TKK Dissertations in Media Technology
Espoo 2010 TKK-ME-D-4

EFFICIENT PHYSICS-BASED ROOM-ACOUSTICS MODELING AND AURALIZATION

Samuel Siltanen

Dissertation for the degree of Doctor of Science in Technology to be presented with due permission of the Faculty of Information and Natural Sciences for public examination and debate in Auditorium T2 at Aalto University School of Science and Technology (Espoo, Finland) on the 15th of January, 2010, at 12 noon.

Aalto University School of Science and Technology
Faculty of Information and Natural Sciences
Department of Media Technology

Aalto-yliopiston teknillinen korkeakoulu
Informaatio- ja luonnontieteiden tiedekunta
Mediatekniikan laitos

Distribution:
Aalto University School of Science and Technology
Department of Media Technology
P.O.Box 5400
FIN-02015 TKK
Finland
Tel. +358-9-470 22870
Fax. +358-9-470 25014
http://media.tkk.fi/

Available in PDF format at http://lib.tkk.fi/Diss/2010/isbn9789522482645/

ISBN 978-952-248-264-8 (print)
ISBN 978-952-248-265-5 (online)
ISSN 1797-7096 (print)
ISSN 1797-710X (online)

Yliopistopaino
Helsinki 2010

ABSTRACT

Author Samuel Siltanen
Title Efficient Physics-Based Room-Acoustics Modeling and Auralization

The goal of this research is to develop efficient algorithms for physics-based room acoustics modeling and real-time auralization. Given the room geometry and wall materials, in addition to listener and sound source positions and other properties, the auralization system aims at reproducing the sound as would be heard by the listener in a corresponding physical setup. A secondary goal is to predict the room acoustics parameters reliably.

The thesis presents a new algorithm for room acoustics modeling. The acoustic radiance transfer method is an element-based algorithm which models the energy transfer in the room like the acoustic radiosity technique, but is capable of modeling arbitrary local reflections defined as bidirectional reflectance distribution functions.

Implementing real-time auralization requires efficient room acoustics modeling. This thesis presents three approaches for improving the speed of the modeling process. First, the room geometry can be reduced. For this purpose an algorithm, based on volumetric decomposition and reconstructions of the surface, is described. The algorithm is capable of simplifying the topology of the model and it is shown that the acoustical properties of the room are sufficiently well preserved with even 80 % reduction rates in typical room models. Second, some of the data required for room acoustics modeling can be precomputed. It is shown that in the beam tracing algorithm a visibility structure called "beam tree" can be precomputed efficiently, allowing even moving sound sources in simple cases. In the acoustic radiance transfer method, effects of the room geometry can be precomputed. Third, the run-time computation can be optimized. The thesis describes two optimization techniques for the beam tracing algorithm which are shown to speed up the process by two orders of magnitude. On the other hand, performing the precomputation for the acoustic radiance transfer method in the frequency domain allows a very efficient implementation of the final phase of the modeling on the graphics processing unit. An interactive auralization system, based on this technique is presented.

UDC 534.84, 004.021, 004.92, 004.94
Keywords room acoustics modeling, auralization, virtual reality

TIIVISTELMÄ

Tekijä	Samuel Siltanen
Työn nimi	Tehokas fysikaalinen huoneakustiikan mallinnus ja auralisaatio

Tämän tutkimuksen tavoite on kehittää tehokkaita algoritmeja fysikaaliseen huoneakustiikan mallinnukseen ja reaaliaikaiseen auralisaatioon. Kun huoneen geometria ja seinien materiaalit sekä kuuntelijan ja äänilähteen paikat ja muut ominaisuudet on annettu, auralisaatiojärjestelmä pyrkii tuottamaan äänen, jonka kuuntelija kuulisi vastaavassa fyysisessä asetelmassa. Toissijainen tavoite on ennustaa luotettavasti huoneakustisia tunnuslukuja.

