Audioptimization: Goal–Based Acoustic Design

by

Michael Christopher Monks

M.S. Computer Graphics
Cornell University, 1993

B.S. Computer and Information Science
University of Massachusetts at Amherst, 1987

Submitted to the Department of Architecture
in Partial Fulfillment of the Requirements for the
Degree of

Doctor of Philosophy in the Field of Architecture: Design and Computation
at the

Massachusetts Institute of Technology

June 1999

Author: . . . .
Department of Architecture
April 30, 1999

Certified by: . . . .
Julie Dorsey
Associate Professor of Design and Computational and Building Technology
and Computer Science and Engineering
Thesis Advisor

Accepted by: . . . .
Stanford Anderson
Professor of History and Architecture
Chairman, Department Committee on Graduate Students

Reader: . . . .

Leonard McMillan

Assistant Professor of Computer Science

Reader: . . . .

Carl Rosenberg

Lecturer of Architecture

Audioptimization: Goal–Based Acoustic Design

by

Michael Christopher Monks

Submitted to the Department of Architecture
on April 30, 1999, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in the Field of Architecture: Design and Computation

Abstract

Acoustic design is a difficult problem, because the human perception of sound depends on such things as decibel level, direction of propagation, and attenuation over time, none of which are tangible or visible. The advent of computer simulation and visualization techniques for acoustic design and analysis has yielded a variety of approaches for modeling acoustic performance. However, current computer-aided design and simulation tools suffer from two major drawbacks. First, obtaining the desired acoustic effects may require a long, tedious sequence of modeling and/or simulation steps. Second, current techniques for modeling the propagation of sound in an environment are prohibitively slow and do not support interactive design.

This thesis presents a new approach to computer-aided acoustic design. It is based on the inverse problem of determining material and geometric settings for an environment from a description of the desired performance. The user interactively indicates a range of acceptable material and geometric modifications for an auditorium or similar space, and specifies acoustic goals in space and time by choosing target values for a set of acoustic measures. Given this set of goals and constraints, the system performs an optimization of surface material and geometric parameters using a combination of simulated annealing and steepest descent techniques. Visualization tools extract and present the simulated sound field for points sampled in space and time. The user manipulates the visualizations to create an intuitive expression of acoustic design goals.

We achieve interactive rates for surface material modifications by preprocessing the geometric component of the simulation, and accelerate geometric modifications to the auditorium by trading accuracy for speed through a number of interactive controls.

We describe an interactive system that allows flexible input and display of the solution and report results for several performance spaces.

Thesis Supervisor: Julie Dorsey
Title: Associate Professor of Design and Computational and Building Technology and Computer Science and Engineering

Acknowledgments

As with any endeavor of this magnitude, many people have contributed to its achievement. By far the most important influence on this work comes, not surprisingly, from my wife Marietta. Were it not for her continuous encouragement and support, this effort would have been abandoned well short of its fruition. My son Andrew James has also provided me with the perspective and outlook that enabled me to press on, without being consumed by the gravity of the project.

Aside from my life support, of the players that filled the most prominent roles in the work itself, my advisor and friend Julie Dorsey assumed the lead. Without her creativity, enthusiasm, and persuasion my involvement in this project would not have begun. Julie generated the driving force to keep the project vital when it appeared to be headed for difficulties. She never lost sight of what the project could be, while I was focused on what it was.

The project itself is composed of joint work done by my research partner Byong Mok Oh and myself. Mok and I shared responsibilities for all aspects of the original work, although like any good team, we each deferred to the other's strengths when disagreements arose. I would like to thank my additional committee members, Carl Rosenberg and Leonard McMillan, for their guidance through numerous helpful suggestions and remarks.

My thanks to Alan Edelman and Steven Gortler for sharing their insights in optimization, graphics, and overall presentation of early drafts of papers related to this work. I am also indebted to a variety of people from the acoustics community. First, thanks to Leo Beranek, for inspirational discussions. I am also grateful to Amar Bose, Ken Jacob and Kurt Wagner for sharing their insights on the project. Thanks goes to Mike Barron, John Walsh and Mendel Kleiner for their help gathering background information.

Thanks to the UROPs who participated in this project, including John Tsuchiha, Max Chen, and Natalie Garner, for modeling Kresge Auditorium and Jones Halls, writing numerous software modules, and experimenting with the system.

Finally, let me acknowledge the lab mates who helped create an enjoyable environment in which to work. Among the most responsible are Mok, Sami Shalabi, Jeff Feldgoise, Osama Tolba, Kavita Bala, Justin Legakis, Hans Pedersen, Rebecca Xiong, and Jon Levene. Thanks people.

This work was supported by an Alfred P. Sloan Research Fellowship (BR-3659) and by a grant from the National Science Foundation (IRI-9596074). The Oakridge Bible Chapel model was provided by Kurt Wagner of the Bose Corporation.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Acoustic design is a difficult process, because the human perception of sound depends on such things as decibel level, direction of propagation, and attenuation over time, none of which are tangible or visible. This makes the acoustic performance of a hall very difficult to anticipate. Furthermore, the design process is marked by many complex, often conflicting goals and constraints. For instance, financial concerns might dictate a larger hall with increased seating capacity, which can have negative effects on the hall's acoustics, such as excessive reverberation and noticeable gaps between direct and reverberant sound; fan shaped halls bring the audience closer to the stage than other configurations, but they may fail to make the listener feel surrounded by the sound; the application of highly absorbent materials may reduce disturbing echoes, but they may also deaden the hall. In many renovations, budgetary, aesthetic, or physical impediments limit modifications, compounding the difficulties confronting the designer. In addition, a hall might need to accommodate a wide range of performances, from lectures to symphonic music, each with different acoustic requirements. In short, a concert hall's acoustics depends on the designer's ability to balance many factors.

In 1922, the renowned acoustic researcher, W. C. Sabine, had the following to say about the acoustic design task:

> The problem is necessarily complex, and each room presents many conditions, each of which contributes to the result in a greater or less degree according to circumstances. To take justly into account these varied conditions, the solution of the problem should be quantitative, not merely qualitative; and to reach its highest usefulness it should be such that its application can precede, not follow, the construction of the building. [45]

Various tools exist today to assist designers with the design process. Traditionally, designers have built physical scale models and tested them visually and acoustically. For example, by coating the interiors of the models with reflective material and then shining lasers from various source positions, they try to assess the sight and sound lines of the audience in a hall. They also might attempt to measure acoustical qualities of a proposed environment by conducting acoustic tests on a model using sources and receivers scaled in both frequency and size. Even water models are used sometimes to visualize the acoustic wave propagation in a design. These traditional methods have proven to be inflexible, costly, and time consuming to implement, and are particularly troublesome to modify as the design evolves. They are most effectively used to verify the performance of a completed design, rather than to aid in the design process.

The advent of computer simulation and visualization techniques for acoustic design and analysis has yielded a variety of approaches for modeling acoustic performance [12, 22, 34, 39, 13]. While simulation and visualization techniques can be useful in predicting acoustic performance, they fail to enhance the design process itself, which still involves a burdensome iterative process of trial and error. Today's CAD systems for acoustic design are based almost exclusively on *direct* methods—those that compute a solution from a complete description of an environment and relevant parameters. While these systems can be extremely useful in evaluating the performance of a given 3D environment, they involve a tedious specify-simulate-evaluate loop in which the user is responsible for specifying input parameters and for evaluating the results; the computer is responsible only for computing and displaying the results of these simulations. Because the simulation is a costly part of the loop, it is difficult for a designer to explore the space of possible designs and to achieve specific, desired results.

An alternative approach to design is to consider the *inverse problem*—that is, to allow the user to create a target and have the algorithm work backward to establish various parameters. In this division of labor, the user is now responsible for specifying objectives to be achieved; the computer is responsible for searching the design space, i.e. for selecting parameters optimally with respect to user-supplied objectives. Several lighting design and rendering systems have employed inverse design. For example, the user can specify the location of highlights and shadows [42], pixel intensities or surface radiance values [47, 43], or subjective impressions of illumination [28]; the computer then attempts to determine lighting or reflectance parameters that best match the given objectives using optimization techniques. Because sound is considerably more complex than light, an inverse approach appears to have even more potential in assisting acoustic designers.

In this thesis, I present an *inverse*, interactive acoustic design system. With this approach, the designer specifies goals for acoustic performance in space and time via high level acoustic qualities, such as "decay time" and "sound level." Our system allows the designer to constrain changes to the environment by specifying the range of allowable material as well as geometric modifications for surfaces in the hall. Acoustic targets may be suitable for one type of performance, or may reflect multiple uses. With this information, the system performs a constrained optimization of surface material and geometric parameters for a subset of elements in the environment, and returns the hall configuration that best matches performance targets.

Our audioptimization design system has the following components: an acoustic evaluation module that combines acoustic measures calculated from sound field data to produce a rating for the hall configuration; a visualization toolkit that facilitates an intuitive assessment of the complex time-dependent nature of sound, and provides an interactive means to express desired acoustic performance; design space specification editors that are used to indicate the allowable range of material and geometry modifications for the hall; an optimization module that searches for the best hall configuration in the design space using both simulated annealing for global searching and steepest descent for local searching; and a more geometrically accurate acoustic simulation algorithm that quickly calculates the sound field produced by a given hall configuration.

This system helps a designer produce an architectural configuration that achieves a desired acoustic performance. For a new building, the system may suggest optimal configurations that would not otherwise be considered; for a hall with modifiable components or for a renovation project, it may assist in optimizing an existing configuration. By using optimization routines within an interactive application, our system reveals complex acoustic properties and steers the design process toward the designer's goals.

1.1 Thesis Overview

The remainder of this thesis is organized as follows. Chapter 2 provides background in room acoustics. Chapter 3 presents previous research in acoustic simulation algorithms. I survey sound characterization measures in Chapter 4 and present visualization tools used to display these acoustic qualities. Chapter 5 details our new simulation algorithm that addresses both computation speed and geometric accuracy. Chapter 6 defines the new components introduced by the inverse design approach—the specification of the design space, the definition of the objective function, and the optimization strategy used to search the design space for the best configuration. I describe the implementation details of these new components in the context of our audioptimization design system in Chapter 7. I demonstrate the audioptimization system through several case studies of actual buildings in Chapter 8.

Chapter 2

Acoustics

This chapter presents background on the phenomenon of sound, from both a physical and psychological viewpoint. In addition, room acoustics is introduced, and a short survey of position independent acoustic measures is presented. Given this background information, the limitations of the approximations and simplifying assumptions made by the acoustic evaluation and simulation algorithms presented in the following chapters can be better understood.

2.1 Physics of Sound

Sound is the result of a disturbance in the ambient pressure, P_0, of particles within an elastic medium that can be heard by a human observer or an instrument [8]. Since my focus is room acoustics, I will be addressing sinusoidal sources propagating sound waves through air. As a sound wave travels, its energy is transferred between air particles via collisions. The result is a longitudinal pressure wave composed of alternating regions of compression and rarefaction propagating away from the source. Although the wave may travel a great distance, the motion of particles remains local, as particles oscillate about their ambient positions relative to the wave propagation direction. While each particle will undergo the same motion, the motion will be delayed, or phase shifted, at different points along the wave. Particle

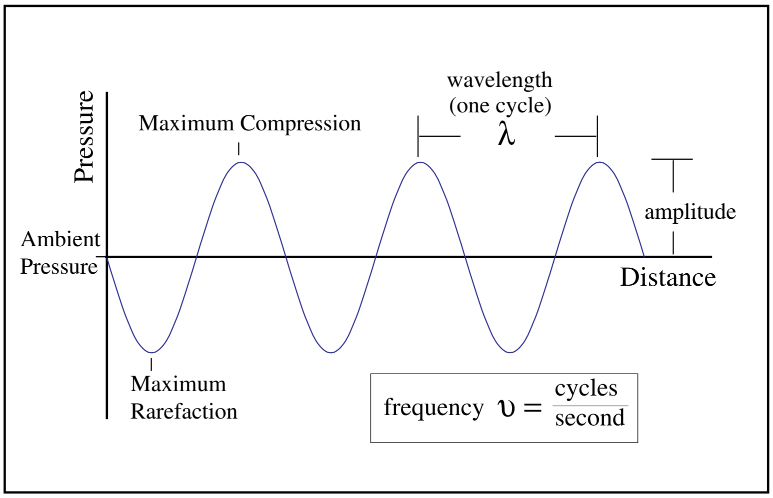

Figure 2-1: Wave terminology.

velocity $u(t)$ and sound pressure $p(t)$ are orthogonal, cyclic functions such that when particle density is greatest, particle velocity is at its ambient level, and when particle velocity is greatest, particle density is at its ambient level [8].

Sound waves share the same characterization as waves of other types, summarized briefly as follows [20]. The length of one cycle of a sound wave—the distance between corresponding crest points, for example—is its *wavelength*, λ, shown in Figure 2-1. The time it takes for the wave to travel one wavelength, is its *period*, T. Its *frequency*, ν, is the number of cycles generated by the sound source in one second, measured in Hertz (Hz). The speed of sound c is independent of λ, but depends on environmental conditions such as air temperature, humidity, and air motion gradients. I assume static conditions for all acoustic modeling presented in this thesis, setting the speed of sound to 345 meters per second.

When the displacement of air by a disturbance takes place fast enough, that is, so fast that the heat resulting from compression does not have time to dissipate between wave cycles, the phenomenon is called *adiabatic* compression, and produces sound waves [8]. The resulting pressure change is related to the change in volume given by the equation $PV^\gamma = constant$, where the gas constant γ is 1.4 for air. Under these adiabatic conditions, the relationship between pressure p, radius r, and time t for spherical propagation can be expressed by the wave equation

$$\frac{\partial^2 p}{\partial r^2} + \frac{2}{r}\frac{\partial p}{\partial r} = \frac{1}{c^2}\frac{\partial^2 p}{\partial t^2}, \tag{2.1}$$

with a solution of the form

$$p(r,t) = \sqrt{2}\, p_r\, e^{j\,\omega\,t}, \tag{2.2}$$

where p_r is the complex root mean square (rms) pressure at a distance r from the source, and j is defined as $\sqrt{-1}$ [8]. This solution yields the (rms) pressure, or effective sound pressure, which monotonically decreases through time, not the instantaneous pressure, which is an oscillating function.

Intensity, I, is defined as the average rate at which power flows in a given direction through a unit area perpendicular to that direction. The relation between intensity and effective sound pressure is given by

$$I = \frac{p^2}{\rho_0\, c}, \tag{2.3}$$

where ρ_0 is the density of air and c is the speed of sound, and the product $\rho_0\, c$ is the characteristic impedance of the medium [8]. For a spherical source, intensity in the radial direction can also be expressed in terms of source power W, neglecting attenuation due to intervening media, as follows:

$$I = \frac{W}{4\,\pi\, r^2}. \tag{2.4}$$

The audible range of sound for the typical human observer encompasses frequencies between 20 Hz and 20 kHz [20]. The perceptible range of intensity covers fourteen orders of magnitude [18]. Instead of working with intensity values directly, it is more convenient to convert these values to various log scale measures. These measures typically use decibel (dB) units, where a decibel is ten times the log of a ratio of energies. Intensity Level (IL) is given in decibels and is defined as

$$IL = 10 \log \frac{I}{I_{ref}}, \tag{2.5}$$

where I_{ref} is $10^{-12} \frac{watt}{m^2}$, the weakest sound intensity perceptible to the typical human listener [8]. A similar measure, Sound Pressure Level (SPL) is also given in decibels and defined as

$$SPL = 20 \log \frac{p}{p_{ref}}, \tag{2.6}$$

where p_{ref} is $2 * 10^{-5} \frac{newton}{m^2}$ [8].

While these measures do not directly map to the human perception of loudness—a doubling of the sound level does not correspond to a doubling of the perception of loudness—a 10 dB increase in level roughly maps to a doubling in loudness. In fact, a doubling of the intensity of sound, or a 3 dB change in level, is just noticeable to the human observer [18]. Our perception of loudness is frequency dependent as well. A sound of 50 dB SPL will fall below our threshold of hearing at 50 Hz, while at 1 kHz it will be easily heard. Equal loudness curves have been established that relate frequency and sound level to loudness [31].

Most sound sources produce complex sounds, composed of a rich spectrum of frequencies. As sound propagates throughout an enclosure, sound waves exhibit frequency dependent behavior. In the following sections we will take a look at the interaction of sound with surfaces in an enclosure, with intervening media, and with other sound waves, and discuss the role that frequency plays in these situations.

2.1.1 Sound Attenuation

Many aspects of our perception of sound within an enclosure are directly related to its decay through time. As it propagates, a spherical wave decays in intensity due to distance, air, and surface attenuation, as follows:

- *Distance Attenuation*: As can be seen from Equation 2.4, intensity decays with the square of distance. As the spherical wavefront propagates and expands, its power is

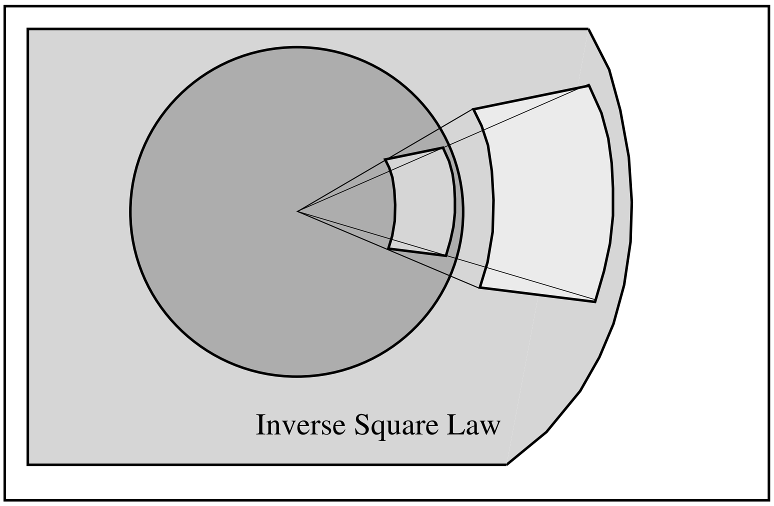

Figure 2-2: Distance attenuation obeys inverse-square law.

distributed over a larger area. Figure 2-2 shows this relationship.

- *Air Attenuation*: Sound intensity also decays by absorption as the sound wave passes through air. While this effect is negligible over short distances under normal atmospheric conditions, it is more significant in large acoustic spaces like concert halls and auditoria. The effect of air absorption on sound intensity is modeled by the following equation

$$I = I_0 * e^{-md}, \tag{2.7}$$

where d is distance traveled in $meters$, and m is the frequency dependent energy attenuation constant in $meters^{-1}$ [8]. The value of m depends on atmospheric conditions such as relative humidity and temperature. Air attenuation is more pronounced for higher frequency sound. For example, after traveling a distance of 345 meters, (one second), the intensity of a 2 kHz wave will decrease roughly 43%, while that of a 500 Hz wave will decrease only 15% under certain atmospheric conditions. In this work, we fix the values for temperature at $68°$F and relative humidity at 50%. Figure 2-3 shows attenuation curves for three frequencies of sound under these atmo-

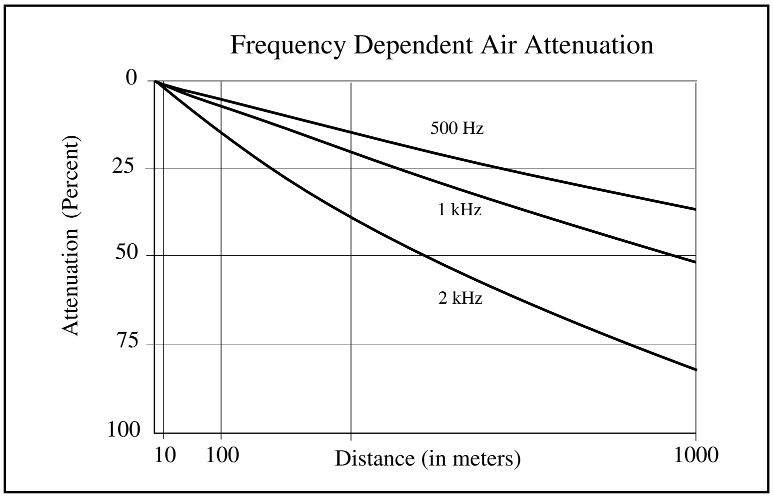

Figure 2-3: Air attenuation curves for three frequencies of sound at room temperature and 50% relative humidity.

spheric conditions.

- *Surface Attenuation*: Surface absorption has the most significant effect on sound decay. Whenever a wave front impinges upon a surface, some portion of its energy is removed from the reflected wave. This sound energy is transferred to the surface by setting the surface into motion, which in turn may initiate new waves on the other side of the surface, accounting for transmission. The extent to which absorption takes place at a surface depends upon many factors, including the materials that comprise the surface, the frequency of the impinging wave front, and the angle of incidence of the wave front.

  The absorption behavior of a surface is commonly characterized by a single value, the *absorption coefficient*, α, which is the ratio between energy that strikes the surface and energy reflected from the surface, averaged over all incident angles [8]. While representing the behavior of sound reflecting from a surface with a single constant term only loosely approximates the actual behavior, it is a useful indicator of the effect that the surface will have on the overall acoustics of an enclosure. The non-uniformity of the reflection curves shown in Figure 2-4 gives an indication of the limitation of

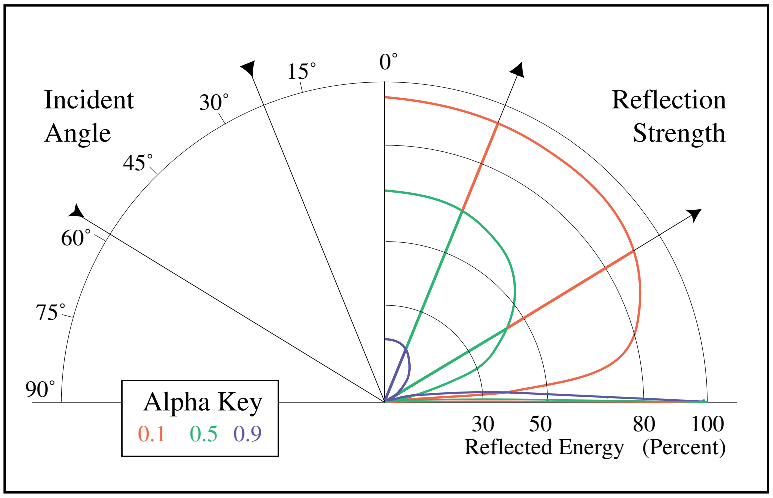

Figure 2-4: Angle dependent surface reflection plots for three materials with average absorption coefficients of 0.9 (blue,) 0.5 (green,) and 0.1 (red.)

this approximation method [35].

