

Efficient Wave-Based Acoustic Material Design Optimization

Nicolas Morales^a, Dinesh Manocha^a

^aUNC Chapel Hill

<http://gamma.cs.unc.edu/AcousticOpt>

Abstract

We present a novel approach for optimizing acoustic parameters using sensitivity analysis for computer-aided design and analysis of architectural models. Our approach builds on recent low-dispersion wave-based acoustic solvers that can accurately compute the pressure field in large models. We present an efficient technique to compute the gradient of the pressure field using automatic differentiation and combine that with a quasi-Newtonian optimization method to automatically compute the optimal material properties. We highlight the performance on many complex CAD models to optimize the strength and clarity acoustic parameters, and thereby improve the acoustic characteristics of large models. To the best of our knowledge, this is the first practical and accurate approach for acoustic material optimization of large indoor CAD models.

1. Introduction

Architectural and engineering design of structures often requires the incorporation of various design goals, such as functionality, reliability, operation, and aesthetics. Moreover, the design of these structures is often governed by specific constraints, such as performance, cost, maintainability, testability, and so on. Of particular importance is the interaction between sound waves and the structure. These sound waves are typically produced by human speech and noise, machines, musical performances, etc. The acoustic characteristics of a space can have an effect on the perception of that space, human communication, and behavior. These characteristics are usually determined by the shape, topology, structure and surface materials, and objects inside the acoustic space. Figure 1 shows an example of how the sound Strength, or sense of fullness of the sound, is affected by the materials inside the space.

The acoustic characteristics of architectural models are measured in terms of sound clarity, strength, delay, reverberation, etc. Different architectural models impose varying requirements on these acoustic characteristics. For example, the premium seats in concert halls often require a sound clarity measure (C80) of between -2dB and 4dB. Other constraints or standards are imposed due to health or environmental factors. The WHO recommends that the equivalent continuous noise level from the environment in hospitals during the night should not exceed 30dB [5]. The implementation of noise minimization procedures in hospitals can result in a significant drop in medical errors [25]. Studies have also shown that poor acoustics can have a negative effect on classrooms [6].

Recent trends in computer-aided design for acoustic design have focused on simulation technologies for prototyping architectural and engineering structures. For example, acoustic models are used in the design of airplanes to predict the noise caused by engine vibration and the propagation of acoustic waves throughout the aircraft cabin [18]. Additionally, manufacturers use large noise engineering laboratories for measuring airframe and aircraft noise [43]. Urban habitation designers [42] and automobile manufacturers [14] have also used acoustic simulation for prototyping designs. However, current acoustic simulation tools are limited in their accuracy and domain capabilities. Often they do not provide reliable solutions. A direct consequence of this lack of reliability is that acoustic design is frequently performed by human designers with a limited set of acoustic simulation tools. It is also common for acoustic engineers to build physical prototypes to validate the acoustic characteristics of their design. This can cause long design cycles or even non-optimal acoustic designs [3, 30, 28].

Main results: We present a novel approach for optimizing the acoustic material properties of CAD models for the purpose of simulation-based acoustic engineering design. Our approach is designed for large architectural models and is based on accurately computing the acoustic pressure field as a function of the material properties.

Our formulation is based on using a low-dispersion numerical solver for the acoustic wave equation, called adaptive-rectangular decomposition (ARD). We present a novel and efficient sensitivity analysis of the ARD wave solver with respect to the input acoustic materials. The resulting sensitivities are used to drive a gradient-based iterative algorithm that can automatically optimize the material properties to satisfy the acoustic design criteria of the space. Overall, the three novel components of our

Email address: nmorales@cs.unc.edu (Nicolas Morales)

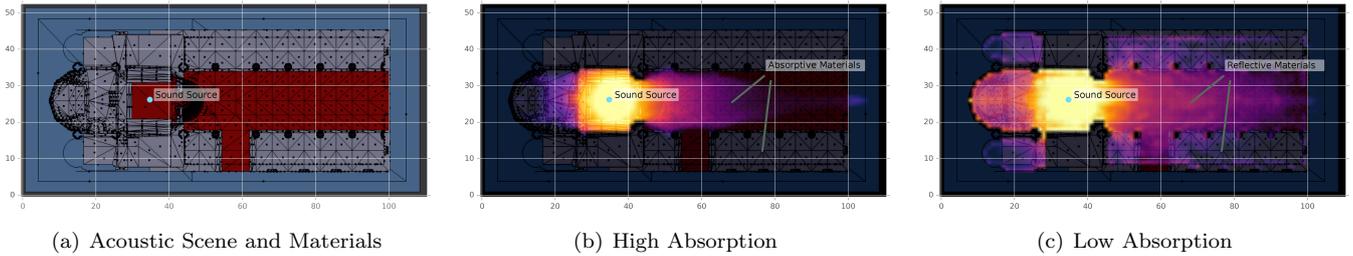


Figure 1: An example of how the materials in a building can affect the acoustic properties of that building. In this figure the materials are each assigned a different color in (a). Parts (b) and (c) show the acoustic Strength (G), or feeling of loudness and fullness of sound, at each point in the scene for different material configurations. The brighter color corresponds to a higher Strength. Part (b) shows the scene with almost fully absorptive materials, where sound waves are mostly absorbed by the walls and floor. Note how overall, high acoustic Strength values are limited to those areas with a direct line of sight to the sound source. Part (c) shows a low absorption scene, where sound waves are reflected and keep most of their energy. In this case, listeners can get a better feeling of fullness and loudness without being directly in front of the sound source.

work include:

- Wave-based sensitivity analysis of the acoustic pressure field for design optimization.
- A fast acoustic design optimization algorithm based on the ARD acoustic wave solver.
- A multi-objective acoustic material optimization that can simultaneously optimize for various acoustic properties including strength and clarity.

We show the results of optimizations using our algorithm on large scale 3D scenes, including CAD models of well-known architectural models. We automatically compute the material properties to give tight bounds on the acoustic characteristics of the resulting models. The overall approach is general purpose and takes many tens of minutes on large architectural models of about $20\,000\text{ m}^3$ volume on a desktop computer. To the best of our knowledge, this is the first method able to perform accurate sensitivity analysis for acoustic material design optimization of large CAD models.

The rest of the paper is organized as follows. Section 2 gives an overview of prior work in acoustic simulation and application. Section 3 introduces the acoustic wave solver used in our algorithm and the acoustic metrics used to evaluate the design. Section 4 describes our algorithm for acoustic material optimization and our method for performing sensitivity analysis on the acoustic pressure field. Section 5 presents the efficient implementation of our algorithm. Section 6 provides an overview of our results and analysis of those results. Finally, Section 7 highlights the benefits of our algorithm.

2. Prior Work

Computational acoustics is an area of active research in engineering design and scientific computing, and is also studied in seismology, geophysics, and meteorology. In this section, we limit ourselves to computational acoustic methods for large architectural models.

2.1. Simulation and Computer Aided Design

Extensive research and software development has focused on vibration analysis, interior and exterior acoustic radiation computation, vibro-acoustics, and aero-acoustic modeling [19, 8]. Much work has focused on modelling of Noise, Vibration, and Harshness (NVH) measurements for car interiors [13]. Additionally, many commercial tools are available for acoustic analysis of objects, structures, or small spaces. Often, these tools are not sufficiently accurate or applicable to large acoustic spaces prevalent in architectural design such as auditoriums, concert halls, or outdoor environments where the volume may exceed $10\,000\text{ m}^3$ to $100\,000\text{ m}^3$.

