no. 3, pp. 117-129, March 1992.
- [10] RPG Diffusor Systems, Inc. (2000) New Diffusion Paradigm. [Online]. <http://www.rpginc.com/research/ndp.htm>
- [11] T.J. Cox, "The optimization of profiled diffusers," *J. Acoust. Soc. Am.*, vol. 97, no. 5, pp. 2928-36, 1995.
- [12] P. D'Antonio and T.J. Cox, "Two Decades of Sound Diffusor Design and Development, Part 1: Applications and Design," *J. Audio Eng. Soc.*, vol. 46, no. 11, pp. 955-976, November 1998.
- [13] P. D'Antonio and T.J. Cox, "Two Decades of Sound Diffusor Design and Development, Part 2: Prediction, Measurement, and Characterization," *J. Audio Eng. Soc.*, vol. 46, no. 12, pp. 1075-1091, December 1998.
- [14] T.J. Cox, "Designing Curved Diffusers for Performance Spaces," *J. Audio Eng. Soc.*, vol. 44, no. 5, pp. 354-364, May 1998.
- [15] T.J. Cox and P. D'Antonio, "Optimized Planar and Curved Diffusers," in *Proc. Audio Eng. Soc. 107th Conv.*, New York, 1999, p. 5062.
- [16] J.A.S. Angus and P. D'Antonio, "Two Dimensional Binary Amplitude Diffusers," in *Proc. Audio Eng. Soc. 107th Conv.*, New York, 1999, p. 5061.
- [17] P. D'Antonio, "Technical Bulletin on the Evaluation of the Kinetics Tuned

- Absorber/Diffuser Panel," RPG Diffusor Systems, Inc., Upper Marlboro, 2005.
- [18] H.D. Luke, H.D. Schotten, and H. Hadinejad-Mahram, "Binary and quadriphase sequences with optimal autocorrelation properties: a survey," *IEEE Trans. Inf. Theory*, vol. 49, no. 12, pp. 3271-3282, December 2003.
 - [19] T.J. Cox, J.A.S. Angus, and P. D'Antonio, "Ternary and quadriphase sequence diffusers," *J. Acoust. Soc. Am.*, vol. 119, no. 1, pp. 310-9, 2006.
 - [20] E.C. Payne-Johnson, G.A. Gehring, and J.A.S. Angus, "Improvements to Binary Amplitude Diffusers," in *Proc. Audio Eng. Soc. 122nd Conv.*, Vienna, Austria, 2007, p. 7143.
 - [21] T.J. Cox and P. D'Antonio, "Thirty years since "diffuse sound reflection by maximum length sequences": Where are we now?," *J. Acoust. Soc. Am.*, vol. 118, no. 3, pp. 2016-2016, 2005.
 - [22] J. Redondo, B. Pico, B. Roig, and M.R. Avis, "Time domain simulation of sound diffusers using finite-difference schemes," *Acta Acustica uw Acustica*, vol. 93, no. 4, pp. 611-22, 2007.
 - [23] C.K.W. Tam and T. Auriault, "Time-Domain impedance boundary conditions for computational aeroacoustics," *AIAA Journal*, vol. 34, no. 5, pp. 917-23, 1996.
 - [24] B.D. Treebly and B.T. Cox, "k-Wave: Matlab toolbox for the simulation and reconstruction of photoacoustic wave-fields," *J. Biomed. Opt.*, vol. 15, no. 2, p. 021314, 2010.
 - [25] AES-4id-2001, "AES Information document for room acoustics and sound reinforcement systems - characterisation and measurement of surface scattering uniformity," *J. Audio Eng. Soc.*, vol. 49, no. 3, pp. 149-65, 2001.
 - [26] R. Oldenheis and J. Vandekerckhove. (2009, August) Global Optimum Determination by Linking and Interchanging Kindred Evaluators. Matlab code. [Online]. <http://www.mathworks.com/matlabcentral/fileexchange/24838>
 - [27] M. Long, "Design of Studios and Listening Rooms," in *Architectural Acoustics*. Burlington, MA: Elsevier Academic Press, 2006, ch. 21, pp. 741-779.
 - [28] M. Noble, D. Kennedy. (2008, July) "Acoustic Design Performance in Green Buildings", Sustainable Architecture & Building Magazine (SABMag). [Online]. <http://www.sabmagazine.com/blog/2008/07/21/acoustic-design/>
 - [29] P.R. Newell, *Recording Studio Design*, 2nd ed. New York: Focal Press, 2007.
 - [30] M. Monks, B.M. Oh, and J. Dorsey, "Audiioptimization: Goal-Based Acoustic Design," *IEEE Comput. Graph. Appl.*, vol. 20, no. 3, pp. 76-91, June 2000.
 - [31] R. Bolejko and P. Pruchnicki, "Sound Diffusers Based on Number Theory with Random Variations of Surface Acoustic Impedance," in *Proc. Audio Eng. Soc. 104th Conv.*, Amsterdam, 1998, p. 4711.