2.1.2 Wave Properties

If an obstacle is encountered whose surface is much more broad in dimension than the wavelength of impinging sound, the surface will reflect the sound. The manner in which the sound is reflected depends upon the scale of surface roughness with respect to the wavelength of sound. Smooth surfaces will reflect sound geometrically, such that the reflection angle is equal to the angle of incidence. Surfaces with roughness of comparable or greater scale than the wavelength of sound will diffuse sound [31].

An obstacle whose scale is small compared to the wavelength of an approaching sound wave will have little effect on the wave. The wave will be temporarily disrupted after passing the obstacle, and reform a short distance beyond it. This behavior occurs because the air particles that carry the sound energy are moving in all directions, and will spread the wave energy laterally if unrestricted. This behavior is called *diffraction*, and accounts for related behaviors such as sound turning corners, or fanning out as it passes through openings [20].

If one imagines a volume of air space through which many sound waves pass—from many different directions, covering a spectrum of frequencies—the particles within that air space respond to all sound waves at once. The human listener is able to separate and discern the effects on those particles due to different frequencies, imposing some level of order to the chaos. If waves of the same frequency are traveling in the same direction through that air space, they will either reenforce each other if they are in phase, or interfere destructively if they are out of phase. In the following section I begin to discuss the degree to which various properties of sound are taken into account in the context of room acoustics.

2.2 Room Acoustics

One goal of room acoustics is to predict various characteristics of the sound that would be produced in an acoustic enclosure, given a description of the room and a sound source. All the information required to characterize the room's acoustics is generated by computing the sound field created by such a source within that room. Theoretically the response of the room could be solved exactly using the wave equation, given the myriad boundary conditions as input. Practically, however, the complexities introduced by any but the simplest conditions make this calculation intractable. Additionally, even if practical, such a calculation would yield far more information than necessary to characterize the behavior of sound in the room to the level of detail that is useful [31].

Fortunately, much insight can be gained by considering various degrees of approximation of the sound source and the room, and making simplifying assumptions about the behavior of sound transport and sound-surface interactions. I summarize a few of the most simple, yet useful approximations below.

2.2.1 Reverberation: Sabine's Formula

For many years the only measure of the acoustic quality of a space was its *reverberance*, or the time it took for a sound to become inaudible after the source was terminated.

The acoustic pioneer, W. C. Sabine, was the first to explore and document the effects that the material qualities of various surfaces and objects in the room have on the reverberance of the room. While a researcher at Harvard University, Sabine was given the task of determining why the lecture room of the Fogg Art Museum had such poor acoustics, and to suggest modifications that would solve the problem.

Sabine observed that, while the room was modeled after its neighbor, the acousticly successful Sanders Theatre, the materials of the surfaces were very different, with Sanders faced with wood and the lecture room with tile on plaster. He conducted an experiment in which he measured the reverberation times for various hall configurations. When empty, the hall's reverberance measured 5.62 seconds. Sabine gradually added a number of cushions from Sanders Theatre, arranged throughout the room, taking additional measurements. The reverberation time was reduced by the addition of cushions, reaching a low of 1.14 seconds.

Sabine then tested and catalogued a number of materials and fixtures (including people,) calculating the absorption characteristics of each with respect to that of the seat cushions used previously.

By fitting the data to a smooth curve, Sabine realized that the minor discrepancies would vanish if reverberation was plotted against the total exposed surface area of the cushions, instead of the running length of cushions. From these and other experiments, Sabine arrived at a formula, derived empirically, to predict the reverberation time T of a room given its volume V in cubic feet and the total room absorption a in *sabins* within the room. A sabin is defined as one square foot of perfect absorption. His formula follows:

$$T = 0.05\frac{V}{a}, \tag{2.8}$$

where the number of sabins is found by summation, over all surfaces S in the room, of the product of surface area S_n and the absorption coefficient α_n [18]. Just as a is frequency dependent, so too is T. While this formula does not account for the effects of such factors as the proximity of absorptive materials to the source, for example, Sabine himself noted that "it would be a mistake to suppose that …[the position of absorption within the room] is of no consequence." [45]

2.2.2 Statistical Approximations

The rate of decay of sound density D in a room can be approximated by calculating statistical averages of the room's geometry and material characteristics [8]. The two values needed are the mean time between surface reflections, and the average absorption at reflection. From this data the envelope of decay due to surface absorption is calculated.

The first simplification is introduced by replacing the geometric description of the room with its *mean free path*, d, defined as the average path length that sound can travel between surface reflections. d is approximated by the formula

$$d = \frac{4V}{S}, \tag{2.9}$$

where V is the volume of the room, and S is the total surface area within the room [8]. Given d, the *mean time*, t', between reflections is simply $t' = \frac{d}{c}$, where c is the speed of sound.

Complex absorption effects at each surface are replaced by the *average absorption coefficient* for the entire enclosure, $\bar{\alpha}$, which is the weighted average of surface absorption defined as

$$\bar{\alpha} = \frac{\alpha_1 S_1 + \alpha_2 S_2 + \cdots + \alpha_n S_n}{S}. \tag{2.10}$$

Absorption due to the contents of the room is accounted for by appending their absorption values to the numerator, although their additional surface area is typically neglected in the denominator [8].

Statistically speaking, the sound energy remaining after the first reflection at time t' is given by $D_{t'} = D_0(1 - \bar{\alpha})$, and at time $2t'$ by $D_{2t'} = D_0(1 - \bar{\alpha})^2$. Beranek shows that by converting from a discrete into a continuous formulation, the following function gives a statistical approximation to the energy density remaining at any point in time:

$$D_t = D_0(1 - \bar{\alpha})^{\frac{cS}{4V}t}. \tag{2.11}$$

He solves for the reverberation time T by rearranging this equation to isolate t, and substituting 60 dB for the ratio of D_0 to D_t, arriving at the following formula:

$$T = \frac{60V}{1.085ca'}, \tag{2.12}$$

where a' is termed *metric absorption units*, and is defined as $-S\ln(1-\bar{\alpha})$, given in square meters. The effects of air attenuation may be accounted for as well by replacing a' with a'_{air}, where $a'_{air} = a' + 4mV$, and m is the air attenuation constant discussed above [8].

While these tools give us some indication of the character of the sound field created by a source within an enclosure, they are limited by the simplifying assumptions implicit in their derivations. Energy density is not uniform throughout the enclosure, due in part to the irregularity of geometry and the non-uniform distribution of absorptive material. Further, during the last few decades a number of acoustic measures have been developed, which require more detailed information about the sound field than these calculations provide. The values for many of these measures vary for different positions throughout the hall, and require directional and temporal data along with the intensity of sound for each passing wavefront. Fortunately, a variety of sophisticated simulation algorithms have been developed in recent years. In the following chapter I survey these algorithms and discuss their strengths and weaknesses.

Chapter 3

Survey of Acoustic Simulation Approaches

In order to evaluate the acoustics of a virtual enclosure it is necessary to simulate the sound field it produces and extract the data required to calculate various acoustic measures. There has been a large amount of previous work in acoustic simulation [7, 31]. These approaches can be divided into five general categories: ray tracing [30], statistical methods [22], radiant exchange methods [34, 48, 52], image source methods [12], and beam tracing [21, 34, 19]. There are also a variety of hybrid simulation techniques, which typically approximate the sound field by modeling the early and late sound fields separately and combining the results [22, 34, 39]. I survey the major simulation algorithms below.

3.1 Ray Tracing

The ray tracing method propagates particles along rays away from the source, which reflect from the surfaces they strike. Ray information is recorded by the receivers that these rays encounter within the enclosure. Since the probability is zero that a dimensionless particle will encounter a dimensionless receiver point in space, receivers are represented as volumes instead of points. Because of this approximation, receivers may record hits from rays that could not possibly reach them, as shown in Figure 3-1a. Errors in the direction and arrival time of the sound will also result.

Another shortcoming of this method is the immense number of rays that are necessary in order to insure that all paths between the source and a receiver are represented. As an illustration, consider a receiver point represented with a sphere of radius one meter, located ten meters from the source in any arbitrary direction. At this distance, one must shoot about 600 rays, uniformly distributed, to insure that the receiver sphere is struck. Now consider the case where we are interested in modeling sound propagation through one full second. Given that sound will travel 345 meters in that amount of time, one would need about 100 times as many rays to insure that a receiver at that distance would be hit. At this sampling density, however, our receiver at ten meters would be struck by 100 redundant rays, all representing the same path to the source. The problem is compounded when considering that the projected angle of any given surface, not the receiver sphere, may determine the minimum density of rays. While ray subdivision may address some of these issues, no guarantees can be made that all paths will be found using ray tracing.

3.2 Statistical Methods

By relaxing the restriction that rays impinging upon a surface must reflect geometrically, and exchanging the goal of finding all possible paths between sources and receivers for the goal of capturing the overall character of sound propagation, ray tracing techniques have been imbued with statistical behavior. Diffusion is modeled with this method by allowing rays to reflect from surfaces in randomly selected directions based upon probability distribution functions. Furthermore, instead of continuing to trace rays as their energy decreases, rays may be terminated at reflection with a probability based on the absorption characteristics of the reflecting surface. Diffraction effects might also be modeled using statistical methods, perhaps by perturbing ray propagation directions mid flight between reflections. While the range of behaviors that can be captured with statistical approaches is quite broad, its appli-

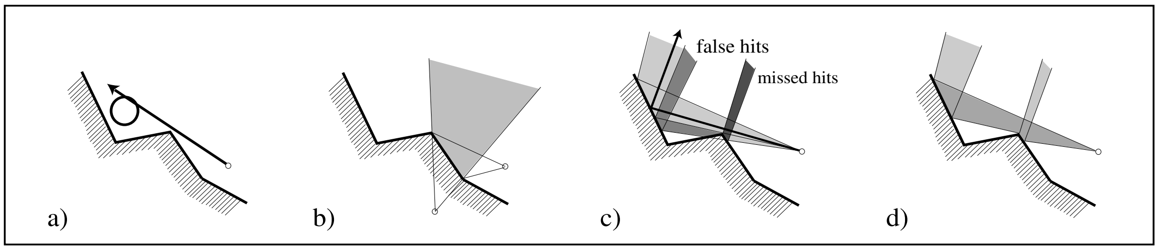

Figure 3-1: a) Ray tracing: positional errors. b) Image-source: limited scope of image source. c) Beam tracing with fixed profiles: false and missed hits. d) General beams: correct hits.

cation might be better suited to modeling later sound, when the direction of propagation is less critical.

3.3 Radiant Exchange Methods

Recently, the radiant exchange methods used for modeling light transport have begun to be applied to acoustic modeling. Assuming diffuse reflection, the percent of energy that leaves one surface and arrives at another is calculated for all pairs of surfaces in a polygonalized environment. Energy is then radiated into the environment, and a steady state solution is calculated. Two issues emerge when applying this technique to sound transport. First, while the arrival time can be ignored for light, it is critical for sound. Second, the diffuse reflection assumption is invalid for modeling most interactions between sound and surfaces. It is particularly troublesome for modeling early sound, where the directional effects are especially important. While it is more acceptable to model late sound than the early sound using diffuse reflection assumptions, more cost effective approaches may suffice for modeling late sound.

3.4 Image Source

The image source method is based on geometric sound transport assumptions. In the ideal case, the approach works as follows. In a room where the surfaces are covered with perfect mirrors, a receiver would be affected by every image of the source, whether direct or reflected, that is visible at the receiver location. No temporal or directional errors would result. While this is the ideal outcome, problems arise in the implementation of the method. Image sources are constructed by mirror reflecting the source across all planar surfaces, as shown in Figure 3-1c. This process recurses, treating each image source as a source. Each resulting image source could potentially influence each receiver. Unfortunately, many image sources constructed in this way are not realizable, and the valid ones are often extremely limited in their scope, as shown in Figure 3-1b. That is, they may only contribute to a fraction of the entire room volume. Recursive validity checking is required to ensure receiver-source visibility. Unless preemptive culling is performed during the construction of image sources, a geometric explosion of image sources results.

3.5 Beam Tracing

A variation on the ray tracing method is beam tracing. Here the rays are characterized by circular cones or polygonal cones of various preset profiles [57]. Receivers are represented as points. As these beams emanate from the source, receivers that are enclosed within the volume of the beam are updated. As they propagate, beams reflect geometrically from surfaces encountered by their central axis. The geometric explosion that characterizes the image source method is avoided here because the number of beams does not grow during propagation. However, the method suffers from various shortcomings. If circular beams are packed such that they touch but do not overlap each other, then gaps result, leaving regions within the room erroneously uneffected by the sound. Conversely, if the beams are allowed to overlap, removing the gaps, then regions are double covered, causing simulation errors [33]. These systematic errors are eliminated for the direct field by using perfectly packed triangular profiles, for example, which attain full spherical coverage of the source. However, the errors reemerge at reflection since beams are not split when they illuminate multiple surfaces, striking an edge or a corner. As a consequence, false hits and missed hits result, as shown in Figure 3-1c [34].

3.6 Hybrid Methods

Many hybrid approaches have emerged that incorporate the best features of these methods. The image source method is best used for modeling early reflections where directional and temporal accuracy is critical. The method may be paired with ray tracing, which is used to establish valid image sources [39, 58]. The image source method is rarely used for later reflections due to the exponential increase in cost. Lewers models late sound with the radiant exchange method [34]. Others use ray tracing, randomizing the reflection direction to attain a diffusing effect [39]. Still others use a statistical approach based on the results of an earlier, or ongoing ray tracing phase [57, 22]. Heinz presents an approach in which surfaces in the enclosure are assigned a wavelength dependent diffusion coefficient, which is used to transfer energy from the incoming ray to the diffuse field [22]. Various approaches are used to combine the early and late response simulations.

3.7 Summary

In the context of acoustic design, it is not necessary that the simulations achieve audio–realistic results. In fact, such a level of detail would detract from the process. While the high computation costs required by radiant exchange methods may be worthwhile for applications requiring high fidelity reproduction, geometric assumptions suffice for our purposes. In Chapter 5 I describe our new acoustic simulation algorithm, which builds upon the strengths of the methods just described, making improvements in both geometric accuracy and computation speed. In addition, I present a more complete set of acoustic performance evaluation criteria, derived from the sound field data, in Chapter 4.

Chapter 4

Acoustic Evaluation

This chapter presents the set of objective measures used by our system to evaluate the acoustic quality of a performance space. I define the measures in the first section and introduce visualizations of the measures in the second section. These visualizations and associated interactive tools give the user of the system an intuitive way to quickly assess acoustic quality.

4.1 Characterization of Sound

Traditionally, reverberation time and other early decay measurements were considered the primary evaluation parameters in acoustic design. However, in recent years, researchers have recognized the inadequacies of using these criteria alone and have introduced a variety of additional measures aimed at characterizing the subjective impression of human listeners [7]. For example, in 1991 Wu applied fuzzy set theory, noting that the subjective response of listeners is often ambiguous [63]. In 1995, Ando proposed another approach showing how to combine a number of orthogonal objective acoustic measures into a single quality rating using his subjective preference test results [5]. In 1996, Beranek built on Ando's work by linearly combining six statistically independent objective acoustic measures into an evaluation function that gives an overall acoustic rating [10]. In this research, we employ Beranek's evaluation approach, known as the *Objective Rating Method (ORM)*.

Below I define the six acoustic measures and introduce visualization techniques used to evaluate them.

Interaural Cross-Correlation Coefficient (IACC). The Interaural Cross-Correlation Coefficient is a binaural measure of the correlation between the signal at the two ears of a listener. It characterizes how surrounded a listener feels by the sound within a hall. If the sound comes from directly in front of or behind the listener, it will arrive at both ears at the same time with complete correlation, producing no stereo effect. If it comes from another direction, the two signals will be out of phase and less correlated, giving the listener the desirable sensation of being enveloped by the sound. I use the following expression to calculate IACC from computer simulated output [5]:

$$\mathrm{IACC} = \Phi_{lr}^{(P)} = \frac{\sum_{p=0}^{P} A_p^2 \Phi_{lr}^{(p)}}{\sqrt{\sum_{p=0}^{P} A_p^2 \Phi_{ll}^{(p)} \sum_{p=0}^{P} A_p^2 \Phi_{rr}^{(p)}}}, \tag{4.1}$$

where $\Phi_{lr}^{(p)}$ is the interaural cross correlation of the pth pulse, $\Phi_{ll}^{(p)}$ and $\Phi_{rr}^{(p)}$ are the autocorrelation functions at the left and right ear, respectively, and A_p is the pressure amplitude of the pth pulse from the set of P pulses. The correlation values depend on the arrival direction of the wave with respect to the listener's orientation. The numerator is greater for highly correlated frontal signals than for less correlated lateral signals. Since the amplitude of sound decreases rapidly as it propagates, the sound waves that arrive the earliest generally have far greater effect on IACC.

Early Decay Time (EDT). The Early Decay Time measures the reverberation or liveliness of the hall. Musicians characterize a hall as "dead" or "live," depending on whether EDT is too low or high. The formal definition of EDT is the time it takes for the level of sound to drop 10 decibels from its initial level, which is then normalized for comparison to traditional measures of reverberation by multiplying the value by six. As Beranek suggests, I determine EDT by averaging the values of EDT for 500 Hz and 1000 Hz sound pulses. The best values of EDT range between 2.0 and 2.3 seconds for concert halls.

Bass Ratio (BR). The Bass Ratio measures how much sound comes from bass, reflecting the persistence of low frequency energy relative to mid frequency energy. It is what musicians refer to as the "warmth" of the sound. BR is defined as:

$$\mathrm{BR} = \frac{\mathrm{RT}_{125} + \mathrm{RT}_{250}}{\mathrm{RT}_{500} + \mathrm{RT}_{1000}}, \tag{4.2}$$

where RT is the frequency dependent reverberation time. RT is the time it takes for the sound level to drop from 5 dB to 35 dB below the initial level, which is then normalized for comparison to traditional measures of reverberation by multiplying by two. For example, for a 100 dB initial sound level, RT would be the time it takes to drop from 95 dB to 65 dB multiplied by the normalizing factor. The ideal value of BR ranges between 1.1 and 1.4 for concert halls.

Strength Factor (G). The Strength Factor measures sound level, approximating a general perception of loudness of the sound in a space. It is defined as follows:

$$\mathrm{G} = 10 \log \left(\frac{\int_0^\infty i(t)dt}{\int_0^\infty i_A(t)dt} \right), \tag{4.3}$$

where t is the time in seconds from the instant the sound pulse is initiated, $i(t)$ is the intensity of a sound wave passing at time t, and $i_A(t)$ is the free field (direct) intensity ten meters from the source. The numerator accumulates energy from the pulse as each propagated wave passes a receiver, until it is completely dissipated; the denominator receives only a single energy contribution. The division cancels the magnitude of the source power from the equation, allowing easy comparison of measured data across different halls. I average the values of G at 500 Hz and 1000 Hz. The preferred values for G range between 4.0 dB and 5.5 dB for concert halls. In general, G is higher at locations closer to the source.

It is instructive to see how the sound level changes through time, as well as location. We perceive a reflected wave front as an echo—perceptibly separable from the initial sound—if it arrives more than 50 msec. after the direct sound and it is substantially stronger than its neighbors. The time distribution of sound also affects our perception of clarity. Two locations in a hall may have the same value of G, but if the energy arrives later with respect to the direct sound for one location than the other, speech will be less intelligible, and music less crisp.

Initial-Time-Delay Gap (TI). The Initial-Time-Delay Gap measures how large the hall sounds, quantifying the perception of intimacy the listener feels in a space. It depends purely on the geometry of the hall, measuring time delay between the arrival of the direct sound and the arrival of the first reflected wave to reach the listener. In order to make comparisons among different halls, only a single value is recorded per hall, measured at a location in the center of the main seating area. It is best if TI does not exceed 20 msec.

Surface Diffusivity Index (SDI). The Surface Diffusivity Index is a measure of the amount of sound diffusion caused by gross surface detail, or macroscopic roughness of surfaces within a hall. SDI is usually determined by inspection, and it correlates to the tonal quality of the sound in a hall. I compute the SDI index for the entire hall by summing the SDI assigned to each surface material, weighted by its area with respect to the total surface area of the hall. SDI can range between 0.0 and 1.0, with larger values representing more diffusion. The preferred value of SDI is 1.0. For example, plaster has a lower index than brick, which has a lower index than corrugated metal.

These six statistically orthogonal acoustic measures form the basis for our analysis and optimization work. While two of the measures, SDI and TI, are single values representing the entire hall, I compute the others by averaging the values sampled at multiple spatial positions, and, in the case of G, multiple points in time. Refer to Appendix A for pseudocode describing the calculation of each measure.

4.2 Visualization of Acoustic Measures

Work has been done in representing sound field data with both visualizations and auralizations. Stettner presented a set of 3D icons to graphically convey the behavior of sound within an enclosure [50]. He also showed the accumulation of sound energy through time by animating pseudo-colors applied to enclosure surfaces. The Bose *Auditioner* system [13] provides auralizations from simulation data at specific listener positions within a modeled hall; these auralizations approximate what the hall might sound like. Our system provides visualizations for both the sound field and a collection of acoustic measures that describe the character of the sound field as it varies in space and time within an environment. These visualizations are used both to analyze the behavior of a given design and to interactively specify desirable performance goals. In this section, I describe these visualizations and associated interactive tools.