Väitöskirjassa esitellään uusi algoritmi huoneakustiikan mallinnukseen. Akustinen radianssinsiirtomenetelmä on elementtipohjainen algoritmi, joka mallintaa energiansiirtoa huoneessa akustisen radiositeettialgoritmin tapaan, mutta kykenee mallintamaan mielivaltaisia paikallisia heijastuksia, jotka on määritelty kaksisuuntaisina heijastusjakaumafunktiona.

Reaaliaikaisen auralisaation toteutus vaatii tehokasta huoneakustiikan mallinnusta. Tämä väitöskirja esittää kolme lähestymistapaa mallinnusprosessin nopeuden parantamiseksi. Ensiksi, huoneen geometriaa voidaan yksinkertaistaa. Tätä tarkoitusta varten on kuvattu algoritmi, joka perustuu tilavuushajotelmaan ja pinnan uudelleenrakennukseen. Algoritmi kykenee yksinkertaistamaan mallin topologiaa ja on näytetty, että akustiset ominaisuudet säilyvät jopa 80 % pelkistysasteilla tyypillisten mallien tapauksessa. Toiseksi, osa huoneakustiikan mallintamiseen tarvittavasta datasta voidaan esilaskea. On osoitettu, että keilojenseuranta-algoritmissa näkyvyysrakenne nimeltään “keilapuu” voidaan laskea tehokaasti niin, että jopa liikkuva äänilähde on mahdollinen yksinkertaisissa tapuksissa. Akustisessa radianssinsiirtomenetelmässä huoneen geometrian vaikutus voidaan esilaskea. Kolmanneksi, ajonaikaista laskentaa voidaan optimoida. Väitöskirjassa kuvataan kaksi optimointitekniikkaa keilojenseuranta-algoritmiin. Näiden on osoitettu nopeuttavan prosessia kaksi kertaluokkaa. Toisaalta, suorittamalla esilaskenta taajuustasossa akustisessa radianssinsiirtomenetelmässä, tulee mahdolliseksi mallinnuksen viimeisen vaiheen erittäin tehokas toteutus grafiikkaprosessorilla. Tähän tekniikkaan perustuen esitellään vuorovaikutteinen auralisaatiojärjestelmä.

UDK	534.84, 004.021, 004.92, 004.94
Avainsanat	huoneakustiikan mallinnus, auralisaatio, virtuaalitodellisuus

PREFACE

The research work presented in this thesis was carried out at the Telecommunications Software and Multimedia Laboratory, Department of Media Technology, Helsinki University of Technology, Espoo, during 2006-2009.

I want to express my gratitude to Prof. Lauri Savioja, the supervisor of this thesis, and to Dr. Tapio Lokki, the instructor of this thesis, for support and encouragement. In addition, I would like to thank for our almost weekly discussions and collaboration with the research papers which made the progress of the thesis work much easier.

I am also indebted to M.Sc. Sami Kiminki who derived the room acoustic rendering equation in his master's thesis that provided the foundation for some of the results presented in this thesis. I would like to thank Dr. Samuli Laine for the development of the beam tracing algorithm, which is also analyzed in this thesis. In addition, I thank M.Sc. Claus Lynge Christensen for the ODEON tests runs which were required when validating the geometry reduction algorithm.

I am also grateful for the effort of the pre-examiners, Dr. Seppo Uosukainen and Dr. Jason Summers. Their expertice was required to assure the quality of the presented research work.

I thank my other co-workers (Sampo, Raine, Jukka, Sakari) who, during the years of writing this thesis, shared my office room and provided a supportative and relaxed atmosphere.

I would like to thank my family and friends who tolerated me while I was stressed because I had to finish a research paper in time or because the results were not always what I had hoped for. It was essential to have someone to remind me that there are more important things in life than the thesis work.

Finally, I am grateful to the organizations that have supported this work financially: the Graduate School of Computer Science and Engineering at the Helsinki University of Technology, Tekniikan edistämissäätiö, and the Nokia Foundation. The research leading to these results has also received funding from the Academy of Finland, project no. [119092] and the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [203636].