Analysis of the acoustic characteristics of architectural spaces is often studied in the context of room acoustics [21]. Work in this field has been done by Sabine dating back to the early 1900s. This work was conducted through ray-based acoustics and Sabine and Eyring’s reverberation time formula for rectangular rooms. More recently, geometric acoustic techniques based on ray-tracing [20, 44, 37] have become prevalent in the evaluation of indoor acoustic designs such as concert halls, theatres, and auditoriums. Geometric acoustic methods are used in several commercial packages but suffer accuracy issues because of the underlying assumption that sound propagates as rays rather than as waves. Therefore, wave effects of sound that are prevalent at lower frequencies, such as diffraction and scattering, are often neglected.

Wave-based methods, on the other hand, directly solve the acoustic wave equation and do not suffer from these accuracy issues. Numerical solvers for wave-based methods include finite difference methods [40], finite element methods [41], and boundary element methods [15]. However, the computational and memory requirements for these methods are much higher and, as a result, wave-based techniques are usually limited in practice to small spaces (less than 1000 m^3 , for example) and low frequencies (less than 2 kHz). Recent advances in wave-based methods have reduced the computational and memory complexity of these algorithms. These works include low dispersion methods such as the Adaptive Rectangular Decomposition (ARD) method [34, 27, 31] and the equivalent source method [26].

2.2. Acoustic Optimization

A common problem in mechanical and architectural design is the automatic optimization of a set of parameters on the computational domain. This field, part of Multi-disciplinary Design Optimization (MDO), is widely used in many areas including the aerodynamic optimization of wings and entire aircraft, architectural features such as bridges and buildings, railway cars, microscopes, automobiles, turbines, and ships [24]. Many commercial design optimization tools are available for building information modeling (such as Autodesk Revit building design software tools) and are frequently used for lighting analysis, structure analysis, energy analysis, and so on. Often these tools can take advantage of the computational capabilities of large distributed clusters or large-scale cloud computing. However, the state of the art in acoustic design for large architectural spaces is still at its infancy.

The acoustic optimization problem is a subset of MDO, and is useful for designing acoustic spaces or engineering structures according to certain target acoustic metrics. These metrics can range in complexity from sound intensity minimization to sound clarity or reverb time to binaural acoustic evaluations [17]. Much of the prior work in acoustic optimization has focused on Noise, Vibration, and Harshness measurements (NVH). A more limited set of methods has focused on space optimization, such as methods targeting architectural acoustics. We introduce two terms to describe these models. The first term, *object models*, refers to sound and vibration traveling through the objects. *Spatial models*, on the other hand, refers to sound waves propagating through the atmosphere in the space within a structure.

2.2.1. Object Models

Object models of acoustic optimization deal with the analysis of vibrations or sound propagation through individual objects rather than through acoustic spaces. Most of these models have been used for the design of automobiles, engines, and architectural support structures such as beams. Maressa et al. [23] introduce a method for NVH measurements for a car interior that is represented by a finite element mesh. This method explicitly studies the relationship between the structure of the object and the acoustics in the object. Other techniques that study the acoustic-structure interaction are based on a unified approach [45], which uses a mixed formulation to represent both the acoustic propagation and the elastic displacement of the structure. Nandy et al. [32] bypass acoustic simulation entirely in the optimization process. Du, Song, and Olhoff [9] present a method of acoustic-structure interaction at a finite boundary around the vibrating object. Shu et al. [38] use a level set based topology optimization to minimize sound resulting from vibrations in an outer structure. The topology optimization allows certain structures, such as beams, to be generated in order to reduce the sound.

2.2.2. Spatial Models

Our goal is to optimize the acoustic material properties of architectural models to satisfy some constraints on the acoustic characteristics. This is an example of *spatial acoustic optimization* and the driving application is acoustic design and optimization of large architectural models. In this context, there are three main subproblems. The first, material optimization, involves the modification of the materials of the scene as parameters of the optimization problem. The second, shape optimization, deals with the modification of dimensions of certain aspects of the scene, but not the actual topology. The third and last one deals with topology optimization.

Material optimization problems limit the number of parameters being optimized to a discrete set of materials for each surface or to an arrangement of material placements. Work on this has been done by Saksela et al. [36] and Monks [30] using geometric acoustic solvers. Shape optimization, on the other hand, reduces the optimization parameters to a smaller set of dimensions or measurements. For example, Robinson et al. [35] parameterize the shape of balconies in a concert hall in order to determine the best shape for sound clarity and strength in all areas in the concert hall. Other shape optimization approaches [12] minimize the acoustic pressure in a specific section of a room by modifying a series of columns at the top of the room. Floody and Venegas [11] modify the dimensions of a rectangular room in order to reduce resonance frequencies. In the area of topology optimization, Dühring et al. [10] minimize the noise in a target area in an output domain by modifying the topology of the design domain. The modification is performed by discretizing the design domain and assigning each element a value between 0 and 1, where 0 corresponds to air and 1 represents an aluminum material. They use an adjoint sensitivity approach and use the Method of Moving Asymptotes to solve the optimization problem.

3. Background and Acoustic Wave Solver

In this section, we give a brief overview of room acoustics and the wave propagation solver. We also highlight various symbols and notation used in the rest of the paper in Table 1.

Our main goal is to optimize for a set of acoustic properties for the scene. Some commonly used acoustic properties are: Onset Delay (Onset), Onset Direction (Dir), Reverberation (RT60), Definition (D), Clarity (C80), and Strength (G) [17, 21]. These properties are all derived from the impulse response (IR) of the scene, as it is excited by an impulse sound. The acoustic characteristics are also dependent on the scene configuration, Ω . The scene configuration includes the acoustic material properties of objects in the scene in addition to the geometric representation and layout of the scene. The actual process of computing a single acoustic parameter is represented as $f(\Omega)$, where

c	the constant speed of sound
t	time
Δt	time step
k	material index
Ω, Ω_k	scene configuration and materials
i	acoustic metric index
$f_i(\Omega)$	acoustic metric of the domain Ω
Z_i	target acoustic metric
$\text{IR}(\Omega)$	impulse response on the domain Ω
$g(\Omega)$	acoustic solver output
\vec{x}	spatial position
$p, p(\vec{x}, t)$	acoustic pressure in Pascals
$F, F(\vec{x}, t)$	acoustic forcing term
w_i	weighting function for acoustic metric i
z_i	target value for acoustic metric i
ω	characteristic frequency
M	pressure in mode-space
\tilde{F}	forcing term in mode-space
S	FDTD stencil operator
Ω_i, Ω_j	subdomain in the rectangular decomposition
Γ_{ij}	interface between subdomains Ω_i and Ω_j

Table 1: Notation used for the acoustic solver and optimization algorithm.

$f()$ is an acoustic wave propagation solver that computes the sound pressure field throughout the space and then evaluates the acoustic metric.

In order to solve the function $f(\Omega)$, we want to use an efficient and accurate acoustic solver. At a broad level, prior work on acoustic simulation can be divided into two categories: geometric solvers and wave-based solvers. The former can compute efficient and often real-time evaluations of the impulse response of a scene for different sources. However, they are only accurate for high frequencies, and do not accurately capture the wave nature of sound, which includes such properties as diffraction and scattering. Wave-based numerical methods on the other hand directly solve the wave equation, and compute the pressure field throughout the space. However, their complexity increases as a linear function of the volume of the acoustic space and the fourth power of the frequency.

In our optimization method, we use an accurate wave-based solver that can accurately compute the pressure field. In particular, we use a low-dispersion method (ARD) that can handle large acoustic spaces efficiently, and is more than an order of magnitude faster than prior time domain solvers such as FDTD [34]. Furthermore, ARD has been parallelized for computational efficiency on both GPUs and distributed memory CPU clusters [27, 31]. As a result, it can also be used for higher frequency simulations. The main advantage of using ARD is computational efficiency. Since our acoustic optimization algorithm requires several evaluations of the optimization function, the computational efficiency of evaluating this function is important.