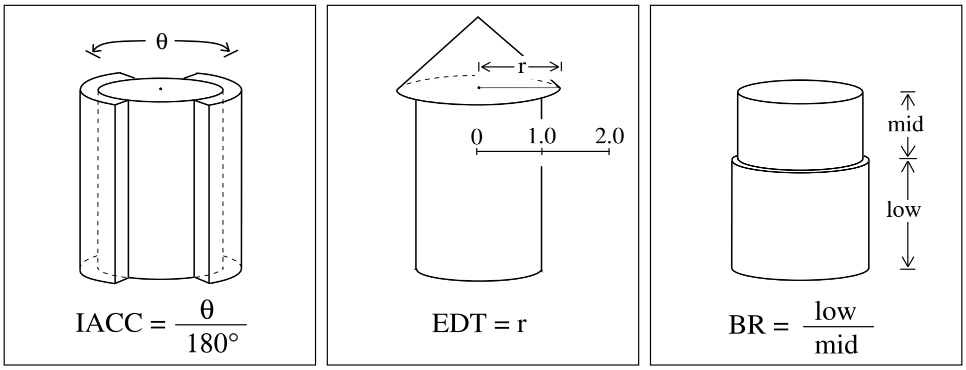

Figure 4-1: Graphical icons representing (from left to right): IACC, EDT, and BR.

The six acoustic measures used to evaluate the quality of sound fall into three categories. The first group includes those measures calculated directly from the configuration of the enclosure, requiring no sound field simulation data. TI and SDI are in this category, and their values are displayed in a text field. Members of the second group share the characteristic that their values differ throughout the enclosure. The measures IACC, EDT and BR are of this type, and are represented using icons, as shown in Figure 4-1. The third type of measure is derived from data that not only contains a spatially varying component, but a temporal component as well. Sound Strength G is of this type.

When designing icons for IACC, EDT and BR, we tried to leverage the intuition of the user whenever possible to help convey meaning. We chose representations that would all be clear from the same view direction. A cylinder is placed at each listener position, or sample point, upon which the icons are placed. A description of the visualization tool used for each measure in the second and third category follows.

- *IACC*: IACC is represented as a shell surrounding the listener icon, which illustrates the degree to which the listener feels surrounded by the sound. The greater the degree of encirclement of the icon by the shell, the more desirable the IACC value.

- *EDT*: EDT is represented graphically as a cone, scaling the radius by decay time and fixing the height. The slope of the cone gives the viewer an intuition for the rate of decay of sound. For a value of 2.0, the cone is twice the width of the listener icon.

- *BR*: Figure 4-1 shows the graphical icon we use for BR, composed of two stacked concentric cylinders of different widths. The top cylinder represents the mid frequency energy and the bottom cylinder represents the low frequency energy. The height of each cylinder represents the relative values in the ratio, with constant combined height. A listener icon representing a desirable BR value of 1.25 would have the top of the lower cylinder just above the halfway mark, as depicted in the figure.

- *G*: The remaining view space real estate is utilized by using color to represent relative scalar values of sound level data, sampled over selected surfaces within the enclosure. The user may choose to view G over all surfaces, as shown in Figure 4-2, or over seating regions only, as in Figure 4-3. In the latter case, we include the ability to view the accumulation of sound level through time (see Figure 4-4), simply by moving a slider. The sampling density over the seating regions is user controlled. This feature gives the user another way to set the balance between the accuracy and speed of the acoustic simulation, as discussed at length in Chapter 5.

I have presented the set of objective acoustic measures used by our goal–based acoustic design system to evaluate the acoustic quality of a performance space, and I have described visualizations used to communicate their values. In the following chapter I will describe the simulation algorithms that generate the data required for the calculation of these acoustic measures.

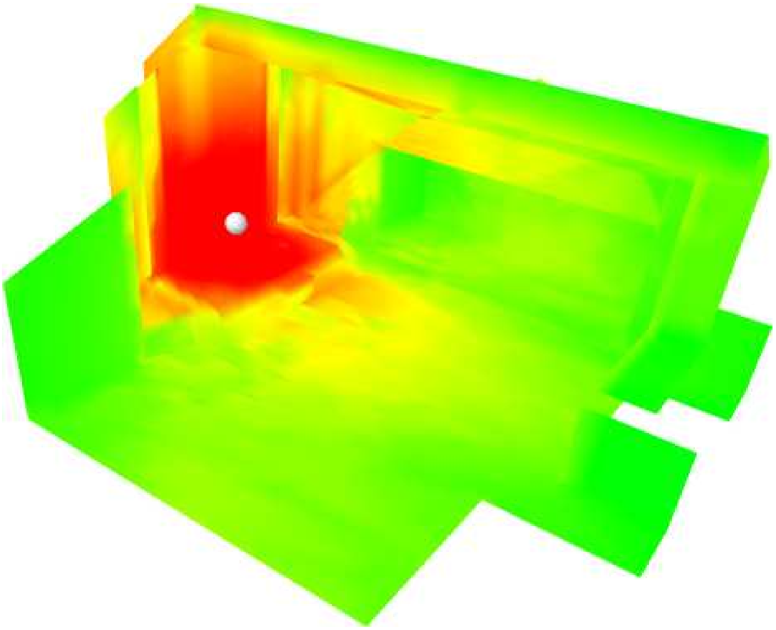

Figure 4-2: Scalar values of sound level data are represented with color over all surfaces of an enclosure.

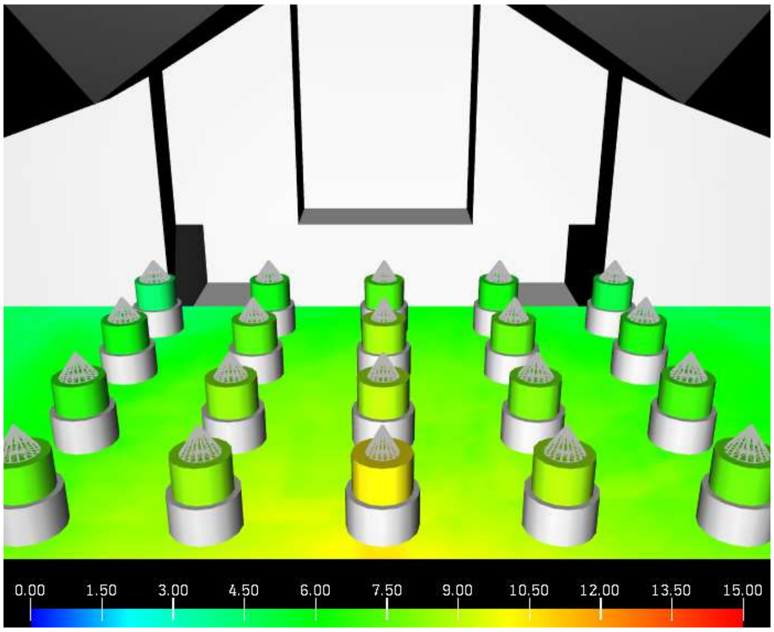

Figure 4-3: Visualization showing scalar values of sound level data represented with color for the seating region along with EDT and BR values, represented with icons at a grid of sample points within an enclosure.

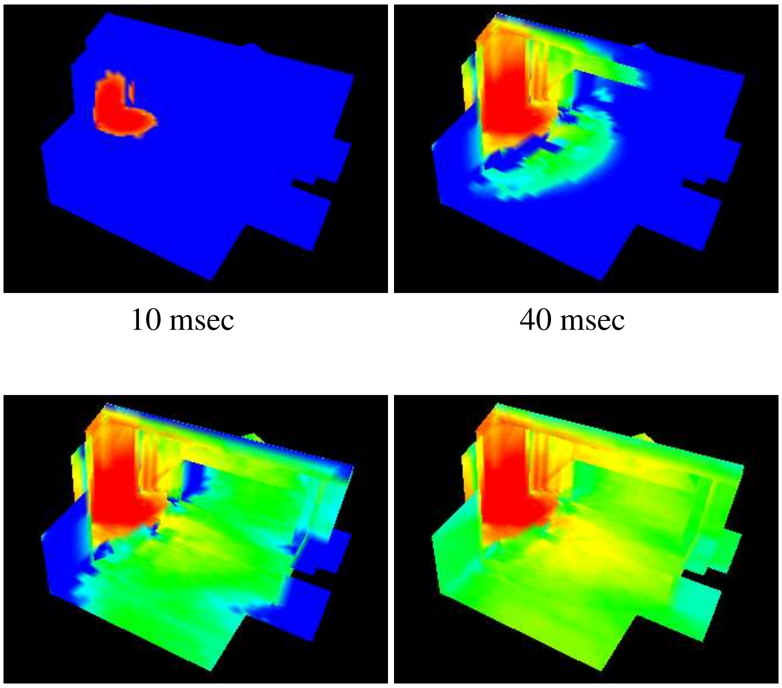

10 msec 40 msec

80 msec 120 msec

Figure 4-4: Color indicates sound strength data at four time steps.

Chapter 5

Acoustic Simulation: A Unified Beam Tracing and Statistical Approach

In order to evaluate the acoustics of a virtual enclosure it is necessary to simulate the sound field it produces and extract the data required to calculate various acoustic measures. In this chapter, I introduce a new hybrid simulation algorithm [37] with two components: a sound beam generator that generalizes previous beam tracing methods for increased geometric accuracy [16], and a statistical approximation of late sound that leverages the beam tracing phase for increased speed. As with previous methods, I make a number of simplifying assumptions about sound transport, and replace the description of the enclosure with a planar representation. In the test models used in this thesis, this representation typically consists of a few dozen to a few hundred polygons.

I employ the *geometric* model of sound transport, which assumes that sound travels in straight lines, and reflects geometrically from surfaces. Wave effects such as diffraction, destructive interference, and other phase related effects are not considered. The angle dependent absorption characteristics of surfaces are not modeled, although frequency dependent characteristics are included.

5.1 Generalized Beams

In beam tracing methods, beam profiles are generally predetermined. Here, beam profiles are dictated strictly by the geometry of the enclosure. The direct field is bounded radially from the source by the portions of each surface that can see the source directly. I refer to these regions as occluders, shown in red in Figure 5-1a, c. The source is mirror-reflected across each occluder as in the image source method, shown in Figure 5-1b, spawning new beams. If we think of a parent beam's occluder as a window into the enclosure for its child beam, then a beam is described as the volume within the enclosure that would be directly illuminated from outside the window by a light located at the point of its image source(see Figure 5-1c. The portions of surfaces illuminated by such a light source compose the beam's mosaic of occluders, which in turn become windows for another generation of beams, as shown in Figure 5-1d, e. The process continues until some termination criteria are reached. The profiles of every generation of beams completely and exactly tile an arbitrary sphere surrounding the source when projected onto it, shown for the first two reflections in Figure 5-2. Refer to Figure 3-1d to see that perfect coverage is attained with generalized beams, unlike fixed profile beam methods.

The consequences of this method are the following. First, a new beam is spawned for each portion of each surface that is visible from the beam source through its window polygon. Since beams are allowed to subdivide at reflection surfaces, no false hits or missed hits can occur. The number of beams is no longer predetermined; hence, no virtual sources are overlooked. Since beam reflection is not approximated using the central axis, the accompanying errors are avoided as well. Generalized beams thus retain all of the benefits—but none of the limitations—of previous methods.

I store the 3D geometric description of each beam in a data structure, along with a 2D representation of its window and mosaic to accelerate point to beam intersection calculations [39]. The reflection history is stored as well, along with accumulated surface attenuation from previous reflections.

In order to determine if a receiver is within a beam, it is not sufficient to simply project

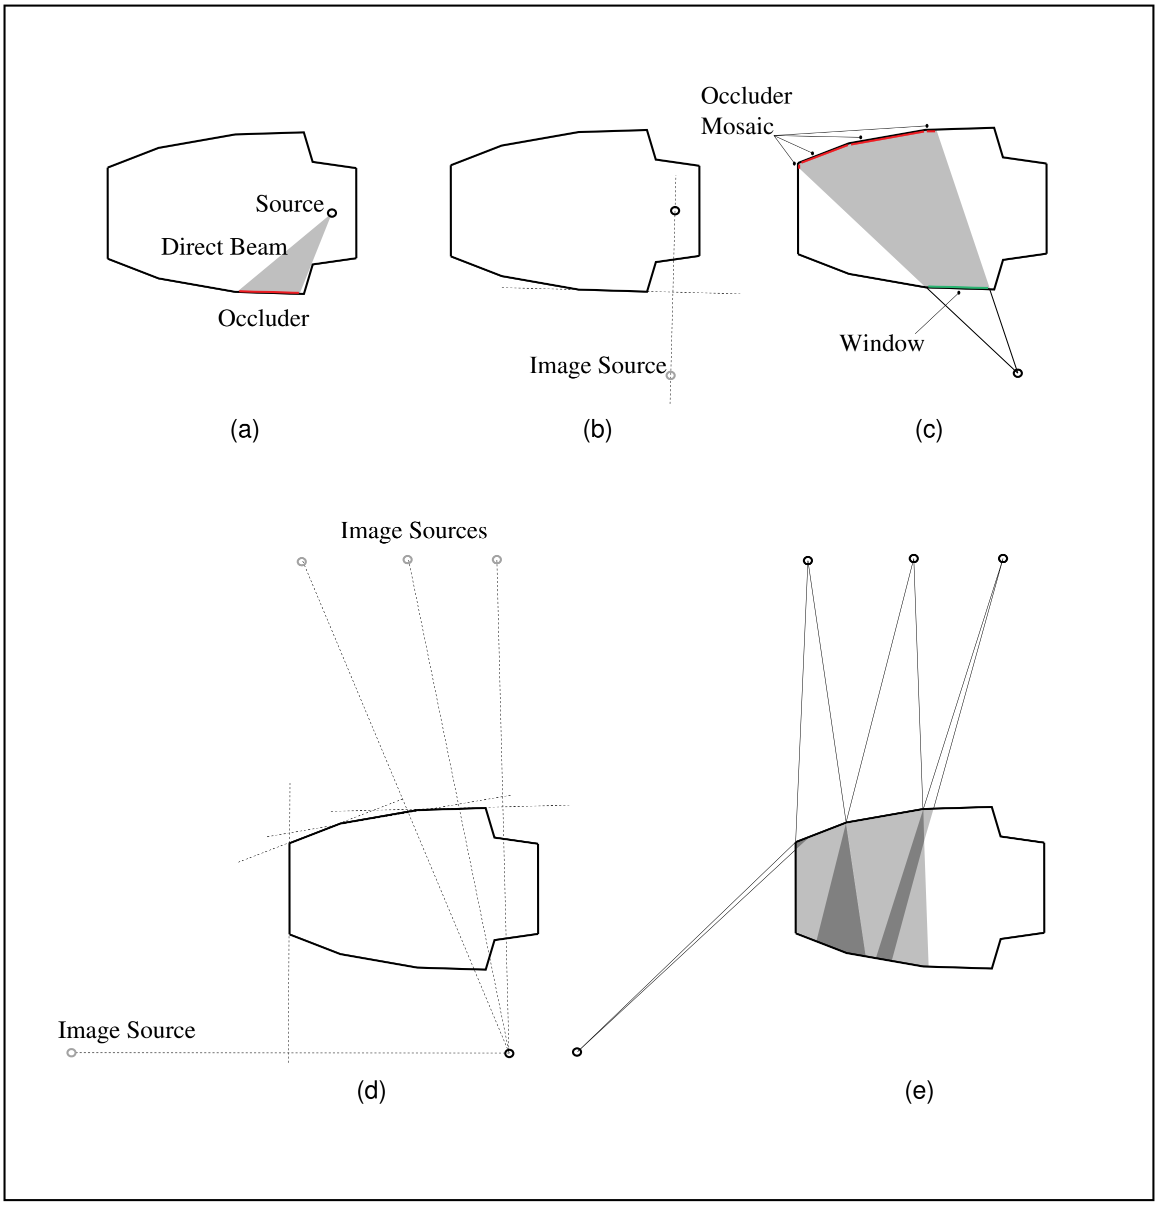

Figure 5-1: Construction series for three generations of beams. Occluders are shown in red, and windows are shown in green. Dotted lines indicate construction lines for mirrored image source location.

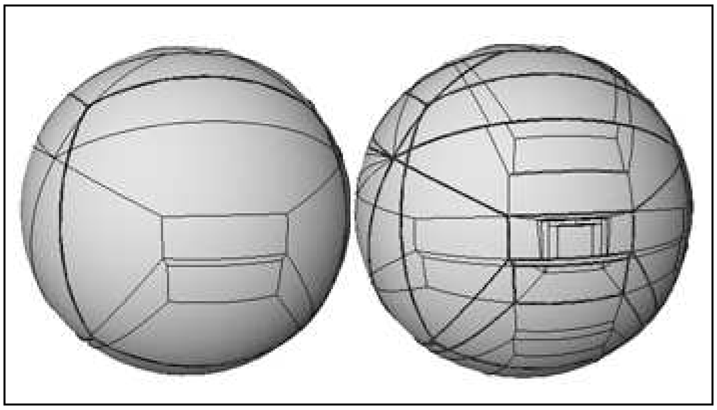

Figure 5-2: Beam profiles mapped onto a sphere centered at the source. a) First generation. b) Second generation.

the point onto the beam window. The point may be shadowed by an intervening surface. The receiver point must lie between the window and the occluder that contains the projection of the receiver point. Figure 5-3 shows one such beam from various views, and a receiver point within it. A ray originates at the image source and extends in the direction of the listener point. The ray pierces the window first, then the listener point, then the occluder. The occluder would shadow the listener point if the ray intersected the occluder first.

Once the hierarchical beam tree is constructed, each beam's energy is calculated given absorption characteristics of the surfaces in its reflection history. It is important to note at this point that even if the absorption characteristics are changed, the beam tree structure remains valid. Only the energy contained by the beams needs to be updated. If, however, the geometry of the enclosure is altered, or the position of the source is modified, part or all of the beam tree structure is invalidated and therefore must be regenerated. Full advantage is taken of these features in our inverse design system.

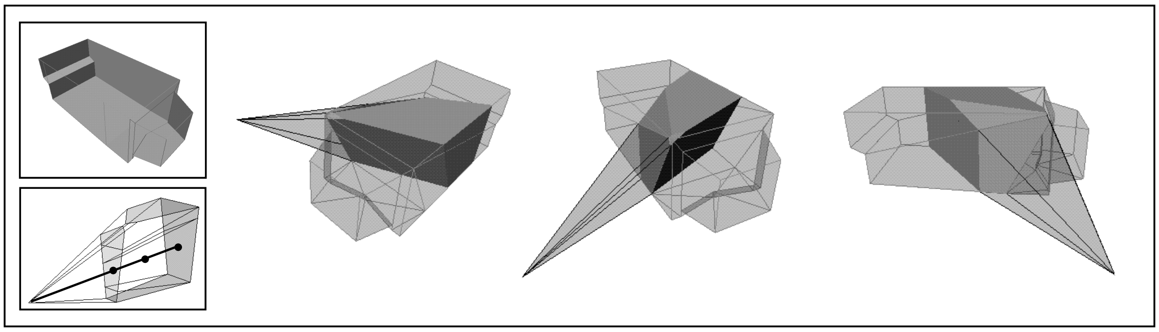

Figure 5-3: Receiver point within a beam projects onto the window and an occluder.

5.2 Improved Statistical Tail Approximation

Shortly after a sound begins to propagate within an enclosure, after a small number of reflections, the behavior begins to resemble an ideally diffuse field, with wave fronts moving in all directions. At this point it is reasonable to neglect the directional component of the sound field, using a statistical model to capture the behavior of the remaining sound. Beranek [8] describes an approximation to the decay of sound in an enclosure as a function of time based on the volume, surface area, and average absorption coefficient over the surfaces. While this serves as a reasonable coarse approximation, it fails to take into account such factors as proximity and orientation of surfaces to the source. Our beam method allows us to improve this approximation by accounting for these factors. After the first reflection, one can determine the remaining sound energy and the statistical distance, and thus the time of reflection, and use these results as input to the standard approximation. In fact, one could determine these values for any reflection depth, further improving the approximation.

This claim is realizable as a direct result of the beam method. The sound energy is calculated at the first reflection as follows:

1. For each first generation beam, determine the projected solid angle and projected center of gravity for each of its occluders.

2. Map the center of gravity back onto the surface and determine its distance and direc-

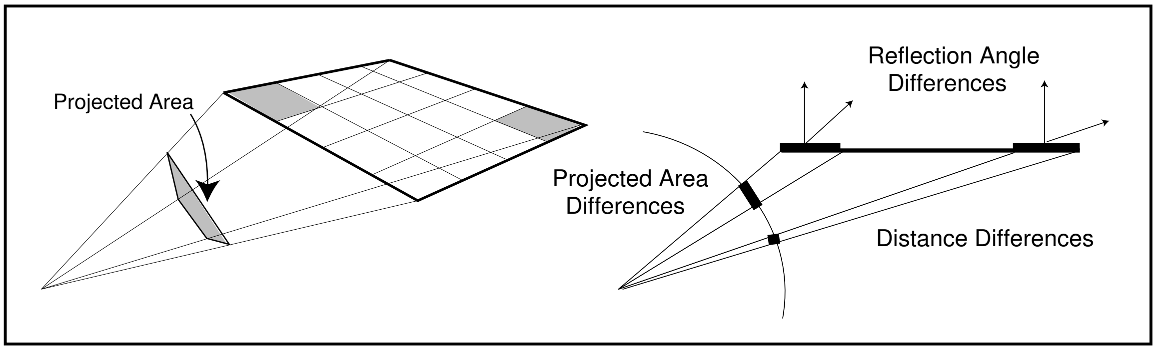

Figure 5-4: Improved energy decay calculation based upon projected areas, effective $\bar{\alpha}$.

tion to the source.

3. Use the angle between this vector and the occluder face normal to determine the angle dependent absorption.

4. Accumulate the energy contribution e_i and distance δ_i for each occluder *i* weighted by projected area a_i. If the occluder extent is too large, as in Figure 5-4, approximation errors will be significant.

5. Recursively subdivide each region to reduce the errors.

The statistical distance Δ and the statistical time Θ are given by the following equations:

$$\Delta_1 = \frac{\sum(\delta_i * a_i)}{\sum a_i}$$

$$\Theta_1 = \frac{\Delta_1}{c}.$$

The energy E remaining at Θ_1 is given by

$$E_1 = \frac{\sum(e_i * a_i)}{\sum a_i}.$$

Air attenuation is applied to the sound energy as well. For later reflections the energy contribution accounts for the absorption due to previous reflections.