- [32] L. Rizzi et al., "Scattering uniformity measurements and first relection analysis in a large non-anechoic environment.," in *Proc. Audio Eng. Soc. 123rd Conv.*, New York, 2007, p. 7241.
- [33] M. Rocha and J. Neves, "Preventing premature convergence to local optima in genetic algorithms via random offspring generation," in *Proc. IEA/AIE 12th Int. Conf.*, Secaucus, 1999, pp. 127-36.
- [34] T. Perry, "Acoustic diffuser design by optimization: I. Literature review.," University of Victoria, Prelude to Honours Thesis 2011.

Appendix A

Finite Difference Time Domain Matlab Code

```
%=====
% dfsrFDTD_lean.m                               Author:   Tim Perry
% Elec498: Diffuser Design by Optimization       Started:  2011-10-20
%
% Finite difference time domain (FDTD) simulation of a rigid-wall acoustic diffuser.
% Excitation source is a Ricker wavelet or a Gaussian pulse.
% Boundary conditions have not been implemented.
%
% REFERENCES:
% [1] Tiny_FDTD_v1.m by Nick Clark, 2007
% [2] T.J Cox, "Designing Curved Diffusers for Performance Spaces", 1996
% [3] T.J Cox and P. D'Antonio, Acoustic Absorbers and Diffusers, 2009
%=====
```

Code removed in this web version of the report.

Appendix B **Prediction of scattering Matlab code**

```
%=====
% diffuserFDTD.m                               Author:   Tim Perry
% Elec498: Diffuser Design by Optimization      Started:  2011-10-20
%
% FDTD prediction model used to evaluate the scattering from a stepped
% diffuser. Can be used as an objective function to be minimized by an
% optimization algorithm. Due to the FDTD mesh spacing, optimization is
% most naturally carried out using an when using an integer genetic algorithm.
%
% INPUT
%   depths0          the well depths of a potential diffuser design.
%                   e.g. [17 2 6 5 6 2 17]. Well depths will be rounded to the
%                   nearest grid step. Grid spacing depends on the boolean
%                   makeFractal. If makeFractal == true, the grid step is
%                   smaller and a fractal formation will be created.
%
% OUTPUT
%   eps1_prime       a single value broadband diffuser parameter to be minimized
%                   by an optimization algorithm.
%                   eps1_prime = eps1_ave + eps1_error, defined in [2]
%
%   eps1_ave         the average diffusion parameter at the frequencies of interest
%                   specified in the vector freqs2opt
%
%   eps1             the spatial-averaged diffusion parameter prior to
%                   frequency averaging (a vector in terms of frequency)
%
% REFERENCES
% [1] B. Treeby and B.T. Cox, 2011. k-Wave Toolbox (http://www.k-wave.org)
% [2] T.J Cox, "Designing Curved Diffusers for Performance Spaces", 1996.
% [3] T.J Cox and P. D'Antonio, Acoustic Absorbers and Diffusers, 2009.
%=====
```

Code removed in this web version of the report.