ARD solves the time-domain wave equation:

$$\frac{\partial^2}{\partial t^2} p(\vec{x}, t) - c^2 \nabla^2 p(\vec{x}, t) = F(\vec{x}, t), \quad (1)$$

where \vec{x} is a 3D position, t is time, $p(\vec{x}, t)$ is the pressure at point \vec{x} and time t , F is a forcing term at point \vec{x} and time t , and c is the speed of sound. ARD assumes a homogeneous environment where the speed of sound c is constant.

Each subdomain in ARD assumes a perfectly reflective boundary condition. The sound field is then propagated across subdomains using an FDTD stencil at the interface. The FDTD stencil is equivalent to a (2,6) FDTD scheme used for full-field acoustic simulation. Therefore, this stencil solves the acoustic wave problem at the interface between rectangular regions. Since it acts as a close approximation to the exact solution, the stencil matches and cancels out the reflections induced by the reflective boundary condition of the subdomain and propagates the pressure field across the interface. Additionally, ARD implements the Perfectly Matched Layer (PML) absorbing boundary condition in special partitions both at the boundary of the scene and at the walls. Variable absorptivity of the walls is done by using an absorptivity factor on the result of the interface to the wall. If this value is zero, then the FDTD stencil is not taken into effect and the wall is regarded as perfectly reflective. On the other hand, if the value is one, the FDTD is taken into full effect and the wave is propagated into the PML partition where it is absorbed. Similarly, values used between zero and one result in partial reflectance [34].

A limitation of this approach is that only a single absorption value is used that does not take into account the dependence on the frequency. For more realistic materials, it is necessary to run the simulator for multiple frequency bands, i.e. multiple bands in the human hearing range 20 Hz to 20 000 Hz, though in practice current solvers are limited to a few thousand kilohertz.

4. Acoustic Optimization and Sensitivity Analysis

In this section, we present our novel approach for wave-based sensitivity analysis for acoustic design optimization. In the process of designing such an approach, many issues arise in terms of automatically determining the best acoustic materials for the desired acoustic characteristic.

4.1. Acoustic Optimization

Our goal is to design an optimizer that can target multiple acoustic design parameters. For example, a designer might want a specific Strength and a specific Clarity value for a concert hall. We use a linear weighted sum optimization function, in which different target characteristics are weighted according to their importance. This allows a designer to put more emphasis on some acoustic characteristics over others. Therefore, our formulation of the optimization problem is expressed as:

$$\min \sum_{i=1}^n w_i \| f_i(\Omega) - Z_i \|, \quad (2)$$

where n is the number of acoustic properties, $w_i \geq 0$ is the weighting for the acoustic property i , $f_i(\cdot)$ is the calculation of the acoustic property on the domain Ω , and Z_i is the target value for the acoustic property i .

The function $f(\Omega)$ is calculated by using an acoustic solver. This solver could be wave-based or geometric, but in our case we use a wave-based simulator. The input to the function is the geometric representation and layout of the scene and the set of acoustic materials in the scene (Ω). The output is a vector of the characteristic acoustic properties mentioned earlier. We use this linear combination in preference to a weightless norm because we wanted designer control over the relative importance of different acoustic characteristics. For example, in a concert hall, the parameter C80 is probably the most important along with RT60 [21]. Moreover, a linear-weighted objective function has been shown to be sufficient for combining various acoustic metrics [30].

Our optimization pipeline (Figure 2) computes the full derivative of the optimization problem. This is used to guide a gradient-based method such as gradient descent, conjugate gradient methods, or quasi-Newtonian methods. As detailed in Section 4.4, the full derivative of the acoustic pressure field is actually computed along with the ARD based acoustic pressure solver.

4.2. Sensitivity Analysis

As part of our optimization algorithm, we need to design a method to efficiently compute the sensitivity of the ARD solver with respect to the input material properties. Our goal in terms of sensitivity analysis is to determine how each acoustic characteristic of the scene changes when these inputs are changed. For example, if we have a scene with the set of material properties Ω , the sensitivity is the full derivative of the function f (Equation 4.1). We assume that the input parameters to this equation (i.e. the acoustic materials) do not occur in any other parameter, so the full derivative is simply the gradient of f . Therefore we have:

$$f^{(k)'} = \nabla_{\Omega_k} f, \quad (3)$$

where $f^{(k)'}$ is the full derivative of f when one of the materials is modified. The calculation of this derivative can be performed using finite differences, but this approach has inaccuracies that can lead to issues in the optimization problem [39] or may take a greater number of steps to converge to the target value.

The computation of this full derivative can be non-trivial. ARD is a domain decomposition method in which each subdomain Ω_i is tightly coupled with a neighboring subdomain Ω_j over the interface Γ_{ij} . In ARD, the analytic solution to the wave equation over each subdomain is used for the pressure computation, but the interfaces use a 6th

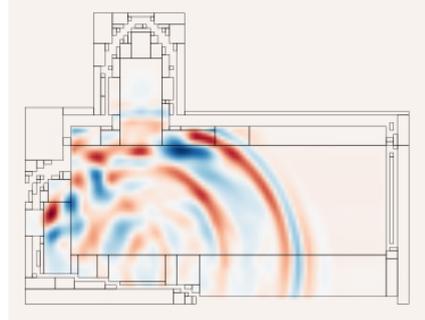


Figure 3: Domain decomposition of a cathedral using the ARD method.

order FDTD stencil for the computation. Figure 3 shows an example of this decomposition and the spatial complexity of a scene with many subdomains and interfaces.

4.3. Computing the Pressure Field Gradient

ARD decomposes the scene into rectangular subdomains which can solve the time-domain acoustic wave equation analytically. The analytical solution inside these domains can be described by the discretization in the modal space of Equation 1 [34]:

$$\frac{\partial^2 M}{\partial t^2} + \omega^2 M = \text{iDCT}(F(t)), \quad (4)$$

where M is the mode coefficient, ω is the characteristic frequency, t is time, and $F(t)$ is the forcing term. This leads to the following update rule for the rectangular subdomains:

$$M^{n+1} = 2M^n \cos(\omega\Delta t) - M^{n-1} + \frac{2\widetilde{F}^n}{\omega^2}(1 - \cos(\omega\Delta t)), \quad (5)$$

where Δt is the time step and \widetilde{F}^n is the forcing term in mode-space [34].

Our goal is to compute the gradient of the pressure field as the acoustic material characteristics in the scene change. Therefore, we propose a modification of the update rule that take the derivatives into account. For a single material, Ω_k , we have:

$$\begin{aligned} \frac{\partial M^{n+1}}{\partial \Omega_k} &= 2 \frac{\partial M_i^n}{\partial \Omega_k} \cos(\omega_i \Delta t) - \frac{\partial M_i^{n-1}}{\partial \Omega_k} \\ &\quad + \frac{2}{\omega_i^2} \frac{\partial \widetilde{F}^n}{\partial \Omega_k} (1 - \cos(\omega_i \Delta t)), \end{aligned} \quad (6)$$

We then apply the chain rule, and can therefore expand the three terms in this equation. $\frac{\partial M_i^n}{\partial \Omega_k}$ and $\frac{\partial M_i^{n-1}}{\partial \Omega_k}$ are similarly updated at different time steps. Of particular interest is $\frac{\partial \widetilde{F}^n}{\partial \Omega_k}$, however. \widetilde{F}^n is the mode-space component of the forcing term F :

$$\widetilde{F}^n = \text{DCT}(F^n). \quad (7)$$

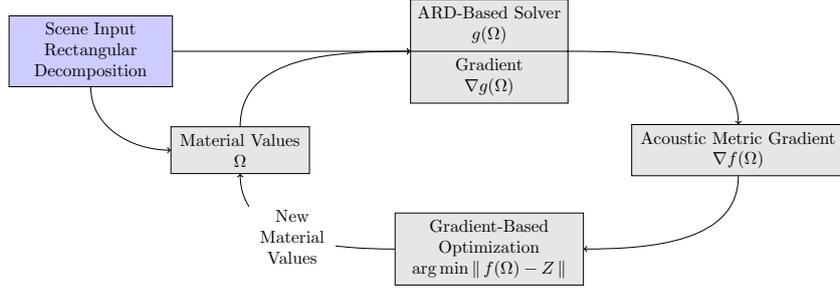


Figure 2: Our optimization pipeline, which computes the full derivative at the same time as it computes the pressure field. The scene input determines the initial material values and geometry. At each optimization step, a new set of material absorption values is computed.