This approximation technique provides the amplitude of the decay curve that is used to determine the reverberation time. However, it does not address the time distribution of hits

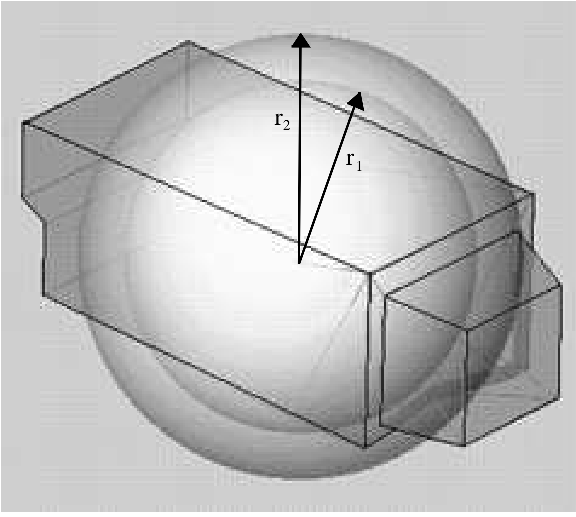

Figure 5-5: Sphere with radius r_1 has the same volume V_1 as the enclosure. Sphere with radius r_2 has volume V_2, twice the volume of the enclosure. The statistical time when a typical receiver will have incurred n hits is $\frac{r_n}{c}$.

that is needed to calculate the acoustic measures used to evaluate the hall. Our approximation is based on the idea that if the total volume of beams propagated through an enclosure equals *n* times the volume of the enclosure, then on average a representative point in the enclosure will be contained by *n* beams. The approximation of the time distribution of hits is as follows.

- Determine the radius r_1 of a sphere whose volume v_1 is equal to that of the enclosure.
- Convert r_1 to time θ_1 to get the statistical time of the first reflection. Statistically, the *n*th hit occurs at a distance equal to the radius r_n of a sphere with volume $v_n = n * v_1$. Likewise, the number of hits h_n by time θ_n is $(\frac{\theta_n}{\theta_1})^3$. This approximation can be improved slightly if the statistical time derived in the amplitude calculation for the first reflection is used. This is accomplished by shifting the time line by the difference between the geometrically determined time and the idealized time.

The statistical tail is constructed by combining the hit times with the decay curve information. Individual hits are recorded from time Θ to 0.5 seconds. The interval between 0.5 and 2.5 seconds is partitioned into ten time spans, each 0.2 seconds in duration and assigned the energy equal to that of the decay curve at its midpoint times the number of hits that occur during the span. The interval from 2.5 seconds to 5 seconds is handled similarly as a single time span.

This approach is sufficient to provide the data necessary to calculate the measures used to evaluate the acoustics of an enclosure. In retrospect, instead of using fixed time points, it might have been preferable to set them relative to the mean free path d of each enclosure. A representative listener location in a large hall will receive far fewer hits by a given time than its counterpart in a small space. By using relative time points, the computational detail and complexity would be much more uniform for spaces of any dimension.

Since this simulation engine is being used in the context of an optimization design system, it should be noted that whenever the absorption characteristics of surfaces within the enclosure are altered, the energy decay envelope is affected, and must be recalculated. Any geometric change will effect the onset and frequency of hits as well. Hence, the tail is simply regenerated for any change to the enclosure description.

5.3 Sound Field Point Sampling

In order to sample the sound field at a point within the enclosure both the beam and statistical tail components of the simulation are considered. A hit list is constructed containing the time, intensity, and direction of every passing beam. Only beam volumes that contain the sample point are considered. The statistical tail data is then included to produce the full compliment of data required for the calculation and visualization of objective acoustic parameters. No directional information is kept for the tail, since its derivation is based on diffuse field assumptions. Figure 5-6 illustrates this approach.

One might assume that a beam would have an identical effect on all sample points it contains. However, this is not the case. For a given sample point, the arrival time and direction are determined directly from the vector connecting the virtual source of the beam and the sample point. The intensity varies throughout the beam as a function of distance

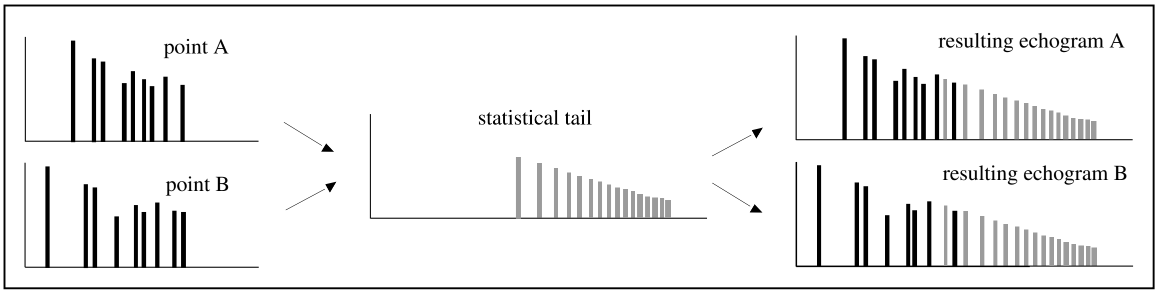

Figure 5-6: Hits from beams are combined with the statistical tail.

from the source. Surface attenuation from previous reflections effects the initial energy of a beam, and air and distance attenuation further reduce this energy as the beam propagates. It is also possible to account for a source with a non-uniform directional energy distribution, angle dependent surface absorption, and phase, by inverse mapping reflections back to the source. I have implemented angle dependent surface absorption, but found that the increase in accuracy did not justify the computation cost and subsequent loss of interactivity.

5.4 Multidimensional Sampling

The system supports the use of higher dimension receivers as well. Various spatial subdivision approaches can increase the efficiency of the sampling by representing large near-uniform regions with a single sample. If the variation in the simulation data within a given region exceeds a limit, the region is divided into two regions along each dimension, and each resulting region is processed similarly. I use a three dimensional subdivision scheme, an octree hit structure for volume receivers, which may be used to aid in positioning diffusers, reflectors, and balcony seats, or to gain an overall understanding of sound transport within the enclosure. A two-dimensional quadtree structure is used for area receivers, covering a seating area, or wall, for example. A row of seats may be represented by a line receiver with a binary tree structure.

The efficiency of these structures increases as receiver points get closer together and virtual sources get farther away. In both cases, the effect of a beam's energy contribution

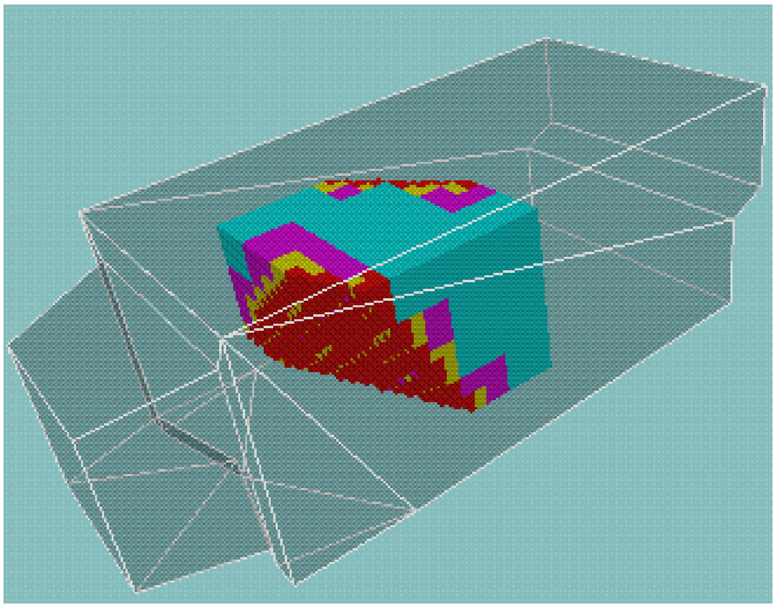

Figure 5-7: Octree representation of a beam intersected with a 3D receiver.

on two adjacent points with respect to their previously accumulated energy is reduced.

The recursive algorithm proceeds as follows:

1. If there is no intersection between the receiver region and the beam, disregard it.
2. If all extrema of the receiver region are not inside the beam, subdivide the receiver region, and for each piece, go to step 1.
3. If the range of energy levels within the receiver region is large, i.e. greater than a given tolerance, subdivide the receiver region, and for each subregion, go to step 1.
4. Record a single hit for the entire receiver region.

This algorithm takes great advantage of spatial and temporal coherence. Figure 5-7 shows the octree representation for the intersection of a beam and a 3D receiver volume.

This acoustic simulation algorithm has several features that work well with the audioptimization design system. The sound field for the entire enclosure volume is pre-calculated, allowing sampling at any location and time without having to reconstruct the beam tree. If only materials are changed between design iterations, then the only aspect of the hit list for a given listener position that changes is the intensity of the hits. The time and direction of the hits are unaffected by these changes. This allows the system to accelerate material optimization, saving the expense of reconstructing the beam tree. The user may also trade off accuracy and speed by specifying the reflection depth, that is, the depth to which beams are traced. When this depth is reached, sound is modeled with statistical methods. The shallower the reflection depth, the faster the simulation calculation, but the lower the accuracy. This flexibility allows the designer to receive more timely feedback during early design, and more accuracy for late or final design.

5.5 Sound Field Visualization

The system provides two very different tools which present the raw sound field data to the designer: the beam animator, which presents the sound field from a position independent viewpoint, and a viewer which presents a series of ”tempolar plots”, visualizations I have designed and named which present detailed listener position dependent data. Both tools use the beam tracing data exclusively, which contains useful directional information for each sound wave.

5.5.1 Beam Animation

In this system the sound field generated by an impulse within an enclosure is represented by a collection of beams, or wave fronts moving through the space. The user is allowed to view these beams though time as they propagate within the enclosure. Figure 5-8 shows a series of snapshots from an animated sequence to illustrate this visualization tool. The tool offers insight into the directional effects that geometric components have on the sound field.

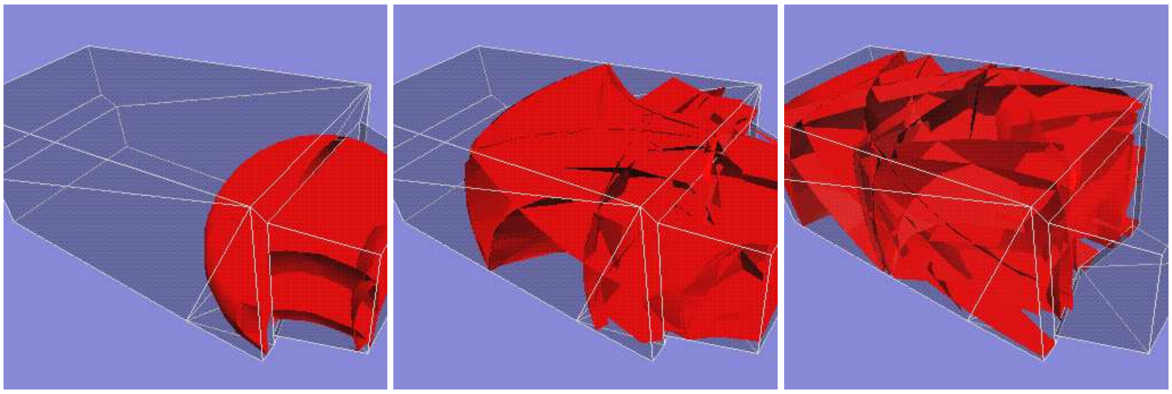

Figure 5-8: Three snapshots of wave front propagation.

5.5.2 Tempolar Plots

Frequently it is effective to represent the sound field from the perspective of a single location, perhaps the location of a listener within the space. The sound field, when point sampled, is represented as a list of hits, one for each wave front that passes the point location. Such a hit has three components: arrival time, arrival direction, and the intensity of the wave front as it passes. Standard point sample visualizations represent only two of the three components. The echogram conveys level as a function of time. The sound rose presents level as a function of direction.

I have designed a visualization, which I call the tempolar plot, that combines all three data components. Each hit is represented as a splotch on a polar plot. The direction of a passing wave front determines the angle the splotch makes with respect to the plot origin. The three dimensionality of the data is reduced to 2D by rotational projection about the lateral axis of the listener. That is, the angle of the impinging sound with respect to the lateral axis is represented, not the actual direction of the wave front. The splotches appear on the left side of the plot if they arrived from the listener's left. Arrival time determines the distance from the origin. Shades of gray are used to represent sound level. This plot combines the advantages of the echogram and sound rose into a single representation.

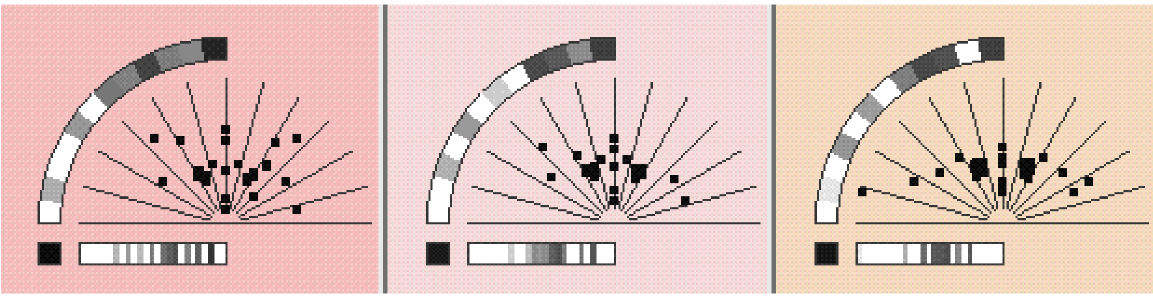

Figure 5-9: Tempolar plots for three hall locations.

The tempolar plot also represents the total energy integrated over angle and time, composing the peripheral strips shown in Figure 5-9. The plot can be viewed through time. The contributions of energy are added to the plot at their respective arrival times, and integrated level indicators are updated through time.

In this chapter I have provided detailed description of the acoustic simulation algorithm used in the acoustic design system presented in this thesis. While this algorithm offers many desirable features which can be leveraged an acoustic design context, It could conceivably be replaced in this design system by an improved algorithm, should one become available.

Chapter 6

Inverse Problem Formulation

In this chapter, I pose the acoustic design task as a constrained, non-linear optimization problem. I begin with a brief discussion of optimization techniques and then move on to the problem formulation.

6.1 Optimization

The role of optimization in a design system is to find the configuration in the feasible design space that best matches desired performance goals. The choice of an optimization technique depends on the nature of the design space and the types of constraints. The goal of optimization is to minimize an objective function of n system parameters while satisfying a set of constraints [11, 51]. Standard non-linear optimization techniques use the gradient and curvature of the objective function to descend to a minimum, or locally optimal configuration [11]. The objective functions encountered in the context of audioptimization have multiple such minima, and therefore require a global strategy. Hence, we first employ a global optimization technique, *simulated annealing*, to locate a more globally optimal neighborhood and then use the *steepest descent* algorithm to descend to the minimum. The existence of constraints implies that not every target is realizable.

In order to accomplish this, I formalize the problem as follows: given a description of a set of desired measures for acoustic performance, determine the material properties and geometric configuration that will most closely match the target. In order to formulate the acoustic design process as a constrained optimization problem, a specification of 1) the optimization variables that express how a hall may be modified, 2) the constraints that must be satisfied, and 3) the objective function must be specified.

6.2 Optimization Variables

In a typical acoustic simulation system, the goal is to compute the sound field in a scene assuming a sound source and a description of the geometry and materials. The measures described in Section 4.1 are the unknowns, which are computed in terms of static material properties and geometry. In the optimization problem, material and geometric properties are no longer fixed but are treated as variables.

Sabine writes…

> Broadly considered, there are two, and only two, variables in a room — shape including size, and materials including furnishings. In designing an auditorium an architect can consider both; in repair work for bad acoustical conditions it is generally impractical to change the shape, and only variations in materials and furnishings are allowable. [45]

A hall is comprised of a collection of polygons, subsets of which may be grouped into geometric *components*. Components are a convenient and natural way to represent entities such as balconies, reflectors, etc. Each component can have associated with it a set of allowable geometric transformations and a set of acceptable materials. Each translation, rotation, or scaling of a component represents a *geometry variable*; a set of possible materials associated with a component is a *material variable*.

6.3 Constraints

There are two types of constraints. *Geometry constraints* are user-specified upper and lower bounds placed on each transformation of a component. Each transformation variable represents a single degree of freedom: translation along a vector, rotation about a vector, or scaling about a point or along a vector. The allowable range of each transformation constraint requires the component to remain within the specified bounds. For example, $T_{low} \leq T_i \leq T_{high}$, requires the transformation i to remain within the bounds T_{low} and T_{high}.

Material constraints are user-specified sets of allowable materials that are assigned to a given component. This subset of materials is selected from a library of materials, which is provided by the system. For example, $\{plaster, concrete, fiberglass\}$ is a set of materials for a component. In addition, the physics of sound propagation imposes physical constraints that govern the behavior of sound as it passes through air, is reflected, and is diffused from surfaces within a hall. These constraints require that all materials have absorption coefficients in the range of 0.0 to 1.0, which prevents the sound energy from becoming negative—keeping the simulation in the realm of physically meaningful solutions.

6.4 Objective Function

Audioptimization problems typically have infinitely many feasible configurations that satisfy the constraints. An *objective function* is necessary to select the optimal configuration from the feasible set. We need to to define a function that produces a single rating for the hall, based on the effect these measures have on the listener.

A successful objective function will consider the following. First, since one acoustic measure may have a more detrimental effect on the experience of the listener than another, each measure must be weighted appropriately. Second, for each measure, different values must be penalized in proportion to their adverse effects on a listener's experience. These two considerations are not independent of each other.

We use Beranek's Objective Rating Method (ORM) as our objective function, which is an application of Ando's Theory of Subjective Preference [5, 10]. Ando found that when m orthogonal objective acoustic measures are given, the following equation based on the responses of human listeners provides the scalar rating of a hall:

$$f(x) = \sum_{i=1}^{m} w_i f_i \tag{6.1}$$

where the multidimensional vector x is the configuration of the hall, the function f_i penalizes the deviation from the target value of each objective acoustic measure, and the weight w_i normalizes the respective functions.

Beranek uses the six objective acoustic measures mentioned earlier, IACC, EDT, BR, G, TI, and SDI, and provides the following values for their weights:

$$\begin{aligned}
w_{\text{IACC}} &= 1.2 \\
w_{\text{EDT}} &= \begin{cases} 9 & \text{if EDT} < 2.0 \\ 12 & \text{if EDT} > 2.3 \\ 0 & \text{otherwise} \end{cases} \\
w_{\text{BR}} &= 10 \\
w_{\text{G}} &= \begin{cases} 0.04 & \text{if G} < 4.0 \\ 0.07 & \text{if G} > 5.0 \\ 0 & \text{otherwise} \end{cases} \\
w_{\text{TI}} &= 1.42 \\
w_{\text{SDI}} &= 1.0.
\end{aligned}$$

The definitions of the penalty functions follow:

$$f_{\text{IACC}} = \left|\text{IACC}(x) - \text{IACC}_{target}\right|^{3/2} \tag{6.2}$$

$$f_{\text{EDT}} = \left|\log\left(\frac{\text{EDT}(x)}{\text{EDT}_{target}}\right)\right|^{3/2} \tag{6.3}$$

$$f_{\text{BR}} = \left|\log\left(\frac{\text{BR}(x)}{\text{BR}_{target}}\right)\right|^{3/2} \tag{6.4}$$

$$f_{\text{G}} = \left|\text{G}(x) - \text{G}_{target}\right|^{3/2} \tag{6.5}$$

$$f_{\mathrm{TI}} = \left|\log\left(\frac{\mathrm{TI}(x)}{\mathrm{TI}_{target}}\right)\right|^{3/2} \quad (6.6)$$

$$f_{\mathrm{SDI}} = \left|\log\left(\frac{\mathrm{SDI}(x)}{\mathrm{SDI}_{target}}\right)\right|^{3/2}, \quad (6.7)$$

where IACC_{target}, EDT_{target}, BR_{target}, G_{target}, TI_{target}, and SDI_{target} are the user-specified target values; and the functions $\mathrm{IACC}(x)$, $\mathrm{EDT}(x)$, $\mathrm{BR}(x)$, $\mathrm{G}(x)$, $\mathrm{TI}(x)$, and $\mathrm{SDI}(x)$ are objective acoustic measures for configuration x.

Finally, the objective function given by Equation 6.1 is minimized to find the hall configuration that best matches the target objective acoustic measures. Note that objectives may be constructed from all or a subset of these terms. Appendix B describes how the objective function presented above can be modified with performance goals that are appropriate for various uses of the hall. I follow this with a discussion of how to construct an objective to handle multiple performance types simultaneously, such as symphonic music and opera.

6.5 Optimization Problem

The optimization problem can now be stated as follows: minimize $f(x)$ subject to $x \in X$, where the constraint set X is the "design space" spanned by feasible hall configurations. The existence of constraints implies that not every target is realizable. The optimization process must identify an optimal point in the design space, i.e. $x^* \in X$, such that $f(x^*) \leq f(x)$, $\forall x \in X$. Simulated annealing and steepest descent techniques are used in combination to search globally for the "best" hall configuration. I discuss these techniques in detail in the following chapter.

Chapter 7

Implementation

In this chapter I describe the implementation of an interactive design system based on audioptimization. The framework of our approach is illustrated in Figure 7-1 and summarized below.

The user provides an initial model that is passed to the simulation algorithm to compute a baseline sound field solution. Then, to specify desired targets for acoustic measures, the user manipulates icons and paints desired sound level values onto a subset of hall surfaces at selected time steps. The user then constrains the design space for the system to search, by indicating the range of modification for variable material and geometric components. The user then optimizes over materials and geometry either separately or simultaneously. Once the optimization process has been initiated, the user can interrupt it to modify goals, add and/or delete variables and modify their constraints, or adjust optimization parameters. Once all the design goals and variables are specified, the optimization process is run until convergence is achieved.