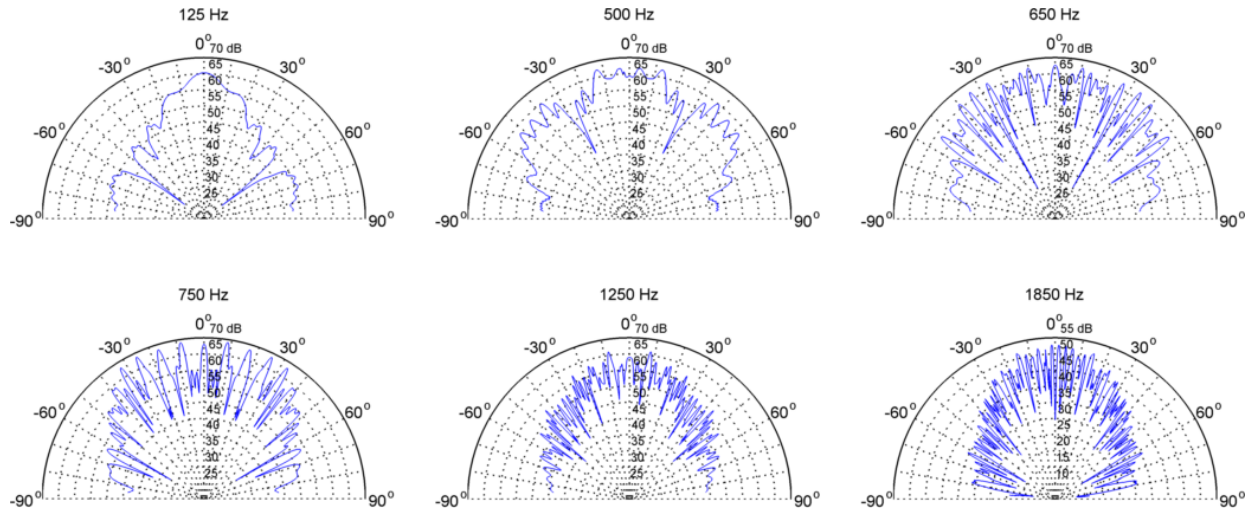
Appendix C **Matlab Scripts for Diffuser Optimization**

```
%=====
% Diffuser_Opti_Driver.m                               Author:   Tim Perry
% Elec498: Diffuser Design by Optimization               Started:  2011-11-16
%
%
% REFERENCES:
% [1] B. Treeby and B.T. Cox, 2011. k-Wave Toolbox (http://www.k-wave.org)
% [2] T.J Cox, "Designing Curved Diffusers for Performance Spaces", 1996
% [3] R. Oldenhuis, 2009 J. Vandekerckhove, 2006
%      GODLIKE - A robust single-& multi-objective optimizer
%      (http://www.mathworks.com/matlabcentral/fileexchange/24838)
%=====
```

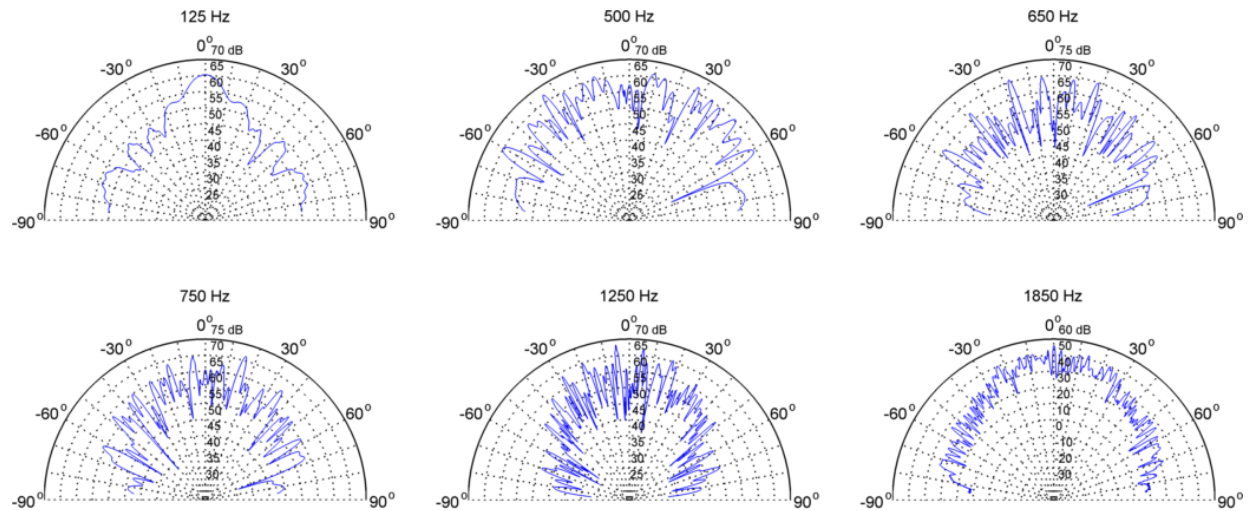
Code removed in this web version of the report.