The DCT is the sum of coefficients of cosines, so the derivative is simply expressed as:

$$\frac{\partial \widetilde{F}^n}{\partial \Omega_k} = \frac{\partial \text{DCT}(F^n)}{\partial \Omega} \quad (8)$$

$$= \text{DCT}\left(\frac{\partial F^n}{\partial \Omega}\right). \quad (9)$$

However, F^n may or may not be dependent on certain materials. If the forcing term is not dependent on any material, such as the forcing term originating from a sound source, then the derivative is zero. However, forcing terms originating from interfaces may be dependent on the material parameter. Finally, forcing terms originating from wall interfaces are directly and possibly indirectly dependent on the wall material. Consider the case where F^n originates from a wall interface. This is governed by the following equation:

$$F^n = \Omega_k c^2 S^n, \quad (10)$$

where S^{n-1} is a (2,6) FDTD stencil applied at the interface. This yields the following derivative:

$$\frac{\partial F^n}{\partial \Omega} = c^2 \left(\Omega_k \frac{\partial S^n}{\partial \Omega_k} + S^n \right). \quad (11)$$

This formulation isolates the material term. However, $\frac{\partial S^n}{\partial \Omega_k}$ may still be dependent on the material, since the stencil covers pressure field locations that may have sound dependent on other walls or even previous steps of the same interface.

These dependencies present problems for computing the derivative. We aim to present a *general purpose* optimization method that does not depend on a specific scene configuration. However, the derivations presented in this section show several cases in which the dependency of various forcing terms or pressure values is unknown. These dependencies are intrinsically related to the geometry and material parameters of the scene, or the pressure field at previous time steps. ARD is composed of many coupled systems including an analytical solver inside the cuboid subdomains, an FDTD solver at the interface, and a PML implementation at the boundary and wall partitions. The way in which these different systems are coupled together depends heavily on the scene configuration. Figure 4 shows the relationship and dependencies among the equations of different subsystems of ARD.

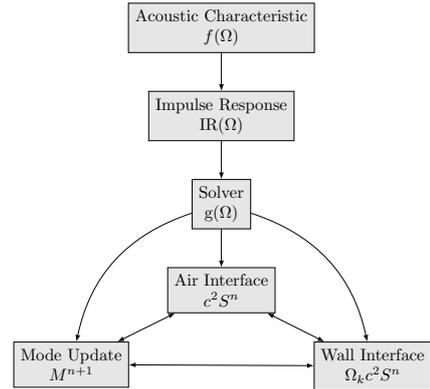


Figure 4: Dependencies of different subsystems of the ARD solver. These dependencies add to the complexity of the derivative calculation. Automatic Differentiation inherently deals with the problems of subsystem dependencies. The impulse response is an acoustic measurement of the room used to calculate various acoustic metrics.

4.4. Automatic Differentiation (AD)

Recently a class of methods called Automatic Differentiation methods (AD) has been developed in order to find the sensitivities of various simulation and engineering codes [2]. The advantage of AD methods is that they compute the analytical derivative of the code (as written). As such, they produce more accurate gradients than numerical methods such as finite differences. Automatic Differentiation works through repeated applications of the chain rule. In this way, a tree of operations and their derivatives can be combined automatically. Two methods of derivative accumulation can be used: forward accumulation and reverse accumulation that determine the order in which the chain rule is applied. Reverse accumulation can allow an AD technique to easily compute a large number of derivatives.

These properties make AD an ideal candidate for a general purpose solver like ARD. AD techniques can deal with the inherent complexity of the system, as dependency information is retained with each operation and application of the chain rule. For example, Equation 8 can either use a forcing term derivative from a wall interface (Equation 11) or an air interface. In Automatic Differentiation, the computation of these derivatives is performed at the same time that the pressure field is computed. If the scene configu-

ration or the rectangular decomposition is changed, the application of AD will yield the correct derivative for that configuration.

Additionally, our AD-based algorithm takes advantage of some of the efficiencies of ARD. Since ARD is a low-dispersion method, we can use a coarser discretization of the spatial domain as compared to other methods such as FDTD. This means that the update equation (Equation 5) needs to be evaluated at fewer points. As a result, the derivative update (Equation 6) can be computed with similar efficiency. This is especially important considering that computing the derivative essentially multiplies the memory requirements of the system by the number of components in the gradient. In our case, the number of components is the number of materials in the scene. The number of components in the gradient has a similar effect on computation, where the time cost of evaluating the gradient is linearly proportional to the number of components in the gradient.

5. Implementation

The following section discusses the implementation details of our algorithm. We also discuss the scientific computing software we use to provide a fast and efficient implementation of both our optimization step and gradient computation step.

5.1. Optimization

For the optimization method, we chose to use a standard BFGS algorithm. Using this method is advantageous because of the non-smooth nature of the optimization function. BFGS is a gradient-based optimization method that can deal well with non-smooth functions [22]. Additionally, BFGS is a subset of quasi-Newtonian methods that are designed to solve multidimensional optimization problems [7]. Because we are interested in architectural models that have multiple materials that we would like to optimize, this is particularly useful.

The implementation we use is the Dlib library which provides both BFGS and LBFGS and algorithms. The optimizer is driven by the acoustic solver ARD (see Section 3) and derivatives produced by Automatic Differentiation on the solver code (see Section 4.2). We use a delta stop condition to detect minima — when there is little change in value the optimizer finishes.

Part of the input of the solver is a set of *material segments*. These are regions of the architectural model that are assigned the same material. For example, if an acoustic engineer wanted to determine the optimal absorption of the carpet for obtaining a certain acoustic strength value, she could mark the floor as a material segment. These segments can be specified in any standard modelling program.

For optimization, material values are directly driven by the optimization algorithm. Some materials must remain

constant, however. These materials are either specified by the user to not be optimized (as would be in the case where parts of the architecture must be made of a certain material) or are part of the free-field boundary condition of the scene. ARD implements a perfectly absorbing boundary condition at the edges of the scene to simulate open areas. Therefore, the solver and optimizer work well with both indoor and outdoor scenes, including hybrid scenes with aspects of both.

The vector of these material segments is initialized with a specified or random starting value per material segment and then modified by the optimizer until convergence. The random starting value allows multiple processes of the optimizer to run in parallel in order to avoid local minima.

5.2. Automatic Differentiation

Our algorithm for acoustic material design optimization uses Automatic Differentiation for computing the sensitivity of the acoustic ARD solver. It takes advantage of ARD’s inherent efficiency and to handle the complexity of the ARD solver.

For computing the sensitivity of ARD, we use an implementation of the Sacado AD library developed by Sandia National Laboratories as part of the open source Trilinos scientific computing package [33]. Sacado is used in many scientific computing applications for calculating the sensitivity of various simulation codes. We chose it because of the inherent efficiency of the library – it uses expression templates and operator overloading in C++ to simplify expressions at compile time.