The following pseudo-code describes the process:

Compute baseline sound field solution.
Establish constraints, objectives, and optimization parameters.
repeat
    Invoke simulated annealing.
    Invoke steepest descent.
    Display results.
    Modify constraints, objectives, and optimization parameters if desired.
until *convergence.*

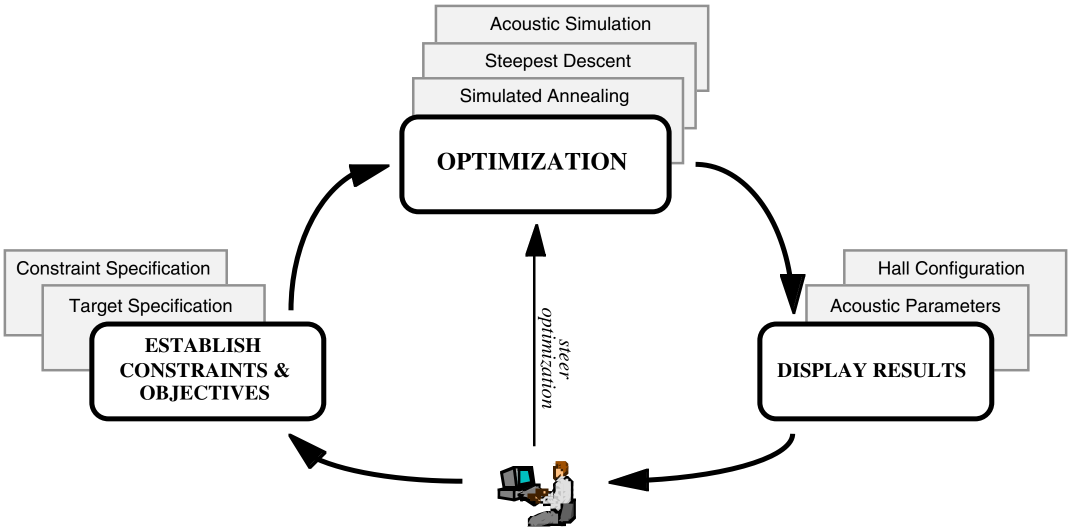

Figure 7-1: Overview of the interactive design process.

7.1 Simulation

The sound field for the initial model configuration is simulated using the hybrid simulation algorithm presented in Chapter 5. A standard source—a single omnidirectional full band spherical impulse—is used, which is simulated and propagated into the environment. The user can trade simulation quality for speed through a number of interactive controls.

Visualizations described in Chapter 4 display acoustic measures derived from the simulated sound field at a set of user-specified locations within the hall. A fine mesh of sample points displays sound level data over time, which is indicated by a color scale as shown in Figure 4-4.

7.2 Constraints and Objectives

Before initiating the optimization process, the user specifies a range of acceptable modifications to the hall. This specification involves selecting component modifications to surface materials and geometric transformations. The user must also set acoustic performance goals so the system can evaluate different hall configurations.

7.2.1 Constraint Specification

The design space describes all the ways the user will allow the enclosure to be modified, and is typically composed of a very large number of configurations. It is this design space that is searched by the optimization process to find the configuration that best meets acoustic goals. The design space is specified interactively, as the user selects the components within the enclosure that can be modified, and the range of modification permitted.

In order to specify a material constraint the user chooses a set of allowable materials with the Material Editor (Figure 7-2) for a component. An array of absorption coefficients corresponding to the frequencies ranging from 31 Hz to 16 kHz, and an SDI value describe each material. Table C.1 lists a subset of the material library. During optimization the component may be assigned any of the selected set of materials. The set is ordered based upon the absorption characteristics of each material. This ordering allows the optimization to utilize a 'difference' metric when selecting from material choices.

The user interactively imposes constraints on the transformation of geometric components using the Geometry Editor. After selecting a component and transformation type, the user indicates the component's positional degree of freedom by placing and orienting a coordinate system axis icon (Figure 7-3a). The user sets boundary constraints by positioning the component at the desired range limits. This range is discretized into a user-specified num-

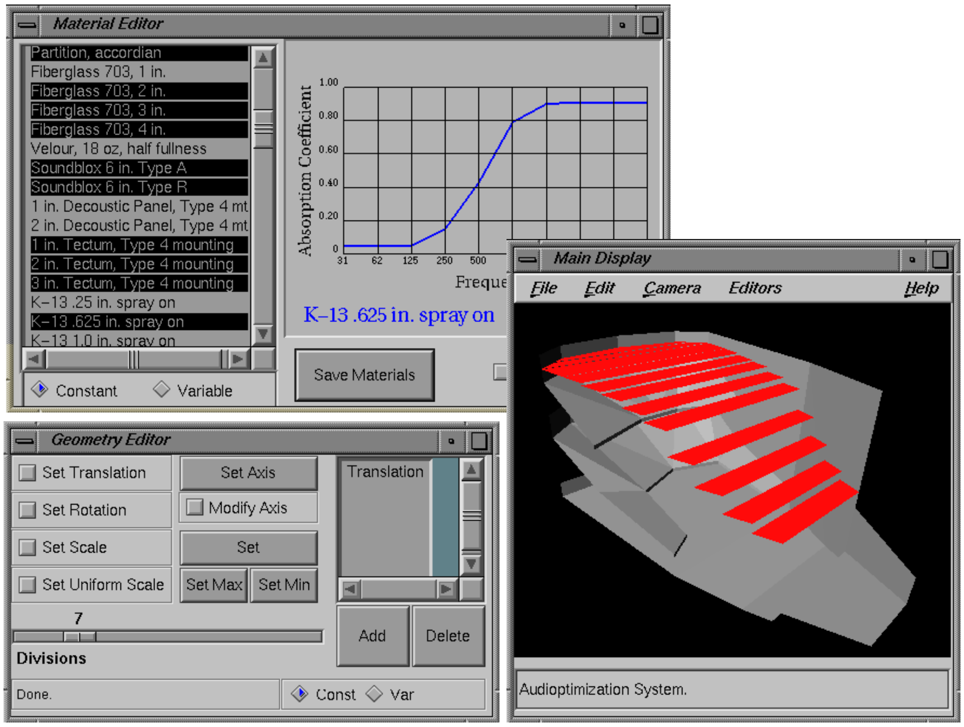

Figure 7-2: Material and geometry editors.

ber of configurations. A rotation constraint is specified by orienting and selecting a rotation axis, then interactively reorienting a component to both minimum and maximum configurations, shown graphically in Figure 7-3b. Figure 7-3c shows the possible states resulting from translation constraint specification for a set of ceiling panels. Figure 7-3d shows a scale constraint specified by indicating a point and direction of scale. Each transformation specified by the user becomes a design variable, and may be set to any of the discretized values during optimization. Conflicting constraints are resolved in the order that they were defined. For example, if the user first scales one wall, then rotates an adjoining wall, the scale operation precedes the rotation operation in constructing the resultant configuration.

During the design of this system the idea of allowing the system to determine the ideal material characteristics for a component was considered, thus removing the constraint that the material must be selected from the constraint set. This approach has merit, and should pursued in the future. However, it seemed more practical to select from a set of real-world

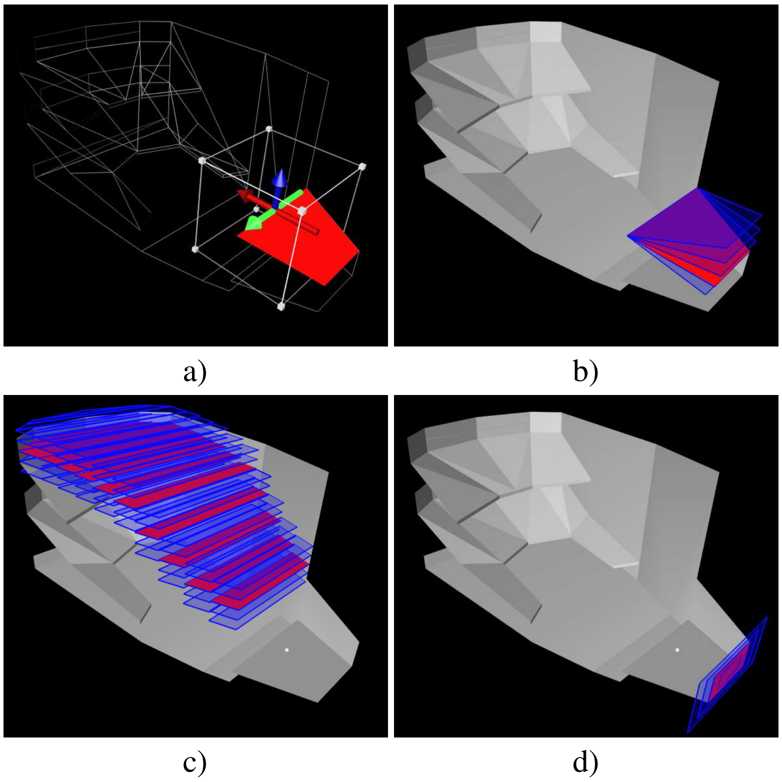

Figure 7-3: Geometry constraint specification. a) Coordinate system axis icon used for transformation specification. b) Rotation constraint specified by orienting and selecting a rotation axis. c) Possible configurations resulting from translation constraint specification for a set of ceiling panels. d) Scale constraint specified by indicating a point and direction of scale. The constraints are discretized according to user specified divisions.

material choices. A reasonable variation might be to let the system determine the ideal material characteristics, then locate the best match from the real-world set of materials.

The definition of a search neighborhood surrounding a geometry variable configuration comes naturally from this discretization process. The distance between configurations is simply the absolute difference between their configuration indices. The definition of a neighborhood surrounding a material variable configuration is less straightforward. Materials are ordered by the system based on their average absorption coefficients. The distance between configurations is then given by the absolute difference between their ordering in-

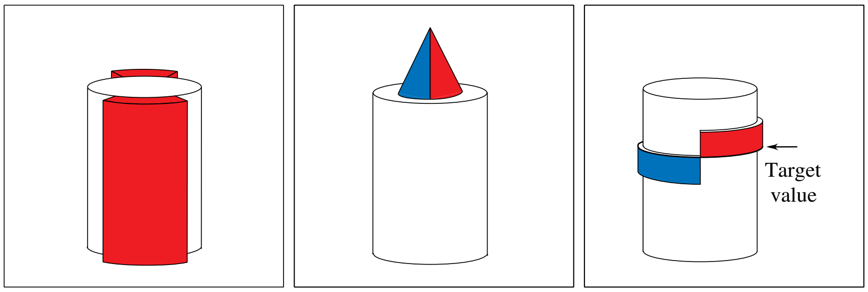

Figure 7-4: Graphical difference icons representing (from left to right) IACC, EDT,and BR.

dices. Because the ranges of materials and geometric transformations are discrete, combinatorial optimization is feasible.

7.2.2 Acoustic Performance Target Specification

The targets for IACC, EDT and BR for point samples can be set interactively by manipulating the point sample icon. The geometric representations of targets are selected, then dragged into positions that represent desired values. Textual feedback is displayed aside the icons during manipulation.

The system provides three viewing modes for iconic representation. The first mode displays the actual value of the three acoustic measures. The second mode displays the target value for the measures. The third shows the difference between the actual and target values. It is easiest to see the positional subtleties of the sound field using this difference mode. Figure 7-4 shows positional difference icons. IACC difference is represented as the region not covered by the actual value. Under ideal conditions the IACC difference shell is absent since it represents directions *not* covered by incoming sound. EDT difference is characterized as a solid cone whose radius indicates absolute difference and whose color indicates whether the actual value is higher than (red), or lower than (blue) its target. BR difference is represented as a ring between the actual and target values, similarly color coded.

The acoustic measure G has both a positional and temporal component. One sets scalar values for sound level targets by selecting a paint color from the color palette and painting

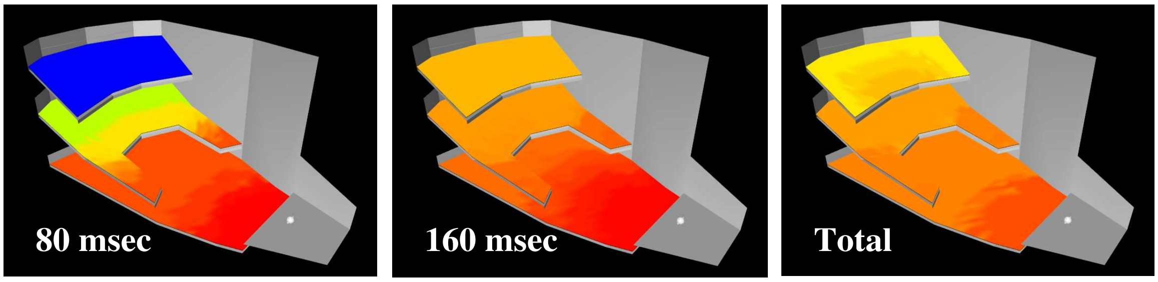

Figure 7-5: Sound strength target specification.

sound level or a change in level directly onto the seating regions. Unlike other measures, sound level may be specified at various time points during the simulation. While the final sound level is important, the time distribution of sound is critical to our perception. Figure 7-5 illustrates three painted time slices. Various brush shapes and sizes are provided.

A scrollable viewer indicates sound level with a column of three thumbnail images for each time slice (Figure 7-6). The top row shows the color coded visualization for the actual sound level. The middle row shows target values specified for a given time slice. At the bottom is a red-blue signed difference image between the actual and target values, where difference is defined as the actual value minus the target value. The user defines the mapping by specifying the range limits represented. Maximum positive difference is mapped to red, maximum negative difference is mapped to blue, and minimum difference is mapped to black, with intermediate values mapped to interpolated colors. This tool provides an easy way to assess and specify desired performance through time.

Discussion

There is a degree of subtlety to specifying sound level targets, particularly for the early time points. One could simply paint the desired sound distribution exactly as desired. However, one might also wish to manipulate these targets as one would manipulate knots on a spline curve. One could think of these painted targets as spline curve knots, that is, control points that influence the shape of a curve without necessarily residing on the path of the curve itself. In this case, the painted targets influence the sound level 'curve'. One can control the

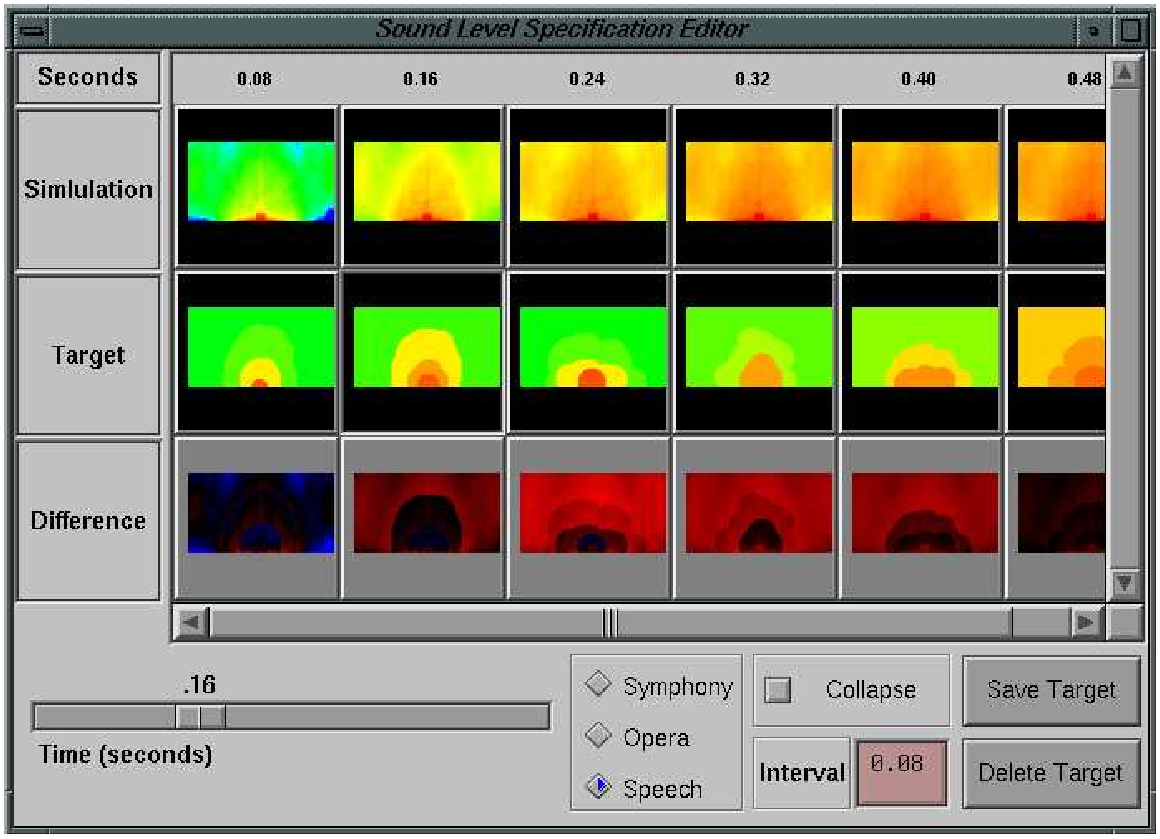

Figure 7-6: Sound level specification editor.

influence of these painted targets by exaggerating their values, pulling them away from the curve, and relying on the optimization process to modify the curve to approach these points.

In order for the early targets to exert any influence on the optimization process, the target must be specified at a time on or after the first reflected sound wave arrives at a point. An earlier target will always produce the same penalty, since the direct sound is all that has arrived by that time, and it is not effected by any modification to the enclosure. If a target is specified too long after the reflected sound starts to arrive, then its influence in tugging at the temporal distribution of sound is lessened. Ideally, the target will be set for a time closely following the onset of reflected sound.

Alternative target setting techniques include textual specification, and preset performance goal selection. Predefined target values are provided for symphonic music, operatic music

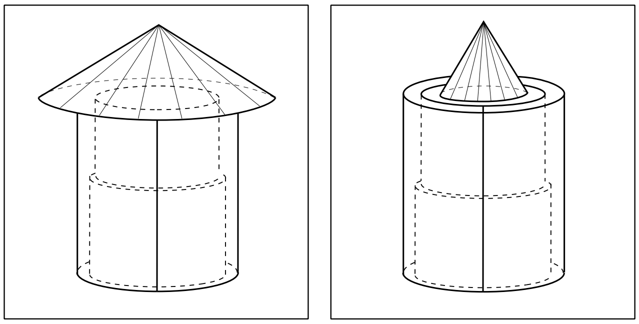

Figure 7-7: These illustrations show ideal values for two different types of hall uses. left: symphonic music. right: speech.

and speech, or combinations of these. Table B.1 summarizes the target values for the preset goals. Figure 7-7 shows preset target values for two cases.

7.3 Optimization

Having specified the evaluation function and the design space in previous sections, it remains to define the final component required by inverse design: the design space search strategy. The role of optimization in a design system is to find the configuration in the design space that best matches desired performance goals.

In general, the audioptimization objective function contains many local minima. Therefore, we first apply a global optimization step using simulated annealing to locate the neighborhood of a good solution and then follow with the steepest descent algorithm, which is more efficient in finding the local minimum.

7.3.1 Global Optimization: Simulated Annealing

Simulated annealing is a combinatorial optimization algorithm that produces a series of transitions between configurations in the design space based on three components: a generation mechanism, a cost function, and an acceptance function [29]. The generation mechanism randomly selects a new configuration x_{new} from within a neighborhood around the current configuration $x_{current}$, where the neighborhood size is determined by ϵ. The cost function f

evaluates x_{new} as defined in the previous section. The acceptance function accepts or rejects x_{new} by comparing the current cost $f(x_{current})$ to the new cost $f(x_{new})$. If x_{new} is accepted, a transition results and x_{new} replaces $x_{current}$. Unlike local optimization methods, which accept only lower cost transitions terminating at the local minimum, the simulated annealing method uses the Metropolis algorithm where the probability of accepting a higher cost configuration is non-zero [1, 41]. This feature allows the search to proceed uphill, away from a local minimum, in search of a more global minimum.

A transition from $x_{current}$ to x_{new} is made with probability M, where:

$$M(x_{new}, x_{current}, \tau) = \begin{cases} \exp\left(\frac{f(x_{current}) - f(x_{new})}{\tau}\right) & \text{if } f(x_{new}) > f(x_{current}) \\ 1.0 & \text{if } f(x_{new}) \leq f(x_{current}). \end{cases} \tag{7.1}$$

τ is a temperature control parameter that slowly drops or "cools" throughout the annealing process, gradually decreasing the probability of accepting a higher cost configuration. A user-defined *annealing schedule* determines the value of both temperature and the neighborhood parameter ϵ used in the generation mechanism. The user specifies the initial values τ_0 and ϵ_0, as well as the initial number of configurations γ_0 to sample, before changing these annealing control parameters. Finally, to complete the annealing schedule, the user specifies the percentage by which each value is changed after sampling γ configurations.

Once the annealing process is initiated, it terminates when satisfying one of two stop criteria: either the process reaches a user-specified limit on the number of configurations to evaluate; or the process converges, which occurs when an exhaustive evaluation and rejection of each configuration in the current neighborhood is achieved.

Once the stop criteria are met, the user may assess the current state of the sound field. The user may then conclude the optimization process by invoking the steepest descent step. Alternatively, acoustic performance goals may be altered, optimization variables may be added or deleted, or their constraints modified, optimization parameters may be modified, then the optimization may be restarted. If the user has chosen not to make any modifications and the optimization process has not yet converged, the optimization may be resumed.

The simulated annealing algorithm is expressed in pseudo-code as follows:

```
procedure SimulatedAnnealing()
    k ← 0
    AnnealingSchedule_InitParams(τ_k, ε_k, γ_k)
    x_current ← InitialRandomConfiguration()
    repeat
        for (i ← 0; i < γ_k; i ← i + 1)
            x_new ← GenerateNeighborConfiguration(ε_k, x_current)
            if (M(x_new, x_current, τ_k) > random[0,1))
                x_current ← x_new
            end if
        end for
        k ← k + 1
        AnnealingSchedule_UpdateParams(τ_k, ε_k, γ_k)
    until Stop Criteria
end procedure
```

7.3.2 Local Optimization: Steepest Descent

Simulated annealing is followed by steepest descent if convergence has not been attained. Starting at $x_{current}$, successive steps are taken between neighboring configurations with decreasing cost. The algorithm terminates when no neighboring configuration has a lower cost than the current configuration.