Appendix D Scattered polar distribution plots.

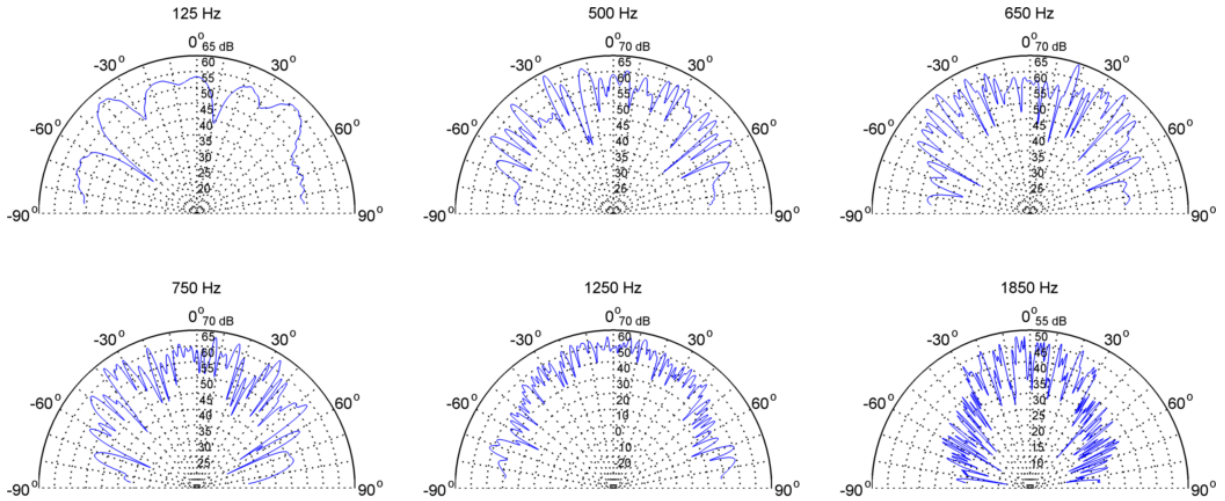
In the following plots, the scattered pressure level was obtained using the finite difference time domain simulation framework *Sim1* (see Chapter 9 for details). Units are dB SPL.