Because Sacado uses operator overloading, we can compute the derivative of the impulse response at the same time we compute the impulse response itself. As a result of the chain rule, further post-processing to derive the acoustic characteristics of the scene (strength, clarity, etc) will also contain the correct derivatives. Furthermore, we only have to evaluate this once per optimization step. This favorably compares to finite difference techniques, where we would need two evaluations of the simulation and acoustic characteristics to compute the derivative.

6. Results and Analysis

In this section we examine the results of our design optimization process for acoustic materials. We show convergence on various scenes and the effect of the optimization, including impulse responses and the value of various acoustic characteristics before and after optimization.

6.1. Benchmarks

A variety of scenes and acoustic materials were used in our experimental setup. All simulations were done with a low frequency (187 Hz) Gaussian derivative pulse as a sound source. Source and listener positions were placed in reasonable locations. We used three benchmark scenes in total: Cathedral, Twilight, and Concert Hall. These

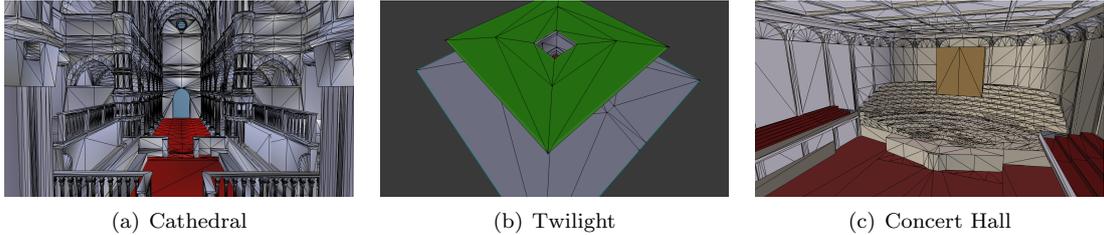


Figure 5: Geometry and material segmentation of the various benchmark scenes. These are specified by the designer. Our algorithm is general purpose and allows arbitrary flexibility in the assignment of these materials. In this figure, different acoustic materials are assigned different colors.

Scene	Volume	Num. Triangles	Num. Materials	Iteration Time
Cathedral	20 686.03 m ³	55665	2	62.0645 s
Twilight	26 759.28 m ³	270	3	85.4107 s
Concert Hall	34 510.03 m ³	5532	5	138.113 s

Table 2: Summary of the various benchmark scenes. The running time of the underlying ARD solver does not depend on the geometry of the scene, but rather the volume of the scene and the number of materials.

provide three locations in which the acoustic properties of the scene are important. Cathedral provides a large, open area that can be particularly reverberant even at higher absorption values. Twilight has a unique architecture including parts of the structure that are open-air. Finally, the Concert Hall model allows us to explore the acoustic characteristics of music hall design. The geometry and material segmentation of each scene is shown in Figure 5. These scenes range in volume from the smallest (Cathedral), which is 20 000 m³ to the largest (Concert Hall), which is 35 000 m³. Table 2 shows a summary of the scenes, their respective volumes, their geometric complexity, the number of material segments used, and the iteration time for a single optimization step.

6.1.1. Acoustic Materials

The acoustic materials in our scene are determined as a result of our optimization process (see Figure 2). Acoustic materials determine the absorption coefficients of different walls or other geometric elements in our scene. These coefficients were constrained to a range of 0.1 to 0.7 (where 0 is fully reflective and 1 is fully absorptive), which limits absorption to a more realistic range.

Some typical absorption values for our range of materials can vary from painted concrete at a coefficient of about 0.1 to materials representing an audience in upholstered seating with a coefficient of about 0.6.

6.1.2. Acoustic Metrics and the Impulse Response

We determine the values of various acoustic metrics by measuring the impulse response at a particular listener location. The impulse response is a measure of sound pressure over time when the room is excited by a sound impulse at a specific source location. We used a Gaussian derivative impulse, which has a zero DC-component (which insures the correctness and numerical stability of the solver). The length of the IR is dependent on how much sound energy remains in the scene.

We measured two different acoustic measurements: Strength (G) and Clarity (C80). These acoustic properties were specifically chosen for their importance in concert hall design [4, 29]. All of these acoustic properties were derived from the impulse response. Acoustic Strength is computed from the IR by comparing the logarithmic ratio of energy in the impulse response with the energy of an impulse response computed with an equivalent sound source 10m away in a free-field condition. Generally, Strength indicates the fullness of the sound. Clarity, on the other hand, gives an idea of how clearly the listener can hear the original sound source. Computing clarity is done by comparing the logarithmic ratio of the first 80 ms of sound to that of the remaining portion of the impulse response. This essentially measures direct sound and early reflections and compares it to the less clear reverberant tail of the impulse response. In all cases, we chose target values that were realistic and represented common practice for the design of concert halls. We chose a target value for acoustic Strength between 3 dB and 5 dB [16] and a target Clarity value between -3 dB and 4 dB [1].

6.2. Convergence

Our approach using Automatic Differentiation compares favorably to computing the derivative via a finite differences technique. The finite differences technique is a simple approach that measures the slope of a function using an epsilon x-component rather than an infinitesimal. In our experiments we picked an epsilon of 10^{-5} . Figure 6 shows the number of optimization steps required when computing sound Strength (G), Clarity (C80), and a combined acoustic measurement of both Strength and Clarity on the cathedral scene. We compare the AD approach with the finite difference approach. The AD approach converges faster, and in some cases the finite difference approach does not converge. In this case, it is stuck in a local minimum that is not the optimal result that the AD approach computes. Overall, using AD is advantageous since it converges to the correct result in addition to converging faster.

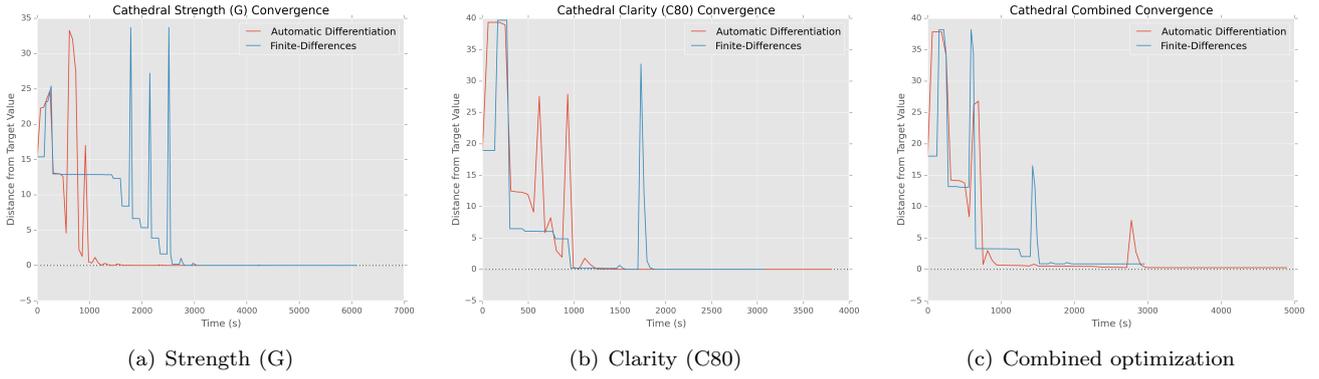


Figure 6: Comparison between different derivation methods: Automatic Differentiation and finite differences at 10^{-5} . We show much faster convergence using Automatic Differentiation for the Strength and Clarity measure. In the combined optimization, the finite differences result finishes early without converging on the target metric. The combined optimization attempts to optimize for both strength and clarity.