Otaniemi, Espoo, 7^{th} December 2009

Samuel Siltanen

LIST OF PUBLICATIONS

This thesis summarizes the following articles and publications, referred to as [P1]–[P4]:

[P1] S. Siltanen, T. Lokki, L. Savioja, and C. L. Christensen. Geometry Reduction in Room Acoustics Modeling, *Acta Acustica united with Acustica*, 94(3):410–418, 2008.

[P2] S. Laine, S. Siltanen, T. Lokki, and L. Savioja. Accelerated beam tracing algorithm, *Applied Acoustics*, 70(1):172–181, 2009.

[P3] S. Siltanen, T. Lokki, S. Kiminki, and L. Savioja. The room acoustic rendering equation, *Journal of the Acoustical Society of America*, 122(3):1624–1635, 2007.

[P4] S. Siltanen, T. Lokki, and L. Savioja. Frequency Domain Radiance Transfer for Real-time Auralization, *Acta Acustica united with Acustica*, 95(1):106–117, 2009.

LIST OF ABBREVIATIONS

2-D	Two-dimensional
3-D	Three-dimensional
BEM	Boundary element method
BRDF	Bidirectional reflectance distribution function
BSP	Binary space partitioning
BTM	Biot-Tolstoy-Medwin diffraction model
C_{80}	Clarity
CAD	Computer-aided design
CPU	Central processing unit
D_{50}	Definition
DFT	Discrete Fourier transform
DWG	Digital waveguide mesh
EDT	Early decay time
EF	Eigentransfer function
FDTD	Finite-difference time-domain
FEM	Finite element method
FFT	Fast Fourier transform
FIR	Finite impulse response
G	Strength
GPU	Graphics processing unit
GTD	Geometrical theory of diffraction
HRTF	Head-related transfer function
HRIR	Head-related impulse response
IIR	Infinite impulse response
ITD	Interaural time delay
LF	Lateral fraction
PC	Personal computer
PCA	Principal component analysis
SH	Spherical harmonics
SPL	Sound pressure level
T_{60}	Reverberation time
TLM	Transmission line matrix
TS	Center time
UTD	Universal theory of diffraction

1 INTRODUCTION

The thesis concentrates on showing how the room acoustics modeling process can be performed efficiently and physically accurately in the context of real-time auralization. This means that when given a room model with a geometry and material description, sound source properties, listener properties and a sound signal, the goal is to reproduce the sound heard by the listener as if he were in the corresponding physical room containing that sound source emitting that sound signal. In addition, the listener should be able to move in the room and be able to hear the changes in the sound in real time, i.e. at a rate that still sounds natural to a human listener.

Since computational resources are always limited, it is impossible to reach the goal described above perfectly. Some simplifying assumptions are required to reduce the task to a level on which the modern computer hardware is efficient enough for real-time computation. On the other hand, since the other goal is to use a physics-based modeling approach, the simplifying assumptions should be kept to a minimum.

The goal of physics-based modeling means that the propagation of sound should be modeled as well as possible. Some acoustical effects that are often ignored by the well-known geometrical room acoustics modeling techniques [162] are:

- diffraction
- non-ideal reflections
- late reverberation.

This thesis shows that these effects can be incorporated in the presented room acoustics modeling techniques without sacrificing the computational performance.

1.1 Scope of This Thesis

Three approaches are taken to improve the efficiency of the acoustics modeling and auralization process:

- simplification of the room geometry
- precomputation of the sound propagation in a room model
- optimization of the run-time computation and auralization.