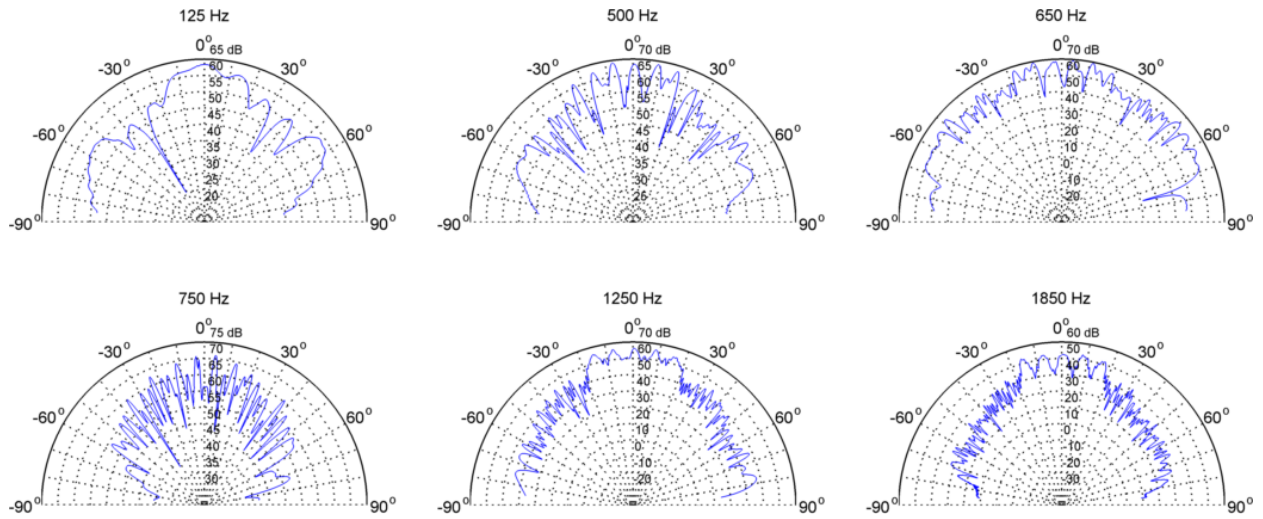
Scattered polar distribution for the optimized diffuser array *A1-LF*.



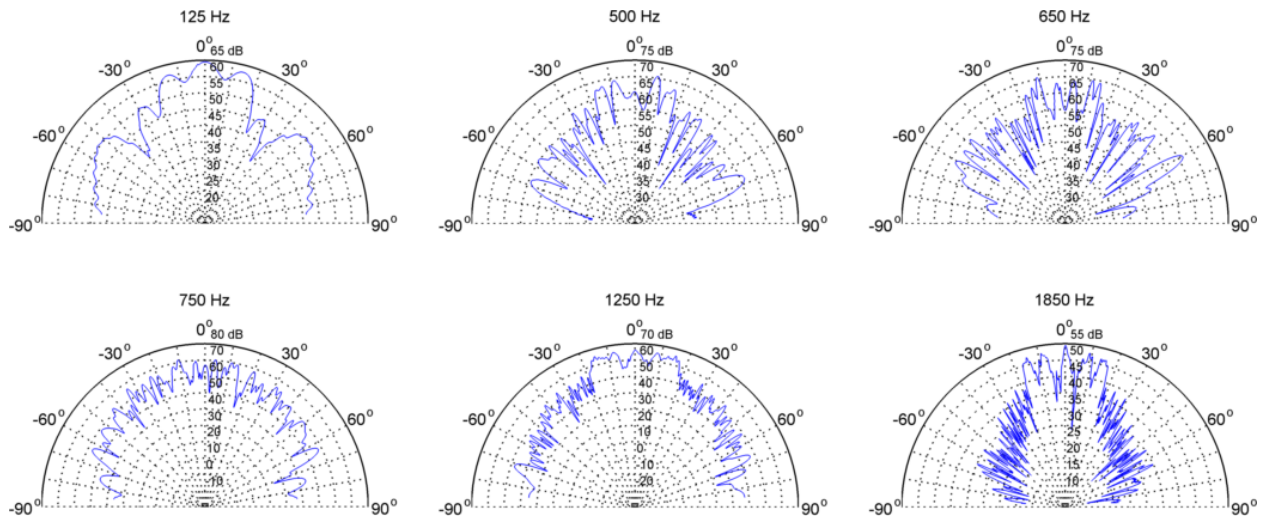
Optimized diffuser array *B2-LF*.



Optimized diffuser array *A3-LF*. This aperiodic design was optimized for diffusion below 1250 Hz.



Five periods of an $N = 7$ stepped QRD.



$N = 7$ stepped QRDs in the aperiodic array $\{1\ 0\ 1\ 1\ 0\}$.