Scene	Metric	Target Value	Initial Materials	Final Materials	Final Metric Values
Cathedral	G	4 dB ($w=1$)	0.3, 0.3	0.0455, 0.2333	4 dB
	C80	1 dB ($w=1$)	0.3, 0.3	0.0472, 0.6814	1 dB
	G, C80	4 dB ($w=0.25$), 1 dB ($w=0.75$)	0.3, 0.3	0.0448, 1	2.9242 dB, 1.0001 dB
Twilight	G	4 dB ($w=1$)	0.5, 0.5, 0.5	0.3884, 0.5, 0.4852	3.9998 dB
	C80	1 dB ($w=1$)	0.5, 0.5, 0.5	0.2725, 0.4998, 0.516	0.9995 dB
	G, C80	4 dB ($w=0.25$), 1 dB ($w=0.75$)	0.5, 0.5, 0.5	0.3288, 0.4988, 0.9756	4.0003 dB, 0.999 dB
Concert Hall	G	4 dB ($w=1$)	0.5, 0.5, 0.5, 0.36	0.1, 0.1, 0.7, 0.1	2.9149 dB
	C80	1 dB ($w=1$)	0.5, 0.5, 0.5	0.1468, 0.4177, 0.3593	1 dB
	G, C80	4 dB ($w=0.25$), 1 dB ($w=0.75$)	0.5, 0.5, 0.5, 0.36	0.1, 0.7, 0.7, 0.1	2.4050 dB, 6.0772 dB

Table 3: The various starting and ending material values for optimization. The value w represents the assigned weight for each value. We show close convergence for all scenes in this example. The material range was limited to absorptions of between 0.1 and 0.7.

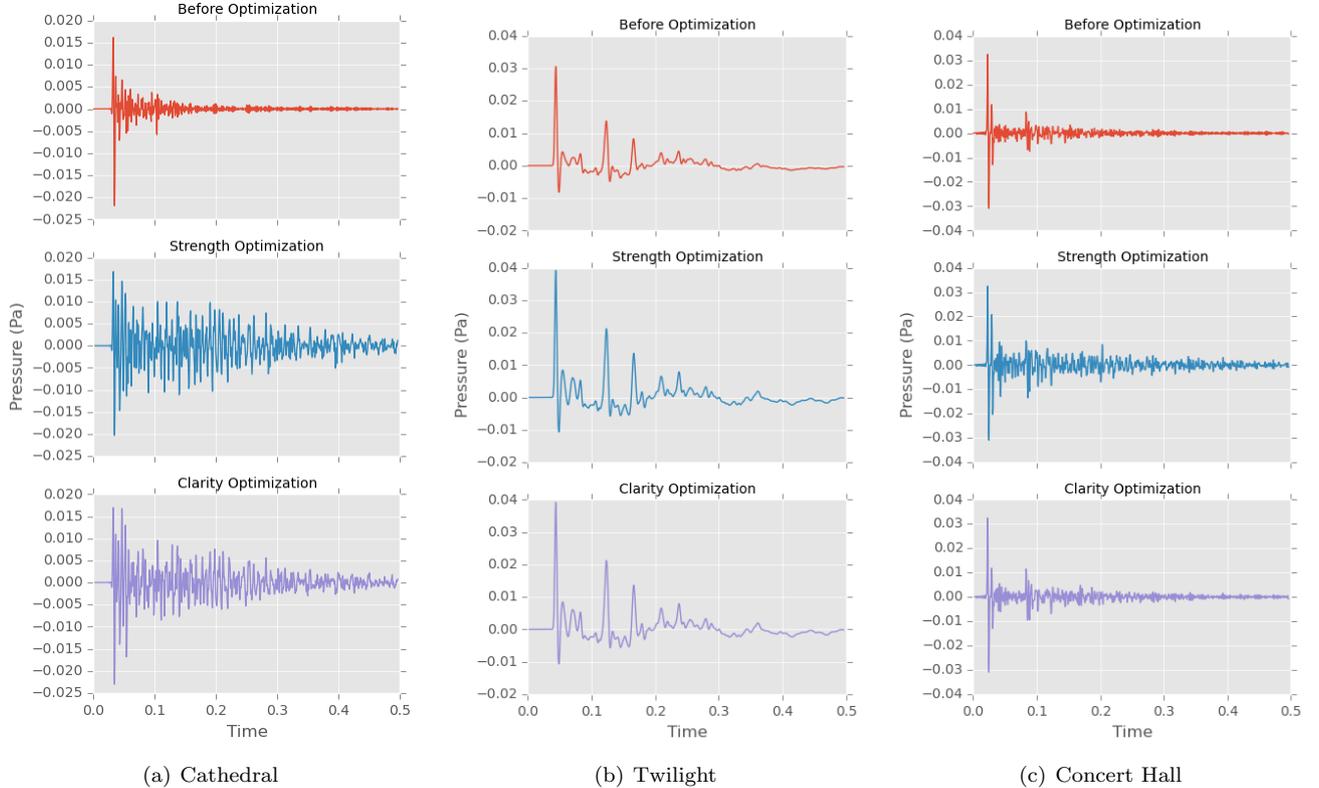


Figure 7: Impulse responses before and after optimization on each scene. Optimizing for Strength increases the overall energy in the response, while optimizing for Clarity increases the ratio energy in the first 80 ms of the impulse response to the remaining energy. These changes can be seen in the figures after optimization is performed.

6.3. Optimization Results

In addition to examining the convergence of our optimization method, we show the effect our optimization has on both the entire pressure field and individual impulse responses at specific listener positions. We ran various experiments one each scene, including one for each acoustic metric and a third experiment for a combined metric that targeted both Strength and Clarity for optimization. Table 3 shows various material values before and after optimization for the various acoustic characteristic metrics and combinations.

Figure 7 shows the impulse responses at each scene before optimization and after optimizing for the different acoustic metrics. Materials with lower absorption values will generally tend to yield impulse responses with more energy after the direct sound impulse (the first maxima in the impulse response). Because acoustic Strength uses the energy of the full impulse response, Strength optimization tended to increase overall energy. On the other hand, Clarity optimization tended to decrease the proportion of the second part of the impulse response compared to the first part, despite overall increasing the reflectivity of the scene materials.

Figure 8, Figure 9, and Figure 10 show the acoustic Strength and Clarity values before and after optimizing for acoustic Strength. These images show how the results for a single source and listener position can be used to drive the acoustic design of a concert hall or other acoustic space.

6.4. Analysis

Our analysis shows that our method is effective in determining the material parameters that effectively yield the desired acoustic characteristics. We are able, in most cases, to complete the optimization in less time than an equivalent finite difference technique for determining the gradient. In the cases in which we do not, our method yields results that are more accurate and closer to the desired acoustic characteristics.

Additionally, our method is general purpose and is capable of working on a multitude of scenes. The only input is the scene with the desired material segment assignments. This is particularly useful for engineers and architects as arbitrarily complex CAD models can be used for optimization.

Finally, our method uses a fast underlying acoustic wave propagation simulation that can give accurate results with much lower computational and memory requirements compared to other standard methods such as FDTD. For example, the clarity optimization on the Cathedral scene took approximately 1 hour to compute. Using FDTD would take around 75 hours to compute [31]. This can make an important difference in turnaround time for architectural acoustic design and development.

7. Conclusion and Future Work

We introduce an efficient wave-based acoustic material design optimizer that is capable of handling multiple material segments and multiple target acoustic properties. We show that using the exact derivatives from Automatic Differentiation helps us converge faster on the target optimization result. Additionally, we take advantage of the performance and memory efficiency of the ARD solver compared to other standard acoustic wave solvers. Finally, we show how our system can be used in the application of designing concert halls or other acoustic spaces.