Figure 1.1 shows how the different approaches contribute to the overall goal. The geometry reduction process simplifies the complex room geometry to a level of detail that can be handled by the acoustics modeling algorithms in a reasonable time. Then, one of the several room acoustics modeling techniques can be applied to produce an intermediate acoustic transfer model for all the data that does not change in run-time. This acoustic transfer model can speed up the real-time processing, since only the computation for the

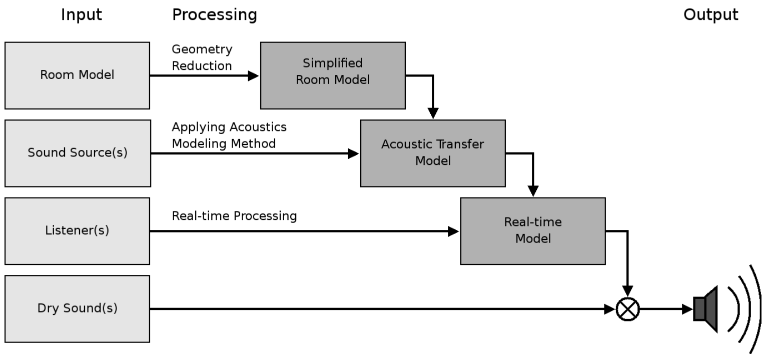

Figure 1.1: Relationships between the different research areas can be seen in the diagram. The input data consists of a room model, one or multiple sound sources, one or multiple listeners and one or multiple dry sound streams. The result of the processing should be spatial sound, which corresponds to the input data.

changing data (i.e. listener) has to be performed. Even the run-time process can be optimized and possibly run on the graphics processing unit (GPU) while the central processing unit (CPU) is free to, e.g., process the signals for the final auralized output.

Modeling sound sources is beyond the scope of this thesis. Simple spherically radiating point-like sources are used when a source model is required. However, generalizations to more complex sources are possible in some cases, which is indicated in the text when appropriate.

Listener modeling and sound reproduction are also beyond the scope of this work. There are several alternative loudspeaker reproduction systems [120, 135] for output as well as head-related transfer function (HRTF) models for headphone listening [53, 88, 121]. When implementing auralization, one HRTF model has been chosen, but any other output system could be applied with small changes.

The room acoustics modeling approaches presented in this thesis are limited to computing the effects of locally-reacting room surface geometry. Modeling structural vibrations is not within the scope of this thesis. The medium is assumed to be non-dispersive (linear). Figure 1.2 illustrates the limits of the research area of this thesis in the context of an auralization system.

1.2 Simplifying the Room Geometry

Since the efficiency of most of the room acoustics modeling methods depends on the complexity of the room geometry, reducing it should produce performance gains in most cases. The geometric models of the rooms used for visualization purposes are often detailed, but the acoustics of the room is not significantly affected by small details. Thus, creating an "acoustic version" of the room makes sense. Preparing such a version manually can be very laborious. An automatic geometry reduction tool could make the task

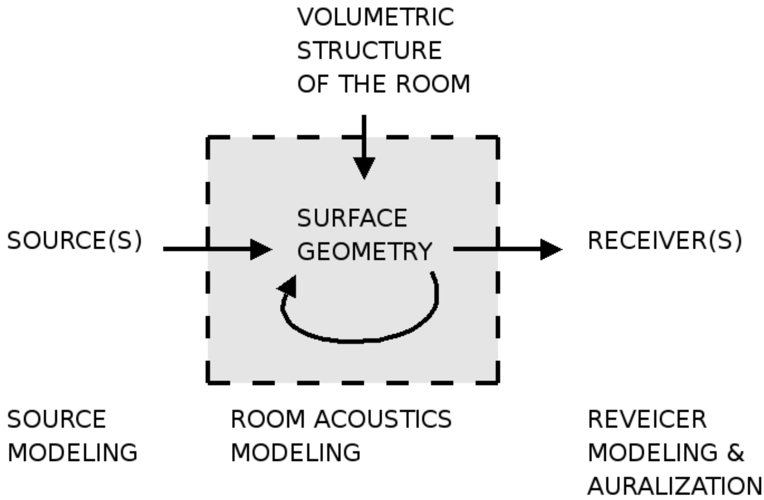

Figure 1.2: The thesis concentrates on the effects of the room surface geometry with a linear medium. More general acoustic modeling would require a volumetric model of the room. Modeling sound sources and receivers is necessary for auralization, since the sound must be emitted into the room and then detected by a listener, but that is not part of this research.

easier.