A descent direction is determined as follows. An n-dimensional direction vector d is computed, where n is the number of optimization variables. Each vector component d_i, where $i = 0, \ldots, n-1$, may assume the value -1, 0, or 1, depending on which neighboring configuration has the lowest cost. The most suitable neighbor is located by defining a unit direction vector e, setting its ith component to 1, adding e to and subtracting e from $x_{current}$ in turn, and comparing the cost of the resulting configurations. This discrete version of the steepest descent algorithm is shown in the pseudo-code below.

```
procedure SteepestDescent()
    repeat
        d ← 0
        for (i ← 0; i < n; i ← i + 1)
            e ← 0
            e_i ← 1
            if ((f(x_current − e) < f(x_current)) and (f(x_current − e) < f(x_current + e)))
                d_i ← −1
            else if ((f(x_current + e) < f(x_current)) and (f(x_current + e) < f(x_current − e)))
                d_i ← 1
            else
                d_i ← 0
            end if
        end for
        if f(x_current + d) < f(x_current)
            x_current ← x_current + d
        end if
    until x_current is optimal
end procedure
```

7.3.3 Discussion

When both materials and geometry are free to change, the optimization algorithm is altered to take advantage of the fact that material modifications execute up to 50 times faster than geometry modifications. Therefore, a full material optimization is run after testing a new geometry configuration and the resulting cost is used in the acceptance test for the geometry optimization. The second example in Chapter 8 uses this approach.

Simulated annealing has both advantages and disadvantages. On the positive side it has a statistical guarantee of locating an optimal solution, and the objective function may be discontinuous and non-differentiable. This approach is well suited to searching discretized design spaces such as ours. On the negative side the control parameters (τ, ϵ, γ) need tuning for each new application to obtain the best results. In our case the cooling schedule required to guarantee statistically locating an optimal solution is impractically slow. As with many applications of simulated annealing, we relax this condition, and decrease τ faster than the formal definition of simulated annealing allows. This variation is called simulated quenching [24], and it provides an acceptable solution within a reasonable time frame.

I have presented the implementation details of a real acoustic design system build on inverse design principles. In the following chapter I will present case studies of existing performances spaces which illustrate the application of this system in the context of real world design.

Chapter 8

Case Studies

To demonstrate the techniques introduced in this thesis, I show results based on three real architectural spaces in which acoustic considerations play a prominent role. A large percentage of acoustic designs involve the renovation of existing spaces to correct acoustic problems; these examples illustrate such scenarios.

The system described in the preceding chapters has been implemented in C++ and run on a Silicon Graphics Onyx Reality Engine 2 workstation with 256MB of memory using a single 195MHz MIPs R10000 processor.

8.1 Oakridge Bible Chapel

The first example is Oakridge Bible Chapel in Toronto, Canada, whose initial design exhibited a fundamental acoustic flaw (see Figure 8-2). The walls and ceiling were faced with highly reflective plaster, which produced excessive reverberation, causing the sound from one spoken syllable to linger and mask the following syllable. Even when sufficiently loud, speech became almost unintelligible.

By simulating the existing environment, we were able to confirm the speech intelligibility problem. Plates 1–3 show the results of the simulation. The width of the cones and height of the bass cylinders indicate that both EDT and BR are too high (see Plate 1a). The

Figure 8-1: Interior photograph of Oakridge Chapel (courtesy of Dale Chote.)

sound strength visualization shown in Plate 3 demonstrates that the total sound level is also too high and that much of the distribution arrives late. The time distribution of sound energy has a significant effect on speech intelligibility; the earlier the sound arrives, the better.

We set out to improve speech intelligibility by reducing the initially high values of EDT and BR while maintaining adequate sound levels. We built the objective function using EDT, BR, and G, which are the three measures relevant for speech. Their respective target values are shown in Table 8.1. Note that these values differ from those used to evaluate symphonic music[7]. Plate 1b shows the target values and Plate 1d shows a difference image relative to the initial configuration. In order to improve speech intelligibility, we painted sound level target distributions on the seating plane for three time slices: 0.08 seconds, 0.16 seconds and total level, as shown in Plate 2. These particular times slices were chosen in order to include reflected sound.

Having set the acoustic goals, we selected modifiable components and specified the range of the modifications. Oakridge is typical of buildings that are constrained by their existing geometry, limiting redesign options. With this in mind, we restricted the optimization to include only changes to materials. We considered two design scenarios: one involving the ceiling surfaces—the most easily modifiable surfaces covering the largest contiguous

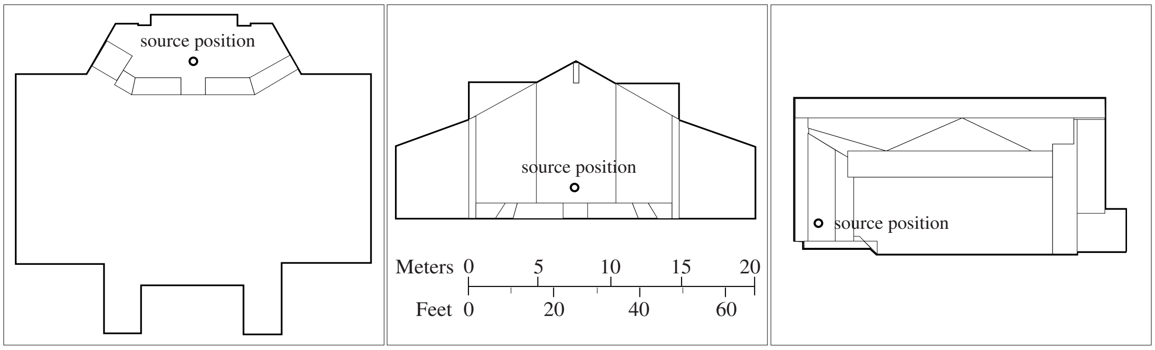

(a) Plan (b) Cross section (c) Longitudinal section

Figure 8-2: Computer model of Oakridge Chapel.

	EDT (sec)	BR	G (dB)	f(x)	Improvement
Target	0.700	1.000	≥ 0.0	0.0	NA
Initial Configuration	1.899	1.631	12.141	4.835	NA
Final (Ceiling Only)	0.784	1.002	7.760	0.120	97.5 %
Final (Walls and Ceiling)	0.651	1.026	6.779	0.098	98.0 %

Table 8.1: Acoustic measure readings for Oakridge Chapel.

area—and the second involving both the ceiling surfaces and walls.

In the first scenario, the system assigned highly absorptive materials to cover the reflective ceiling surfaces. The resulting configuration yielded an objective value 97.5% closer to the goal than the initial configuration. As desired, EDT dropped greatly to a fraction of the original value, and BR dropped significantly as well. Speech would be much more intelligible to the congregation with the resulting configuration.

For the second scenario, we also altered the materials on the walls. The system assigned highly absorptive materials to most surfaces, which improved the objective value by 98.0%. Plate 3 shows the accumulation of sound energy on the seating plane at four time steps for both the initial and final configurations. Note that in the final configuration, the majority of the sound arrives early, resulting in improved speech intelligibility. The system has in effect improved the temporal distribution of sound by increasing the percentage of the early sound and decreasing the percentage of late sound. The results are summarized in Table 8.1. Each

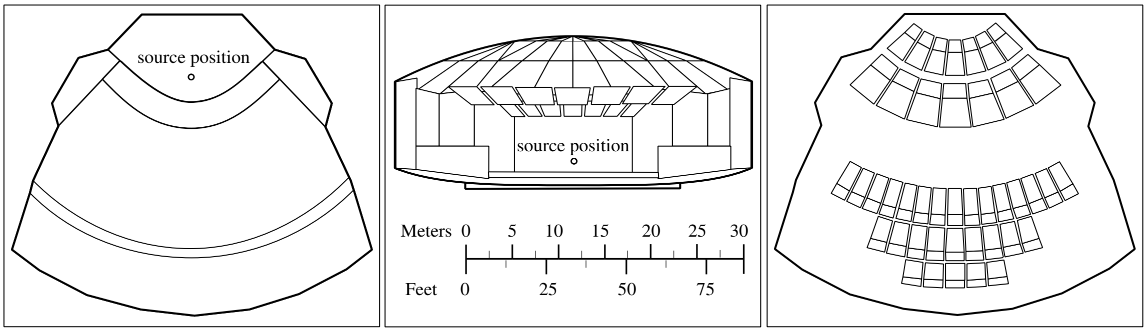

Figure 8-3: Computer model of Kresge Auditorium.

scenario ran in under three minutes.

8.2 Kresge Auditorium

The second example, Kresge Auditorium, is a multi-purpose auditorium at MIT, which is currently undergoing acoustic reevaluation (see Figure 8-3). The Institute utilizes Kresge for everything from conferences to concerts. The hall does not possess reconfigurable elements that would help accommodate such disparate acoustic requirements. Consequently, the auditorium suffers from too much reverberation for speech, although the reverberation is adequate for music, shown graphically in Plate 4 by the EDT cones. Another shortcoming is that the audience does not feel surrounded by the sound, indicated by the IACC shells, which fail to encircle the icons. The average sound level G is 7.18 dB (see Table 8.2). However, the temporal distribution of energy is poor for speech intelligibility with too much energy arriving late.

While the acoustics of the hall will never satisfy all uses without introducing reconfigurable geometry or material elements, our intent was to consider modifications that would improve the acoustics as a whole. In order to reflect this, we modified the objective function to include two sets of targets, one each for speech and symphonic music, equally weighted (see Section B.2). We painted sound level targets for both speech and symphonic music for time slices at 0.08 seconds, 0.16 seconds and total level, as shown in Plate 5. The other

	IACC	EDT (sec)	BR	G (dB)	SDI	TI (sec)	f(x)	Improvement
Speech and Symphonic Music								
Target	0.000	0.850	1.000	4.750	1.000	0.020	0.830	NA
Initial Config	0.620	2.126	0.952	7.184	0.202	0.039	3.142	NA
Materials Only Config	0.723	1.372	1.058	3.012	0.219	0.039	2.292	36.8 %
Geometry Only Config	0.539	2.047	0.958	7.227	0.202	0.028	2.890	10.9 %
Materials & Geometry	0.594	1.462	1.094	4.908	0.211	0.028	2.178	41.7 %
Speech								
Target	NA	0.700	1.000	≥ 0.000	NA	NA	0.000	NA
Initial Config	NA	2.216	0.952	7.184	NA	NA	4.419	NA
Final Config	NA	0.785	1.025	0.150	NA	NA	0.295	93.3 %
Symphonic Music								
Target	0.000	2.150	1.175	4.750	1.000	0.020	0.000	NA
Initial Config	0.620	2.126	0.952	7.184	0.202	0.039	1.866	NA
Final Config	0.590	1.882	1.025	4.365	0.209	0.028	1.543	17.3 %

Table 8.2: Acoustic measure readings for Kresge Auditorium.

targets are given in Table 8.2.

Unlike the optimization of Oakridge Chapel, which we restricted exclusively to material changes, here we allowed geometry optimizations as well. Since construction costs for geometric changes generally exceed costs for material changes, we separated the optimization into three passes, modifying only materials, only geometry, then materials and geometry combined, to compare the effectiveness of each.

Using the visualization tools, we observed the pattern of sound level accumulation as a function of time for the initial hall configuration, noting that the direct sound and the earliest reflections have the greatest effect. As variables for the first example, we selected the materials for the seats, the wall at the back of the hall, and the walls of the stage shell, since reflected sound from these surfaces reaches much of the seating area.

The optimization over materials took 72 seconds to converge, sampling 200 configurations. The system assigned absorptive material to the seats and stage floor, and it assigned reflective materials to the rear wall and remaining surfaces near the stage. The Materials Only entry in Table 8.2 shows that EDT improved substantially. While SDI improved, we did not include surface treatments that would raise SDI much beyond 0.3. TI remained un-

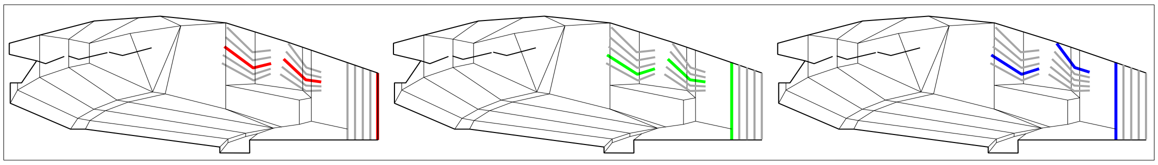

(a) Initial (b) Geometry (c) Material and geometry

Figure 8-4: Variable positions for the rear and forward bank of reflectors and the back stage wall in Kresge Auditorium for the initial (red) configuration, the geometry only (green) configuration, and the combined materials and geometry (blue) configuration.

changed since it is only affected by geometric changes.

In the second example, the geometry modifications included the depth of the center stage wall and the rotation of the two sets of suspended reflector groups above the stage area, illustrated in Figure 8-4. Each component could assume one of five positions, with the initial configuration indicated in red. These geometric components share the characteristic that the ratio between their size and the solid angle they span with respect to the sound source location is small. Further, these modifications would not require expensive alterations to the external shell of the auditorium. In the optimization over geometry, the system left the orientation of the rear reflector group unchanged, but lowered the forward reflector group, and moved the rear stage wall to the position closest to the source, as shown in green in Figure 8-4b. These modifications improved TI by 58%. The optimization took approximately 20 minutes to converge while sampling 100 geometric configurations.

The combined optimization—involving both materials and geometry—altered materials as before, but selected a new configuration for the banks of reflectors, raising the rear reflector group, lowering the forward reflector group, and again moving the rear stage wall to the position closest to the source, as shown in blue in Figure 8-4c. This configuration produced the lowest cost by maintaining the improvements to TI from geometry modifications and improvements to EDT from material modifications. The sound level G dropped to 4.9 dB, and the temporal distribution improved, with a higher percentage of the energy arriving earlier than in the initial configuration. IACC improved somewhat, but remained far from optimal, which is expected for fan shaped halls such as this one. The optimization took

17 minutes to converge while sampling 80 geometric configurations, improving the overall acoustic rating by 41.7%. The results in Table 8.2 and Plate 4 show that performance improved for both uses.

Finally, we examined speech and music objectives separately. we restricted the study to materials for the stage surfaces and the back wall of the auditorium, and used only materials that could be changed between speech and symphonic music performances, such as curtains that could be drawn, or rugs that could be taken up. The geometry configuration remained unchanged for these optimizations, as did the seat material, both set to the result from the previous optimization which combined music and speech objectives. The system assigned highly absorptive materials for the speech configuration, and materials with mid to low absorption for the music configuration. Table 8.2 and Plates 5 and 6 show the improvements for speech and music respectively at 93.3% and at 17.3%.

8.3 Jones Hall for the Performing Arts

Our third example is Jones Hall in Houston, Texas, which is a multi-purpose auditorium (see Figure 8-6). This hall is designed to accommodate different uses by modifying the position of ceiling panels. Figure 8-5 shows two possible ceiling configurations. Like many multi-purpose halls, the quality of the acoustics is less than optimal for any single use (see Table 8.3). Plate 8 shows the results of the simulation for the initial configuration of the hall using the combined objective for symphonic and operatic music. The reverberation is too strong for the combined target, as shown by the EDT cones. IACC exceeds the optimal value, shown by the narrow IACC shells. G is too low for the symphonic use. The temporal distribution of sound is good, but the quantity is too low. Of the global measures, TI is too high and SDI is too low.

We wanted to explore two directions with this case study. First, we wanted to find changes in the existing structure or materials that would improve the acoustics for both uses simultaneously. Second, we wanted to determine the effectiveness of the reconfigurable ceiling

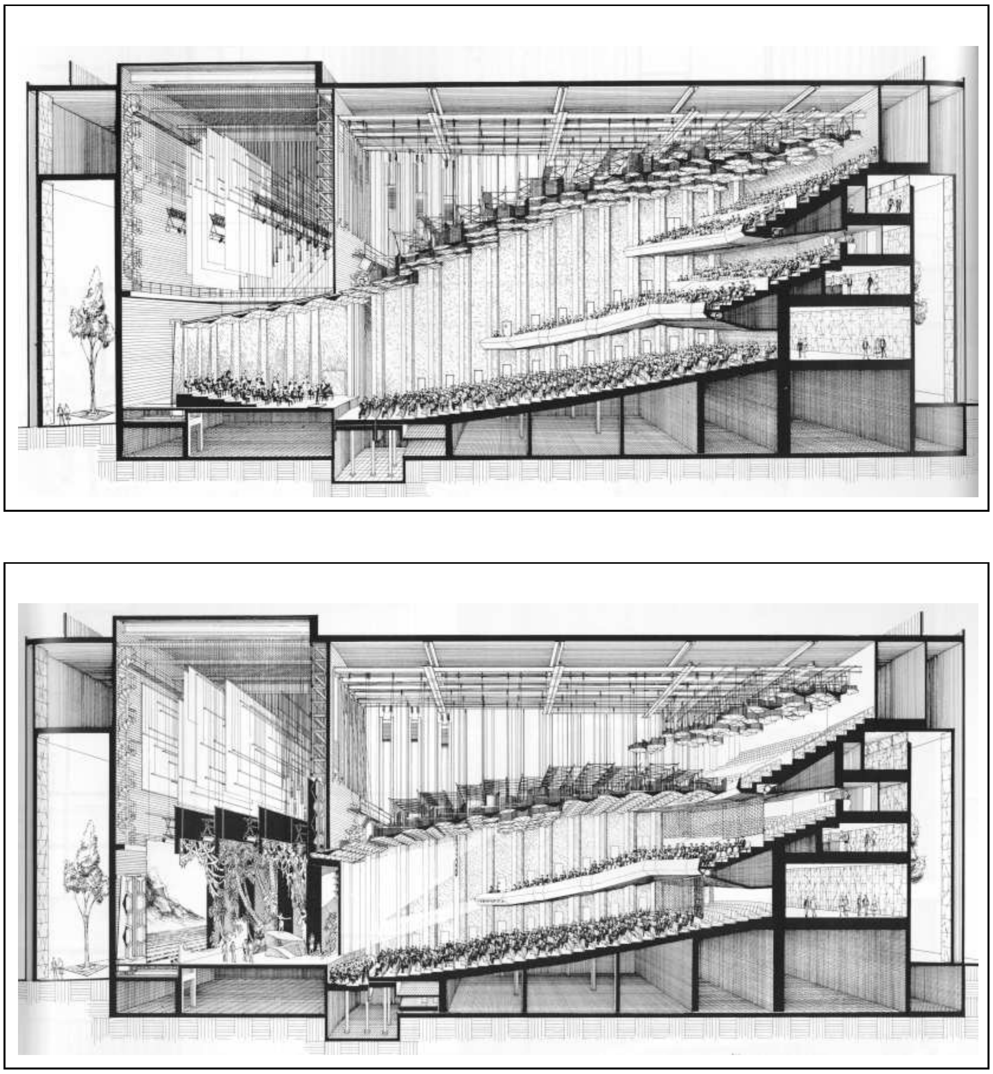

Figure 8-5: Jones Hall: illustration of hall configuration with movable ceiling panels (top) in raised position, and (bottom) in lowered position. Source of figures: [25]

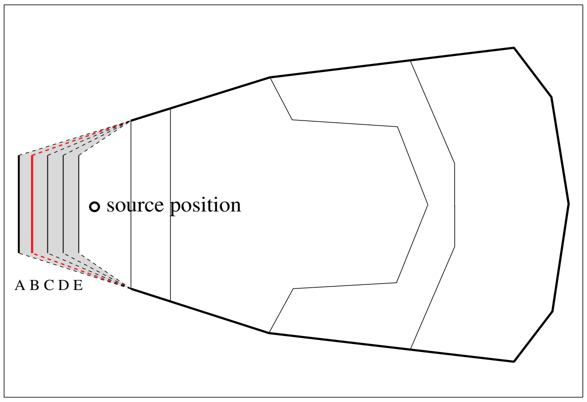

(a) Plan

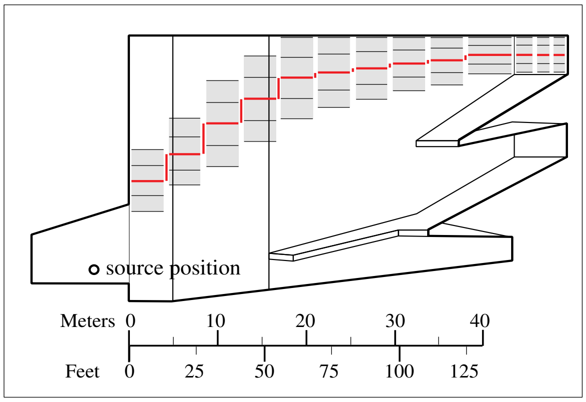

(b) Longitudinal Section

Figure 8-6: Computer model of Jones Hall. The left illustration shows the variable positions for the rear stage wall, and the right illustration shows the variable positions for the ceiling panels. The initial positions are indicated in red.

panels for two disparate target conditions, opera and symphonic music. We built two single-use objectives for these acoustic requirements, which we then combined to create a multi-use objective, as discussed in Section B.2, using the targets shown in Table B.1. Notice how opera targets differ from symphonic music targets, with the target for EDT dropping to 1.3 seconds. The BR target has a lower proportion of bass, and sound strength of 0.0 dB or greater is acceptable, with as much of the energy arriving as early as possible for increased clarity. We painted sound level targets for both symphonic music and opera for time slices at 0.16 seconds, 0.32 seconds, and total level, as shown in Plate 7. The similarity between the early and late painted targets for opera is present to increase intelligibility by maximizing early sound energy.

Optimization Using Combined Symphonic and Operatic Objective

In order to improve the overall acoustics we optimized for both symphonic and operatic uses simultaneously, keeping the ceiling panels fixed in the neutral configuration shown in Figure 8-6b. The variables we selected for this scenario were components that could be modified with the least expense, and have the greatest influence on the seating area, via the early reflections. We included one geometry modification, allowing the rear stage wall to

	IACC	EDT (sec)	BR	G (dB)	SDI	TI (sec)	f(x)	Improvement
Opera and Symphonic Music								
Target	0.000	1.450	1.000	4.750	1.000	0.020	0.411	NA
Initial Config	0.336	1.661	0.977	1.315	0.267	0.023	1.210	NA
Combined Objective	0.194	1.602	0.953	1.773	0.399	0.019	0.968	30.1 %
Operatic Music								
Target	0.000	1.300	1.000	$\geq$ 0.000	1.000	0.020	0.0	NA
Initial Config	0.336	1.661	0.977	1.315	0.267	0.023	0.980	NA
Modified Config	0.194	1.602	0.953	1.773	0.399	0.019	0.669	31.7 %
Final Config	0.187	1.314	0.942	1.726	0.399	0.019	0.363	63.0 %
Symphonic Music								
Target	0.000	2.150	1.175	4.750	1.000	0.020	0.0	NA
Initial Config	0.336	1.661	0.977	1.315	0.267	0.023	1.440	NA
Modified Config	0.194	1.602	0.953	1.773	0.399	0.019	1.267	12.0 %
Final Config	0.198	1.899	0.957	1.899	0.399	0.019	0.885	38.5 %

Table 8.3: Acoustic measure readings for Jones Hall. The modified configuration entries give the individual objective ratings for the configuration resulting from the optimization using the combined objective.

translate. Material modifications included the seats, walls, and the stage surfaces.