In the future we would like to explore methods of applying discrete optimization techniques to the acoustic material optimization problem. While our method can take advantage of the sensitivity of ARD to drive continuous optimization, some of the materials produced may not be physically realistic materials. These continuous values can be discretized into material categories (for example concrete bricks have an absorption between (0.01 and 0.02)). However, discrete optimization approaches could take as input a library of acoustic materials that must be used rather than a continuous curve of absorption values. A further advantage to this approach could be the incorporation of other constraints on the optimization, including material cost or the structural feasibility of using a particular material in a specific location.

Acknowledgements

This research was supported in part by NSF award 1456299. Thanks to Sandia National Laboratories for the Sacado project and the maintainers of the DLib library.

- [1] Glen Ballou. *Handbook for sound engineers*. Taylor & Francis, 2013.
- [2] Michael Bartholomew-Biggs, Steven Brown, Bruce Christianson, and Laurence Dixon. Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124(1):171–190, 2000.
- [3] Alban Bassuet, Dave Rife, and Luca Dellatorre. Computational and optimization design in geometric acoustics. *Building Acoustics*, 21(1):75–86, 2014.
- [4] Leo Beranek. The sound strength parameter g and its importance in evaluating and planning the acoustics of halls for music. *The Journal of the Acoustical Society of America*, 129(5):3020–3026, 2011.
- [5] Berglund Birgitta, Lindvall Thomas, and HS Dietrich. Guidelines for community noise. *World Health Organization, Geneva*, page 21, 1999.
- [6] ARLINE L Bronzaft. A quieter school: An enriched learning environment. *Quiet Classrooms*. Retrieved on, 15, 2012.
- [7] John E Dennis, Jr and Jorge J Moré. Quasi-newton methods, motivation and theory. *SIAM review*, 19(1):46–89, 1977.
- [8] Adrian Doicu. Acoustic and electromagnetic scattering analysis using discrete sources. *Acoustic and Electromagnetic Scattering Analysis Using Discrete Sources*, ISBN: 978-0-12-219740-6, p. ix–xi., 1, 2000.
- [9] JB Du, XK Song, and Niels Olhoff. Topological design of acoustic structure based on the bem-fem format and the mixed formulation. In *Proceeding of the 9th world congress on structural and multidisciplinary optimization, Shizuoka, Japan*, 2011.

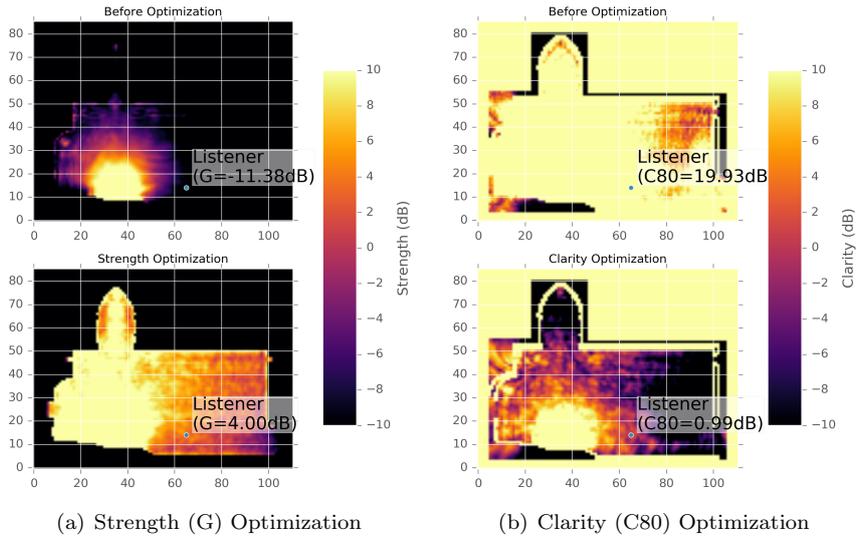


Figure 8: Cathedral field slices before and after optimization. Optimizing the acoustic materials causes a change in the full acoustic field. At the sample listener position we use for the impulse responses in Figure 7, we show the desired target acoustic value after optimization.

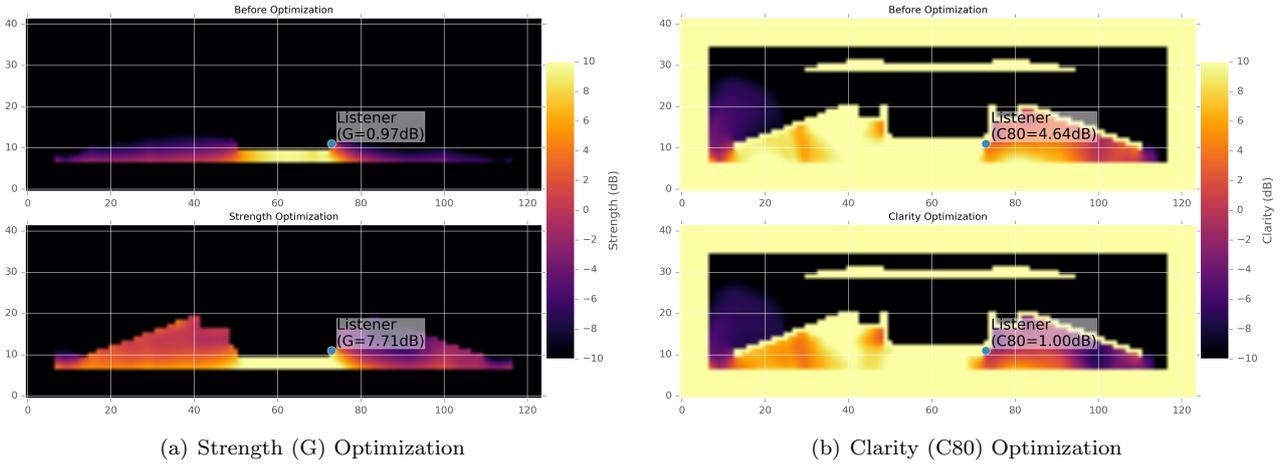


Figure 9: Twilight field slices before and after optimization. In this scene partly open to air, we can see how the strength and clarity values outside of the main structure are affected by the optimization of acoustic materials.

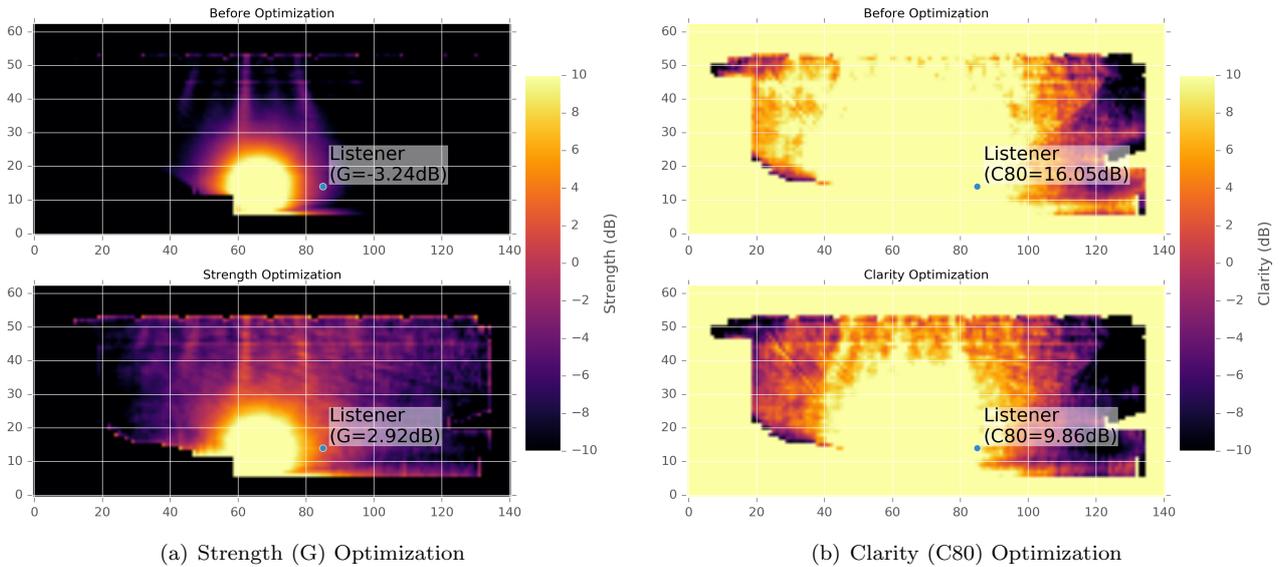


Figure 10: Concert Hall field slices before and after optimization. The listener position in this case is further away from the source and mostly affected by direct sound, yielding a high clarity but lower strength.