The system repositioned the back stage wall closer to the audience (from position B to position D in Figure 8-6a,) forcing the side stage walls to reflect more sound out to the audience. The system assigned highly reflective materials to the surfaces surrounding the stage, and a less absorptive material to the seats. These adjustments significantly reduced TI and IACC. Table 8.3 shows an overall improvement of 30.1% for the combined objective. The optimization took under ten minutes.

Optimization Using Individual Symphonic and Operatic Objectives

We then optimized the ceiling panel positions for each use separately. Starting from the best configuration resulting from the simultaneous optimization, we selected only the ceiling panels as variables, allowing them to translate vertically.

Figure 8-7 shows the best solution found by our system for each objective. Given the symphonic objective, the system raised the panels from their neutral positions to the configuration shown in Figure 8-7a. Table 8.3 shows that both EDT and G improved, although

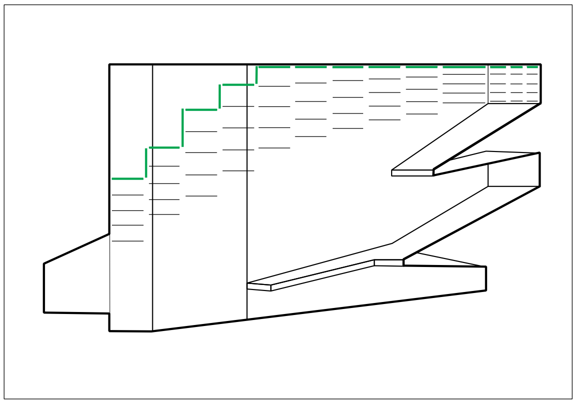

(a) Symphonic Music

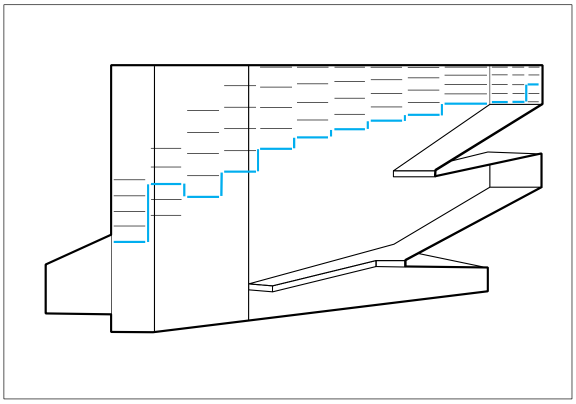

(b) Operatic Music

Figure 8-7: Resulting ceiling panel configurations for Jones Hall: left: symphony configuration (green,) right: opera configuration (blue.)

G remained quite weak, at only 1.9 dB. The system achieved an additional 26.5% improvement for the symphonic objective. Using the operatic objective, the system assigned panels to the positions shown in Figure 8-7b. EDT was the most improved measure, dropping 0.3 seconds, very close to the desired target. IACC improved slightly as well. The system achieved further improvements of 31.3% with these modifications. These optimizations took under 10 minutes each.

Chapter 9

Discussion and Future Work

We have developed an interactive acoustic design system which addresses the major limitations present in today's computer-aided design systems. The tedious sequence of modeling/simulation steps required by the user of a direct design system is replaced by the automated design-optimization approach characterized by inverse design, shifting much of the burden of design space searching from the designer to the computer. In our design system the user interactively specifies desired acoustic performance goals and describes material and geometric variations and constraints for a collection of architectural components in the scene. The system then searches the feasible space for the configuration that "best" meets the specification, using a combination of simulated annealing and steepest descent techniques. Our new acoustic simulation algorithm features increased geometric accuracy and decreased computation time, enhancing the overall performance and usability of our design system. We present the resulting configuration and its time dependent sound field characteristics using visualization tools; these tools are also used to express design goals. This approach can be easier and more intuitive to use than the usual direct edit-simulate cycle.

Our experiences with the system suggest a number of theoretical and practical extensions that might be investigated to enhance the system, and support the design process.

- *Application to preliminary design.* While I have tested the system extensively on existing halls, I am eager to explore the system's utility as a tool in the design of a brand new hall. Because the acoustic design problem is so difficult, designers often feel compelled to use a familiar or proven geometric configuration in order to avoid a potentially costly mistake. Our system seems amenable to assisting in the preliminary design phase—perhaps enabling a designer to not only consider a wide range of possible designs but also to gain new insight into the intricacies of sound propagation.

- *Acoustic evaluation.* Currently the sound field is sampled at a uniform grid of points, each contributing equally to the evaluation of the hall. There may be better ways of combining the sample data, perhaps giving more weight to central seating area locations. Also, acoustic measures could be explored that better reflect the overall quality of the acoustics. For example, an echo metric might be used to automate the search for poor time distributions of energy. It might be interesting to produce a scalar field over the seating areas that represents objective function values. This would indicate the best and worst audience seating areas with respect to acoustics. The acoustic objective could be augmented with viewing objectives as well, producing a "best value" scalar field.

- *Acoustic simulation and visualization.* The area of accurate acoustic simulation is an important one, with a variety of challenges remaining, such as diffraction and other wave effects. Funkhouser has described a method for extending the geometric beam tracing algorithm to model diffraction [19]. Modal analysis might prove useful for locating dead spots caused by the destructive interference of sound waves. Volumetric representations of the sound field might also improve the user's understanding of the space but would require a voxel representation of the 3D scene. In order to increase simulation speed, the model of the space might be represented at various levels of detail, corresponding to the frequencies of simulated sound. With this approach, far fewer beams would be generated for low frequency sound.

- *Auralization.* Some systems play back a prerecorded anechoic sound sample convolved with the sound field signature at a specific location within the hall, permitting the listener to hear what a hall might sound like at that location. This approach could be utilized and extended in our system as follows. The listener could set acoustic targets by modifying the sound field signature interactively while listening to the resulting signal, using graphical tools like sliders or dials hooked to specific sound field characteristics. While one can only listen to the sound produced at a single location in the hall at a time, user-specified targets could be allowed to influence regions of any size. Variable weights could be applied to a region of influence, allowing a feathering effect between targets, or interpolation between them.

A related target specification tool might allow the user to select two predefined sound samples, then adjust a slider that interpolates the samples, selecting an intermediate point with the best sound. The characteristics of the selected sound would then be used as performance goals. This technique could focus on specific characteristics of sound, or consider all characteristics at once.

- *Optimization and design.* While the advent of computer-aided design systems has brought designers tools that assist in predicting and visualizing complex phenomena such as light and sound, it is clear that the computer could play a more significant role in the design process. One intriguing area for future work is in visualizing the optimization process itself—providing the designer with an intuitive representation of the multidimensional search space and the ability to steer the optimization process. Perhaps the greatest shortcoming of the current system is that it cannot directly identify the components that have the strongest adverse effect on acoustic performance. Such variable sensitivity analysis, in addition to boundary analysis and application-specific heuristics based on the insights of experienced acoustical engineers, might be beneficial. Hierarchical grouping of components might allow for faster optimization approaches based on progressive refinement techniques. Finally, the *Design Galleries* [36] approach and our optimization approach might be combined such that the dispersion phase of DG selects configurations that meet not only a difference criteria, but that also satisfy evaluation requirements as well.

Appendix A

Acoustic Measure Calculations

The following pseudo-code segments illustrate our implementation for each acoustic measure calculation. The input data used in these calculations are gathered from two simulation algorithms, over the five frequency bands: 125 Hz, 250 Hz, 500 Hz, 1 kHz, and 2 kHz. The beam tree provides data for the early sound, recording the time, intensity, and direction of each passing wavefront in a hit list unique to each listener point in the hall. The reverberant tail provides non-directional data for the later sound. The same data is used for all points in the hall. The functions that calculate acoustic measures have access to the beam tree data via the helper functions **first(list)**, **last(list)**, **next(hit)** and **prev(hit)**, which return the appropriate hit from the list. Tail energy contributions are accessed using the functions **get_tail_energy**, which returns total energy, and **get_tail_energy_interval**, which returns the energy contribution during a specified interval of time, for a given frequency. In addition, the surface diffusivity calculation has access to the absorption characteristics and surface area data for each surface in the enclosure.

- *Interaural Cross-Correlation Coefficient (IACC).* The function **get_IACC** considers the wavefronts that pass the listener in the first 80 msec interval after the arrival of the direct sound. The angle of impact is defined as the angle between the direction of the incoming wave and the angle perpendicular to the view direction on the horizontal plane. Notice that the contribution of each wave to this measure is proportional to the square of the intensity of the wave, which means that early waves have much greater impact than later waves. In fact, this measure can be reasonably approximated by considering only the earliest waves. Due to the directional dependence of this measure, only the beam data is used.

```
function get_IACC ()
    for each frequency(500 Hz, 1000 Hz, 2000 Hz) do
        initialize llSum, rrSum, and lrSum to 0.0
        for each hit within 80 msec. of direct hit do
            determine angle of incoming sound
            use lookup to set correlation fractions ll, rr, lr based on angle
            increment llSum by intensity^2 * ll
            increment rrSum by intensity^2 * rr
            increment lrSum by intensity^2 * lr
        end do
        set IACC(frequency) to lrSum / sqrt(llSum * rrSum)
    end do
    set IACC to the average of IACC(500 Hz, IACC(1000 Hz), and IACC(2000 Hz)
end
```

- *Early Decay Time (EDT).* The function **get_EDT** calculates the time required for the sound to decay 10 dB. The standard method used to calculate decay is the following. First, calculate the total energy arriving at a listener position, and convert it to sound level. Next, determine the target level, which in this case is 10 dB below the total level. Finally, integrate the sound energy in reverse order, stopping when the accumulated energy attains the target level. The interval between the arrival of the direct sound and the time when the target level is reached is used to calculate the decay measure. By integrating in this way, the sound level increases monotonically, which guarantees that the target level is attained at exactly one point in time.

  One alternate method, which I rejected, does not offer this guarantee, and would proceed as follows. Compare the energy of each reverberant wave in the order of arrival to that of the direct sound, stopping when the difference in level exceeds 10 dB. The problem with this approach is illustrated by considering the early arrival of a very weak wave, reflected from a highly absorbent surface. If its sound level was at least

10 dB below that of the direct sound, the algorithm would terminate, and the resulting interval of time would be uncharacteristically short. Instead, by integrating energy in reverse order, the effects of large variance in the intensity of passing waves is removed, providing a more stable result.more stable result.

The helper function **get_decay_time** takes as input the target level at the frequency of interest, and returns the time that the target level is attained. The early time interval is modeled by beam tracing and the later time interval is modeled by the tail approximation. Due to the statistical nature of the tail, these intervals may overlap. Since the sound energy needs to be integrated in reverse order, the algorithm must alternate between these two data sources during an overlapping interval. The reverberant tail is stored as an array, where the accumulated energy from each successive hit constitutes an entry. By accumulating energy in this way, the contribution of tail energy can be determined in constant time for any interval, by simply taking the difference of energy at the end points of the interval. An internal function takes a point in time as input and calculates the corresponding tail array index.

The integration algorithm works as follows. Determine the energy from the part of the tail that extends beyond the last beam tree hit, if any. If the beam hit list and tail overlap, alternate between them, successively adding tail energy and beam energy. Once the tail interval is passed, continue adding beam energy, if any, until the target level is attained. If the target level is exceeded within a tail interval, then recursively subdivide the interval until the target level is attained. Otherwise, return the time when a beam energy pushes the level to or above the target level.

```
function get_EDT
    for each frequency(500 Hz, 1000 Hz) do
        set total_energy to get_total_energy (frequency)
        set actual_level to convert_energy_to_level (total_energy)
        set target_level to actual_level - 10.0 dB
        set EDT(frequency) to get_decay_time (target_level, frequency)
    end for
    set EDT to the average of EDT(500 Hz) and EDT(1000 Hz)
end
```

```
function get_decay_time (target_level, frequency)
    initialize accumulated_level to 0.0
    set hit to last(hit_list)
    set hit_time to time_of_hit (hit)
    initialize hit_energy to 0.0
    initialize tail_energy to 0.0
    while (accumulated_level ≤ target_level) do
        Integrate energy in reverse order
        set tail_energy to get_tail_energy_interval (hit_time, frequency)
        set accumulated_level to get_level(tail_energy + hit_energy)
        if (accumulated_level > target_level)
        then return find_intermediate_time ()
        increment hit_energy by energy(hit, frequency)
        set accumulated_level to get_level(tail_energy + hit_energy)
        if (accumulated_level > target_level)
        then return hit_time
        set hit to prev (hit)
        set hit_time to time_of_hit (hit)
    end while
    return hit_time
end
```

```
function get_total_energy (frequency)
    set energy to get_tail_energy (frequency)
    for each hit in hit_list do
        increment energy by energy (hit, frequency)
    return energy
end
```

- *Bass Ratio (BR).* The function **get_BR** requires the calculation of RT at four frequency bands: 125 Hz, 250 Hz, 500 Hz, and 1000 Hz, where RT is defined as the difference between 35 dB and 5 dB decay times. The function returns the ratio of low to mid frequency decay measures. As with EDT, the helper function **get_decay_time** is used.

```
function get_BR
    for each frequency(125 Hz, 250 Hz, 500 Hz, 1000 Hz) do
        set total_energy to get_total_energy (frequency)
        set actual_level to convert_energy_to_level (total_energy)
        set target_level to actual_level - 5.0 dB
        set T5(frequency) to get_decay_time (target_level, frequency)
        set actual_level to convert_energy_to_level (total_energy)
        set target_level to actual_level - 35.0 dB
        set T35(frequency) to get_decay_time (target_level, frequency)
        set T30(frequency) to T35(frequency) - T5(frequency)
    end for
    set BR to (T30(125) + T30(250)) / (T30(500) + T30(1000))
end
```

- *Strength Factor (G).* The function **get_Gmid** converts the total accumulated energy at a receiver point to a strength level for the frequency bands of 500 Hz and 1000 Hz, and returns the average strength level. However, in this system the designer is allowed to specify strength targets through time. Strength levels are calculated at each target time point.

```
function get_Gmid
    for each frequency(500 Hz, 1000 Hz) do
        set total_energy to get_total_energy (frequency)
        set G(frequency) to convert_energy_to_strength (total_energy)
    end for
    set Gmid to the average of G(500 Hz) and G(1000 Hz)
end

function convert_energy_to_strength (total_energy)
    level of ratio of total energy to energy at 10 meters from source
    set energy_at_10_meters to initial_energy / 100.0
    return 10.0 * log (total_energy/energy_at_10_meters)
end
```

- *Initial-Time-Delay Gap (TI).* While this measure could be calculated for any point in the enclosure, standard practice dictates that a single representative point be used. The function **get_TI** returns the time between the arrival of the direct sound and that of the first reverberant wave in the hit list.

```
function get_TI
    uses representative point in audience
    set hit to first (hit_list)
    set T_0 to time (hit)
    set hit to next (hit)
    set T_1 to time (hit)
    return (T_1 - T_0)
end
```

- *Surface Diffusivity Index (SDI).* The function **get_SDI** returns the average of diffusion for all surfaces in the enclosure, weighted by surface area. Each material in the material library has an average diffusion coefficient associated with it, which is used in this calculation.

```
function get_SDI
    set total_area to 0.0
    set total_diffusion to 0.0
    for each surface in enclosure do
        increment total_area by area(surface)
        increment total_diffusion by area(surface) * diffusion_index(surface)
    end for
    set SDI to total_diffusion / total_area
end
```

Appendix B

Objective Function Construction

This section describes how the objective function presented in Section 6.4 specifically designed for symphonic music can be modified with performance goals that are appropriate for various other uses of the hall. I follow this with a discussion of how to construct an objective to handle multiple performance types simultaneously, such as symphonic music and opera.

B.1 Single Use Objectives

The target values of certain acoustic measures for opera and speech differ from those of symphonic music, as shown in Table B.1 [8, 7]. In the general case, new objective function terms are produced simply by providing the appropriate target value for a given use. For example, to construct the term for BR in the case of a speech objective, one simply inserts 1.0 for BR_{target}, as shown in the table. In some cases, however, the nature of the objective curve differs between uses. In the case of symphonic music, G_{target} is a single value of 4.5 dB, whereas G_{target} for speech covers a range of values, with no cost imposed for values above 0.0 dB.

	IACC	EDT (sec)	BR	G (dB)	TI (sec)	SDI
Symphonic Music (SM)	0.0	2.125	1.175	4.75	0.02	1.0
Operatic Music (OM)	0.0	1.300	1.000	≥ 0.0	0.02	1.0
Speech (S)	NA	0.700	1.000	≥ 0.0	NA	NA
SM & S	0.0	0.850	1.000	4.75	0.02	1.0
SM & OM	0.0	1.450	1.000	4.75	0.02	1.0

Table B.1: Acoustic measure targets for predefined objectives.

B.2 Multiple Use Objectives

A single hall must often accommodate a number of disparate uses, each of which has its own set of acoustic requirements. I have extended the single use objective, described in Chapter 4, to account for a number of uses simultaneously. Multiple use objective functions are constructed by linearly combining individual objective functions as follows:

$$f_{combined}(x) = a_1 * f_{use_1}(x) + a_2 * f_{use_2}(x) + \cdots + a_n * f_{use_n}(x), \tag{B.1}$$

where a_n is the weighting factor of the nth individual objective function, and the sum of the weighting factors is 1.0.

When functions are combined, the target for an acoustic measure may be found by linearly combining the like terms of each individual objective function, and locating the minimum point or range in the resulting curve. This minimum value will be greater than zero unless there is an acoustic value for which each objective curve is zero. EDT and BR have non-zero minimum costs. These minimum costs are lower bounds on the objective function for the individual measures. The lower bound for the entire objective function is the sum of lower bounds for each term.

In each of the examples of simultaneous optimization presented in this thesis the objective function combines two individual use functions, equally weighted. When combining symphonic music and speech, the targets for BR and EDT differ from individual targets, at 1.0 and 0.85 seconds respectively. The costs associated with these values are computed as follows:

$$
\begin{aligned}
\textit{lower bound cost}_{\mathrm{BR}} &= \left|\log\left(\frac{1.0}{1.175}\right)\right|^{3/2} * 10.0 * 0.5 = 0.093,\\
\textit{lower bound cost}_{\mathrm{EDT}} &= \left|\log\left(\frac{0.85}{2.15}\right)\right|^{3/2} * 9.0 * 0.5 = 1.151,\\
\textit{lower bound } f(x) &= 0.093 + 1.151 = 1.244.
\end{aligned}
$$

When combining symphonic music and speech, the targets for BR and EDT differ from individual targets, at 1.0 and 1.45 seconds respectively. The costs associated with these values are computed as follows:

$$
\begin{aligned}
\textit{lower bound cost}_{\mathrm{BR}} &= \left|\log\left(\frac{1.0}{1.175}\right)\right|^{3/2} * 10.0 * 0.5 = 0.093,\\
\textit{lower bound cost}_{\mathrm{EDT}} &= \left|\log\left(\frac{1.45}{2.15}\right)\right|^{3/2} * 9.0 * 0.5 = 0.318,\\
\textit{lower bound } f(x) &= 0.093 + 0.318 = 0.411.
\end{aligned}
$$

The lower bound on the objective for G remains 0.0 in both cases, since no cost is levied for values above 0.0 dB for speech or opera, nor between 4.0 dB and 5.0 dB for symphonic music. Therefore the combined target spans the range between 4.0 dB and 5.0 dB.

Appendix C

Material Library

Material	Absorption Coefficients				
	125.0 Hz.	250.0 Hz.	500.0 Hz.	1000.0 Hz.	2000.0 Hz.
Seats, Heavy	0.72	0.79	0.83	0.84	0.83
Seats, Medium	0.56	0.64	0.70	0.72	0.68
Seats, Light	0.35	0.45	0.57	0.61	0.59
Carpet on Concrete	0.02	0.06	0.14	0.37	0.60
Carpet, Heavy, on Foam	0.08	0.24	0.57	0.69	0.71
Wood Parquet	0.04	0.04	0.07	0.06	0.06
Linoleum	0.02	0.03	0.03	0.03	0.03
Concrete Block, Unpainted	0.10	0.01	0.02	0.02	0.02
Concrete, Painted	0.10	0.05	0.06	0.07	0.09
Cork on Brick or Concrete	0.02	0.03	0.03	0.03	0.02
Plaster	0.01	0.01	0.02	0.03	0.04
Gypsum over 2x4, 16 o.c.	0.29	0.10	0.05	0.02	0.02
Plywood Paneling	0.28	0.22	0.17	0.09	0.10
Brick, Unglazed	0.03	0.03	0.03	0.04	0.05
Soundblox 6 in. Type A	0.61	0.83	0.36	0.43	0.27
Soundblox 6 in. Type R	0.38	0.99	0.64	0.57	0.43
1 in. Decoustic Panel, Type 4 Mt	0.14	0.36	0.83	0.99	0.99
2 in. Decoustic Panel, Type 4 Mt	0.44	0.80	0.99	0.99	0.96
1 in. Tectum, Type 4 Mt	0.07	0.11	0.23	0.43	0.70
2 in. Tectum, Type 4 Mt	0.11	0.20	0.47	0.80	0.61
3 in. Tectum, Type 4 Mt	0.18	0.34	0.81	0.76	0.75
K-13 .625 in. Spray On	0.05	0.15	0.43	0.79	0.90
K-13 1.0 in. Spray On	0.08	0.28	0.75	0.98	0.93
K-13 1.5 in. Spray On	0.25	0.44	0.86	0.99	0.98
K-13 2.0 in. Spray On	0.40	0.61	0.93	0.99	0.95
Fiberglass 703, 1 in.	0.06	0.20	0.64	0.90	0.95
Fiberglass 703, 3 in.	0.52	0.99	0.99	0.99	0.99
Fiberglass 703, 4 in.	0.99	0.99	0.99	0.99	0.99
Velour, 18 oz, Half Fullness	0.14	0.35	0.55	0.72	0.70
Drapery, Light	0.03	0.04	0.11	0.17	0.24
Drapery, Medium	0.07	0.31	0.49	0.75	0.70
Partition, Accordian	0.33	0.20	0.15	0.10	0.05

Table C.1: An example of a Material Library.