- [10] Maria B Dühring, Jakob S Jensen, and Ole Sigmund. Acoustic design by topology optimization. *Journal of sound and vibration*, 317(3):557–575, 2008.
- [11] Sergio E Floodya and Rodolfo B Venegasb. Shape optimization in rectangular rooms for a correct modal distribution at low frequencies based on psychoacoustical model. *Mecánica Computacional*, 26:81–95, 2007.
- [12] Kazuko Fuchi and Hae Chang Gea. Room acoustic optimization with variable thickness columns.
- [13] Klaus Genuit. The sound quality of vehicle interior noise: a challenge for the nvh-engineers. *International Journal of Vehicle Noise and Vibration*, 1(1):158–168, 2004.
- [14] Klaus Genuit and Wade R Bray. A virtual car: Prediction of sound and vibration in an interactive simulation environment. Technical report, SAE Technical Paper, 2001.
- [15] Nail A Gumerov and Ramani Duraiswami. A broadband fast multipole accelerated boundary element method for the three dimensional helmholtz equation. *The Journal of the Acoustical Society of America*, 125(1):191–205, 2009.
- [16] Jerald R Hyde and Henrik Möller. Sound strength in small halls. *Proceedings of the Institute of Acoustics*, 28(PT 2), 2006.
- [17] ISO. Acoustics – measurement of room acoustic parameters. ISO 3882, International Organization for Standardization, Geneva, Switzerland, 2009.
- [18] J. Johnescu. Modeling aircraft noise. *Occupational Health and Safety*, 72(7):106–108, 2003.
- [19] Andy Keane and Prasanth Nair. *Computational approaches for aerospace design: the pursuit of excellence*. John Wiley & Sons, 2005.
- [20] Asbjørn Krokstad, S Strom, and Svein Sørsdal. Calculating the acoustical room response by the use of a ray tracing technique. *Journal of Sound and Vibration*, 8(1):118–125, 1968.
- [21] Heinrich Kuttruff. *Room acoustics*. CRC Press, 2009.
- [22] Adrian S Lewis and Michael L Overton. Nonsmooth optimization via bfgs. *Submitted to SIAM J. Optimiz*, 2009.
- [23] Antonio Maressa, Bert Plumey, Stijn Donders, and Wim Desmet. Nvh optimization methodologies based on bead modification analysis in vehicle body design. In *Proceedings of the international conference on noise and vibration engineering ISMA*, pages 4319–36. Citeseer, 2010.
- [24] Joaquim RRA Martins and Andrew B Lambe. Multidisciplinary design optimization: a survey of architectures. *AIAA journal*, 51(9):2049–2075, 2013.
- [25] Susan Mazer. Curing the noise epidemic. *The Journal of the Acoustical Society of America*, 118(3):1956–1956, 2005.
- [26] Ravish Mehra, Nikunj Raghuvanshi, Lakulish Antani, Anish Chandak, Sean Curtis, and Dinesh Manocha. Wave-based sound propagation in large open scenes using an equivalent source formulation. *ACM Transactions on Graphics (TOG)*, 32(2):19, 2013.
- [27] Ravish Mehra, Nikunj Raghuvanshi, Lauri Savioja, Ming C Lin, and Dinesh Manocha. An efficient gpu-based time domain solver for the acoustic wave equation. *Applied Acoustics*, 73(2):83–94, 2012.
- [28] Tomas Mendez. *Computational Search in Architectural Design*. PhD thesis, Politecnico di Torino, 2014.
- [29] Andrzej Miśkiewicz, Tomira Rogala, Teresa Rościszewska, Tomasz Rudzki, and Tadeusz Fidecki. Concert hall sound clarity: A comparison of auditory judgments and objective measures. *Archives of Acoustics*, 37(1):41–46, 2012.
- [30] Michael Monks, Byong Mok Oh, and Julie Dorsey. Audio optimization: Goal-based acoustic design. *Computer Graphics and Applications, IEEE*, 20(3):76–90, 2000.
- [31] Nicolas Morales, Ravish Mehra, and Dinesh Manocha. A parallel time-domain wave simulator based on rectangular decomposition for distributed memory architectures. *Applied Acoustics*, 97:104–114, 2015.
- [32] Arup Kumar Nandy and CS Jog. Optimization of vibrating structures to reduce radiated noise. *Structural and Multidisciplinary Optimization*, 45(5):717–728, 2012.
- [33] Eric Phipps and Roger Pawlowski. Efficient expression templates for operator overloading-based automatic differentiation. In *Recent Advances in Algorithmic Differentiation*, pages 309–319. Springer, 2012.
- [34] Nikunj Raghuvanshi, Rahul Narain, and Ming C Lin. Efficient and accurate sound propagation using adaptive rectangular decomposition. *Visualization and Computer Graphics, IEEE Transactions on*, 15(5):789–801, 2009.
- [35] Philip W Robinson, Samuel Siltanen, Tapio Lokki, and Lauri Savioja. Concert hall geometry optimization with parametric modeling tools and wave-based acoustic simulations. *Building Acoustics*, 21(1):55–64, 2014.
- [36] Kai Saksela, Jonathan Botts, and Lauri Savioja. Optimization of absorption placement using geometrical acoustic models and least squares. *The Journal of the Acoustical Society of America*, 137(4):EL274–EL280, 2015.
- [37] Carl Schissler, Ravish Mehra, and Dinesh Manocha. High-order diffraction and diffuse reflections for interactive sound propagation in large environments. *ACM Transactions on Graphics (SIGGRAPH 2014)*, 33(4):39, 2014.
- [38] Lei Shu, Michael Yu Wang, and Zhengdong Ma. Level set based topology optimization of vibrating structures for coupled acoustic–structural dynamics. *Computers & Structures*, 132:34–42, 2014.
- [39] Jaroslaw Sobieszczanski-Sobieski. Sensitivity of complex, internally coupled systems. *AIAA journal*, 28(1):153–160, 1990.
- [40] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Artech House Publishers, 3 edition, 2005.
- [41] Lonny L Thompson. A review of finite-element methods for time-harmonic acoustics. *The Journal of the Acoustical Society of America*, 119(3):1315–1330, 2006.
- [42] Jin-Yeu Tsou, Benny Chow, and Selina Lam. Performance-based simulation for the planning and design of hyper-dense urban habitation. *Automation in construction*, 12(5):521–526, 2003.
- [43] Anna Turnage. Reducing aircraft noise with computer graphics. *Computer Graphics and Applications, IEEE*, 22(3):16–21, 2002.
- [44] Michael Vorländer. Simulation of the transient and steady-state sound propagation in rooms using a new combined ray-tracing/image-source algorithm. *The Journal of the Acoustical Society of America*, 86(1):172–178, 1989.
- [45] Gil Ho Yoon, Jakob Søndergaard Jensen, and Ole Sigmund. Topology optimization of acoustic–structure interaction problems using a mixed finite element formulation. *International journal for numerical methods in engineering*, 70(9):1049–1075, 2007.