Appendix D

Objective Function Values for Case Studies

This chapter contains specific information about the variables selected for each scenario of the three case studies, including the set of materials assigned to each material variable, the resulting acoustic measures, and the individual and total objective values associated with each solution. In each scenario, the initial temperature variable used in the simulated annealing algorithm was initially set to 0.5, and was scheduled to decrease by 20% after testing each set of 20 configurations. The neighborhood variable was initially set to equal the maximum number of items assigned to any optimization variable. The neighborhood variable was also scheduled to decrease by 20% after testing each set of 20 configurations.

D.1 Oakridge Bible Chapel

In this section I list the materials for the surfaces within the Oakridge Bible Chapel, including the initial materials as well as those materials assigned by the design system during optimization. Initially all walls within the Chapel are plaster, and the seats are composed of a lightly absorptive material. Each surface that is selected as a variable during optimization is assigned one material from the following set: { Plaster, Gypsum over 2x4: 16 o.c., Fiber-

Ceiling Surface Variable	Material
far left	Soundblox 6 in. Type A
near left	Soundblox 6 in. Type A
near right	Fiberglass 703, 4 in.
far right	Soundblox 6 in. Type A

Table D.1: Ceiling variable assignments for Oakridge Bible Chapel.

glass 703: 4 in., Velour: 18 oz. Half Fullness, Soundblox 6 in. Type A, Soundblox 6 in. Type R, 2 in. Decoustic Panel: Type 4 Mt, K-13 2.0 in. Spray On, 2 in. Tectum, Type 4 Mt}.

The first optimization for Oakridge Bible Chapel included four ceiling surfaces as material variables. The system selected the materials shown in Table D.1. Orientation information is given for the viewer standing at the altar, facing the rear of the room. The final optimization for Oakridge Bible Chapel included four ceiling surfaces and 13 wall surfaces as material variables. The system selected the materials shown in Table D.2. Table D.3 shows the acoustic measures and associated objective function values for each measure under each condition.

D.2 Kresge Auditorium

I ran a total of five scenarios for Kresge Auditorium. In the first three scenarios, I used an objective that combined music and speech, using only material variables in the first scenario, followed by just geometry variables, then finishing the set by combining material and geometry variables into a single optimization. The final two scenarios I optimized for music and speech separately, using a selected set of material variables only.

The initial materials for Kresge Auditorium are shown in Table D.4. Orientation information is given for the viewer standing on stage, facing the rear of the auditorium. The set of materials for all variables except the seats includes the following: { Carpet on Concrete, Carpet: Heavy: on Foam, Wood Parquet, Concrete Block: Unpainted, Concrete: Painted,

Ceiling Surface Variable	Material
far left	Gypsum over 2x4, 16 o.c.
near left	Soundblox 6 in. Type A
near right	Gypsum over 2x4, 16 o.c.
far right	Gypsum over 2x4, 16 o.c.
Wall Surface Variable	Material
side left	Fiberglass 703, 4 in.
side right	Velour, 18 oz, Half Fullness
rear far left	2 in. Decoustic Panel
rear far right	Fiberglass 703, 4 in.
rear near left	Velour, 18 oz, Half Fullness
rear near right	Gypsum over 2x4, 16 o.c.
rear center	2 in. Tectum, Type 4 Mt
front far left	Soundblox 6 in. Type A
front far right	Soundblox 6 in. Type A
front mid left	2 in. Tectum, Type 4 Mt
front mid right	Fiberglass 703, 4 in.
front near left	Soundblox 6 in. Type R
front near right	Soundblox 6 in. Type A

Table D.2: Wall and ceiling variable assignments for Oakridge Bible Chapel.

	EDT (sec)	BR	G (dB)	f(x)
Initial Configuration				
Target	0.700	1.000	≥ 0.0	0.000
Acoustic Measures	1.899	1.631	12.141	
Objective Values	3.422	0.979	0.434	4.835
Ceiling Variables				
Acoustic Measures	0.784	1.002	7.760	
Objective Values	0.000	0.002	0.118	0.120
Wall and Ceiling Variables				
Acoustic Measures	0.651	1.026	6.779	
Objective Values	0.003	0.011	0.083	0.098

Table D.3: Acoustic measure readings and objective function values for Oakridge Bible Chapel.

Surface Variable	Material
stage walls	Plywood Paneling
stage floor	Wood Parquet
seats	Seats, Medium
main rear wall	Concrete Block, Unpainted

Table D.4: Initial materials for Kresge Auditorium.

Surface Variable	Material
left stage wall	Cork on Brick or Concrete
right stage wall	K-13 2.0 in. Spray On
rear stage wall	Soundblox 6 in. Type R
stage floor	Concrete Block, Unpainted
seats	Seats, Heavy
main rear wall	K-13 2.0 in. Spray On

Table D.5: Variable assignments for the 'Material Only' optimization of Kresge Auditorium.

Cork on Brick or Concrete, Plywood Paneling, Brick: Unglazed, Partition: Accordian, Fiberglass 703: 1 in., Fiberglass 703: 4 in., Velour: 18 oz: Half Fullness, Soundblox 6 in. Type A, Soundblox 6 in. Type R, 1 in. Decoustic Panel: Type 4 Mt, 2 in. Tectum: Type 4 Mt, K-13 .625 in. Spray On, K-13 1.0 in. Spray On, K-13 2.0 in. Spray On, Drapery: Light }.

The 'Materials Only' optimization for Kresge Auditorium included six surfaces as material variables. The system selected the materials shown in Table D.5. The 'Geometry Only' optimization did not include material variables, so no tabular information is included here. The 'Materials and Geometry' optimization for Kresge Auditorium included six surfaces as material variables. The system selected the materials shown in Table D.6.

Information for the final two scenarios using individual music and speech objectives follow. The 'Music' optimization for Kresge Auditorium included five surfaces as material variables. The system selected the materials shown in Table D.6. The 'Speech' optimization for Kresge Auditorium included five surfaces as material variables. The system selected the materials shown in Table D.8.

Surface Variable	Material
left stage wall	Fiberglass 703, 1 in.
right stage wall	Velour, 18 oz, Half Fullness
rear stage wall	Carpet on Concrete
stage floor	Concrete, Painted
seats	Seats, Heavy
main rear wall	K-13 2.0 in. Spray On

Table D.6: Variable assignments for the 'Material and Geometry' optimization of Kresge Auditorium.

Surface Variable	Material
left stage wall	Drapery, Light
right stage wall	K-13 1.0 in. Spray On
rear stage wall	Partition, Accordian
stage floor	Soundblox 6 in. Type A
main rear wall	Brick, Unglazed

Table D.7: Variable assignments for the 'Music' optimization of Kresge Auditorium.

Surface Variable	Material
left stage wall	Fiberglass 703, 4 in.
right stage wall	2 in. Tectum, Type 4 Mt
rear stage wall	K-13 .625 in. Spray On
stage floor	Fiberglass 703, 1 in.
main rear wall	1 in. Decoustic Panel, Type 4 Mt

Table D.8: Variable assignments for the 'Speech' optimization of Kresge Auditorium.

	IACC	EDT (sec)	BR	G (dB)	SDI	TI (sec)	f(x)
Combined Speech and Symphonic Music							
Target	0.000	0.850	1.000	4.750	1.000	0.020	0.830
Initial Configuration	0.620	2.126	0.952	7.184	0.202	0.039	
Objective Values	0.293	2.010	0.154	0.285	0.289	0.111	3.142
Materials Only	0.723	1.372	1.058	3.012	0.219	0.039	
Objective Values	0.369	1.333	0.068	0.143	0.267	0.111	2.292
Geometry Only	0.539	2.047	0.958	7.227	0.202	0.028	
Objective Values	0.238	1.912	0.145	0.275	0.289	0.039	2.890
Materials and Geometry	0.594	1.462	1.094	4.908	0.211	0.028	
Objective Values	0.274	1.393	0.066	0.128	0.277	0.039	2.178
Speech							
Target	NA	0.700	1.000	$\geq$ 0.000	NA	NA	0.000
Initial Configuration	NA	2.216	0.952	7.184	NA	NA	
Objective Values	NA	4.020	0.031	0.368	NA	NA	4.419
Final Configuration	NA	0.753	1.033	0.079	NA	NA	
Objective Values	NA	0.175	0.018	0.103	NA	NA	0.295
Symphonic Music							
Target	0.000	2.150	1.175	4.750	1.000	0.020	0.000
Initial Configuration	0.620	2.126	0.952	7.184	0.202	0.039	
Objective Values	0.586	0.000	0.277	0.203	0.222	0.578	1.866
Final Configuration	0.590	1.882	1.025	4.365	0.209	0.028	
Objective Values	0.544	0.104	0.144	0.112	0.560	0.080	1.543

Table D.9: Acoustic measure readings and objective function values for Kresge Auditorium.

Surface Variable	Material
walls	Plywood Paneling
stage floor	Wood Parquet
seats	Seats, Heavy

Table D.10: Initial materials for Jones Hall.

Surface Variable	Material
main room walls	Partition, Accordian
stage side walls	Gypsum over 2x4, 16 o.c.
stage floor	Cork on Brick or Concrete
stage ceiling	Partition, Accordian
rear stage wall	Partition, Accordian
seats	Seats, Heavy

Table D.11: Variable assignments for the combined objective for Jones Hall.

D.3 Jones Hall

This section provides information about the optimization of Jones Hall. I ran three scenarios, first using an objective that combined symphonic music and opera, followed by individual objectives for each use. I included both material and geometry variables in first scenario, but restricted the last two scenarios to geometry only.

Table D.10 shows the initial materials for Jones Hall. The material variables for the combined objective included the walls {Plaster, Plywood Paneling, Brick: Unglazed, Gypsum over 2x4: 16 o.c., Partition: Accordian, Fiberglass 703: 3 in., 3 in. Tectum: Type 4 Mt, K-13 1.5 in. Spray On, Drapery, Medium }, stage floor { Carpet on Concrete, Linoleum, Wood Parquet, Cork on Brick or Concrete, Concrete: Painted, Carpet: Heavy: on Foam }, and the seats { Seats: Heavy, Seats: Medium, Seats: Light }. From these, the system selected the materials shown in Table D.11. The last two scenarios contained no material variables, and consequently no tabular information is provided.

	IACC	EDT (sec)	BR	G (dB)	SDI	TI (sec)	f(x)
Combined Opera and Symphonic Music							
Target	0.000	1.450	1.000	4.750	1.000	0.020	0.411
Initial Configuration	0.336	1.661	0.977	1.315	0.267	0.023	
Objective Values	0.161	0.377	0.119	0.102	0.433	0.017	1.210
Final Configuration	0.194	1.602	0.953	1.773	0.399	0.019	
Objective Values	0.083	0.360	0.153	0.118	0.252	0.002	0.968
Operatic Music							
Target	0.000	1.300	1.000	$\geq$ 0.000	1.000	0.020	0.0
Initial Configuration	0.336	1.661	0.977	1.315	0.267	0.023	
Objective Values	0.080	0.412	0.010	0.026	0.433	0.017	0.980
Final (Sim. Opt)	0.194	1.602	0.953	1.773	0.399	0.019	
Objective Values	0.042	0.303	0.031	0.040	0.252	0.002	0.669
Final (+Ceiling Panels)	0.187	1.314	0.942	1.726	0.399	0.019	
Objective Values	0.039	0.003	0.041	0.027	0.252	0.002	0.363
Symphonic Music							
Target	0.000	2.150	1.175	4.750	1.000	0.020	0.0
Initial Configuration	0.336	1.661	0.977	1.315	0.267	0.023	
Objective Values	0.241	0.341	0.228	0.179	0.433	0.017	1.440
Final (Sim. Opt)	0.194	1.602	0.953	1.773	0.399	0.019	
Objective Values	0.125	0.417	0.275	0.196	0.252	0.002	1.267
Final (+Ceiling Panels)	0.198	1.899	0.957	1.899	0.399	0.019	
Objective Values	0.126	0.083	0.270	0.156	0.252	0.002	0.885

Table D.12: Acoustic measure readings and objective function values for Jones Hall.

Bibliography

[1] Emile Aarts and Jan Korst. *Simulated Annealing and Boltzman Machines*. John Wiley & Sons, New York, 1989.

[2] Jont B. Allen and David A Berkley. Image method for efficiently simulation small-room acoustics. *Journal of the Acoustical Society of America*, 65(4):943–950, 1979.

[3] Allred and Newhouse. Applications of the monte carlo method to architectural acoustics. *Journal of the Acoustic Society of America*, 30(10):903–904, 1958.

[4] John Amanatides. Ray tracing with cones. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, 3, pages 129–135, July 1984.

[5] Yoichi Ando. *Concert Hall Acoustics*. Springer-Verlag, Berlin, Heidelberg, New York, Tokyo, 1985.

[6] Robert Azencott. *Simulated Annealing, Parallelization Techniques*. John Wiley and Sons, Inc., Paris, 1992.

[7] Michael Barron. *Auditorium Acoustics and Architectural Design*. E and FN Spon, London, 1993.

[8] L. Beranek. *Acoustics*. Acoustical Society of America, 1993.

[9] Leo L. Beranek. *Music, Acoustics & Architecture*. John Wiley & Sons, Inc., New York, London, 1962.

[10] Leo L. Beranek. *Concert and Opera Halls: How they Sound*. Acoustical Society of America, New York, 1996.

[11] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, Massachusetts, 1995.

[12] J. Borish. Extension of the image model to arbitrary polyhedra. *Journal of the Acoustic Society of America*, 75:1827–1836, 1984.

[13] Bose Corporation. Auditioner audio demonstrator technology. http://www.bose.com/, November 1998.

[14] Richard P. Brent. *Algorithms for Minimization without Derivatives*. Prentice-Hall, Englewood Cliffs, New Jersey, 1973.

[15] N. Dadoun, D. G. Kirkpatrick, and J. P. Walsh. Hierarchical approaches to hidden surface intersection testing. *Graphics Interface '82*, pages 49–56, 1982.

[16] N. Dadoun, D. G. Kirkpatrick, and J. P. Walsh. The geometry of beam tracing. In *Proceedings of the Symposium on Computational Geometry*, pages 55–61, May 1985.

[17] Bengt-Inge Dalenback, Mendel Kleiner, and Peter Svensson. A macroscopic view of diffuse reflection. *Journal of the Audio Engineering Society*, 42(10):793–807, 1994.

[18] M. David Egan. *Concepts in Architectural Acoustics*. McGraw-Hill Book Company, New York, 1972.

[19] T. Funkhouser, I Carlbom, G Elko, G Pingali, M Sondhi, and J West. A beam tracing approach to acoustic modeling for interactive virtual environments. In *Computer Graphics Proceedings, Annual Conference Series, 1998*, pages 21–32, July 1998.

[20] David Halliday and Robert Resnick. *Physics*. John Wiley and Sons Inc., New York, 1978.

[21] Paul S. Heckbert and Pat Hanrahan. Beam tracing polygonal objects. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 119–127, July 1984.

[22] Renate Heinz. Binaural room simulation based on an image source model with addition of statistical methods to include the diffuse sound scattering of walls and to predict the reverberant tail. *Applied Acoustics*, 38:145–159, 1993.

[23] Reiner Horst and Hoang Tuy. *Global Optimization, Deterministic Approaches*. Springer-Verlag, New York, New York, 1996.

[24] L. Ingber. Simulated annealing: Practice versus theory. *Mathl. Comput. Modelling*, 18:29–57, 1993.

[25] George C. Izenour. *Theater Design*. Yale University Press, New Haven, 1996.

[26] George C. Izenour. *Theater Technology*. Yale University Press, New Haven, 1996.

[27] T. Kang, J. Seims, J. Marks, and S. Sheiber. Exploring lighting spaces. *Technical Report*, Mitsubishi Electric Research Laboratory MERL TR-95-18, 1996.

[28] John K. Kawai, James S. Painter, and Michael F. Cohen. Radioptimization - goal based rendering. In *Computer Graphics Proceedings, Annual Conference Series, 1993*, pages 147–154, 1993.

[29] S. Kirkpatrick, Jr. C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.

[30] A. Krokstad, S. Strom, and S. Sorsdal. Calculating the acoustical room response by the use of a ray tracing technique. *Journal of Sound Vibration*, 8(1):118–125, 1968.

[31] Heinrich Kuttruff. *Room Acoustics*. Elsevier Applied Science, London, 1991.

[32] K. H. Kuttruff. Auralization of impulse responses modeled on the basis of ray-tracing results. *Journal of the Audio Engineering Society*, 41(11):876–880, 1993.

[33] Hilmar Lehnert. Systematic errors of the ray-tracing algorithm. *Applied Acoustics*, 38:207–221, 1993.

[34] T. Lewers. A combined beam tracing and radiant exchange computer model for room acoustics. *Applied Acoustics*, 38:161–178, 1993.

[35] Albert London. The determination of reverberant sound absorption coefficients from acoustic impedance measurements. *The Journal of the Acoustic Society of America*, 22:263–269, 1950.

[36] J. Marks, B. Andalman, P. A. Beardsley, W Freeman, S. Gibson, J Hodgins, T. Kang, B. Mirtich, H Pfister, W. Ruml, K. Ryall, J. Seims, and S. Sheiber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Computer Graphics Proceedings, Annual Conference Series, 1997*, pages 389–400, August 1997.

[37] Michael Monks, Byong Mok Oh, and Julie Dorsey. Acoustic simulation and visualization using a new unified beam tracing and image source approach. In *Proceedings of the Convention of the Audio Engineering Society*, 1996.

[38] Michael Monks, Byong Mok Oh, and Julie Dorsey. Audioptimization: Goal-based acoustic design. Technical Report MIT-LCS-TM-588, Submitted for Publication, 1998.

[39] G. M. Naylor. Odeon-another hybrid room acoustical model. *Applied Acoustics*, 38:131–143, 1993.

[40] James R. Oestreich. A maligned art (science? sorcery?) gains new respect. New York Times, Nov 9, 1997. Arts & Leisure, Section 2, p. 1.

[41] R. H. J. M. Otten and L.P.P.P. van Ginneken. *The Anealing Algorithms*. Kluwer Academic Publishers, Boston, 1989.

[42] Pierre Poulin and Alain Fournier. Lights from highlights and shadows. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, pages 31–38, April 1992.

[43] Pierre Poulin and Alain Fournier. Painting surface characteristics. In *Eurographics Rendering Workshop 1995*. Eurographics, June 1995.

[44] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing (2nd ed.)*. Cambridge University Press, Cambridge, 1995.

[45] Wallace Clement Sabine. *Collected Papers on Acoustics*. Harvard University Press, Cambridge, 1922.

[46] C. Schoeneman. *Discussions*. 1995.

[47] Chris Schoeneman, Julie Dorsey, Brian Smits, James Arvo, and Donald Greenberg. Painting with light. In *Computer Graphics Proceedings, Annual Conference Series, 1993*, pages 143–146, 1993.

[48] Jiaoying Shi, Aidong Zhang, Jose Encarnacao, and Martin Gobel. A modified radiosity algorithm for integrated visual and auditory rendering. *Computers and Graphics*, 17(6):633–642, 1993.

[49] Adam Stettner. Computer graphics for acoustic simulation and visualization. Master's thesis, Program of Computer Graphics, Cornell University, January 1989.

[50] Adam Stettner and Donald P. Greenberg. Computer graphics visualization for acoustic simulation. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 195–206, July 1989.

[51] Gilbert Strang. *Introduction To Applied Mathematics*. Wellesley Cambridge Press, Wellesley, Massachusetts, 1986.

[52] Nicolas Tsingos and Jean-Dominique Gascuel. A general model for the simulaiton of room acoustics based on hierarchical radiosity. In *SIGGRAPH'97 technical sketch, Los Angeles, USA*, August 1997.

[53] Edward R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, Connecticut, 1994.

[54] Edward R. Tufte. *Visual Explanations*. Graphics Press, Cheshire, Connecticut, 1997.

[55] A. Krokstad U. R. Kristiansen and T. Follestad. Extending the image method to higher-order reflections. *Applied Acoustics*, 38:195–206, 1993.

[56] P. J. M. van Laarhoven and E. H. L. Aarts. *Simulated Annealing: Theory and Applications*. D. Reidel Publishing Company, Dordrecht, Holland, 1987.

[57] Dirk van Maercke and Jacques Martin. The prediction of echograms and impulse responses within the epidaure software. *Applied Acoustics*, 38:93–114, 1993.

[58] Michael Vorlander. Simulation of the transient and steady-state sound propogation in rooms using a new combined ray-tracing/image-source algorithm. *Journal of the Acoustic Society of America*, 86:172–178, 1989.

[59] John P. Walsh. The design of godot: A system for computer-aided room acoustic modeling and simulation. In *Proceedings of the 10th International Conference on Acoustics*, July 1980.

[60] John Patterson Walsh. The simulation of directional sound sources in rooms by means of a digital computer. Master's thesis, The University of Western Ontario, August 1979.

[61] Kevin Weiler. Polygon comparison using a graph representation. In *Computer Graphics (SIGGRAPH '80 Proceedings)*, volume 14, pages 10–18, July 1980.

[62] Kevin Weiler and Peter Atherton. Hidden surface removal using polygon area sorting. In *Computer Graphics (SIGGRAPH '77 Proceedings)*, volume 11(2), pages 214–222, July 1977.

[63] Shuo xian Wu. Applying fuzzy set theory to the evaluation of concert halls. *Journal of the Acoustic Society of America*, 89(2):772–776, 1991.