The simplification process should still preserve the acoustical properties of the room as well as possible. On the other hand, it should be able to produce models with the desired level of detail. Publication [P1] presents one technique for this purpose.

1.3 Precomputing an Acoustical Model

If the geometry of the room is static, i.e. it is not expected to change in run-time, some parts of the computation can be done beforehand and the results of the computation can be stored in an appropriate data structure. If the sound source is also static, even more data can be precomputed. Then, most of the computation has already been performed before the position and orientation of the listener is fixed. Thus, updating the results when changing the listener position can be very efficient.

Visibility is one property that can be precomputed. This is the basic idea behind the beam tracing algorithms [47]. Publication [P2] concentrates on improving the beam tracing technique.

Another approach is to discretize the geometry, i.e. split the model into elements, and precompute the acoustic relationships between the elements. There are already several modeling techniques that use elements, but publication [P3] presents a novel technique: acoustic radiance transfer. The properties of the algorithm are discussed later.

1.4 Optimizing the Run-time Computation

The run-time computation for real-time auralization must be efficient. In addition to the acoustics modeling, listener modeling and signal processing must also be performed. The acoustic model must be updated whenever the listener moves and the lag between the listener movement and the change heard in the output signal should be tolerable. This requires a highly optimized implementation of all parts of the run-time computation.

When using the beam tracing approach, it is likely that the reflection paths for two listener positions close to each other are very similar. This coherence can be utilized for a moving listener, since his position is likely to change only a little in one step. Publication [P2] shows how the beam tracing algorithm can thus be optimized.

On the other hand, modern graphics cards are very efficient in parallel processing and employing them in element-based acoustics modeling can lead to an efficient implementation. Publication [P4] describes how the run-time phase of the acoustic modeling and the listener modeling can be run on the GPU, while the CPU does the signal processing.

While Publication [P4] presents a complete auralization system for the room acoustics modeling method presented in publication [P3], publication [P2] describes only the room acoustics modeling part of the system without the signal processing and listener modeling. However, this room acoustics modeling method can be used for real-time auralization as well by using the techniques presented in publication [P4] or utilizing other auralization frameworks [83, 129]. Validation of the techniques is done by comparing the extracted room acoustics parameters to those measured in a real room. The room acoustics parameter prediction is a secondary goal of the presented methods.

1.5 Organization of the Thesis

This thesis is organized as follows. In section 2 the most relevant previous research is reviewed. The new acoustical models which serve as a foundation for the presented modeling algorithm in publication [P3] is described in section 3. The geometry simplification technique for producing low-detail geometry is presented in section 4. Section 5 covers the precomputation phase algorithms, including beam tracing, and acoustic radiance transfer. Section 6 concentrates on the run-time computation optimizations and the auralization system. Section 7 summarizes the results of the thesis and suggests future research topics.

2 RELATED RESEARCH

The research related to the topic of the thesis is reviewed in this section. First, the many approaches to geometry reduction are surveyed. Then, the various room acoustics modeling methods are discussed. Finally, some research related to real-time auralization is presented.

2.1 Geometry Reduction

The following discussion assumes the original room model to be a polygonal mesh. This representation is general enough for most purposes. Each polygon represents a wall or a part of a wall and some material properties can be linked to it. Most of the previous research on geometry simplification has been in the area of computer graphics, but the algorithms can usually be viewed in a more general context. Figure 2.1 shows a taxonomy of geometry simplification techniques. There are two basic approaches: decimation of the existing surface and reconstructing a new surface. [54, 117]

Decimation Algorithms

The decimation algorithms remove elements of the surface by using some heuristics. The removed elements can be vertices, edges, or triangles (or parts of triangulated polygons). Vertex clustering techniques can also be classified into this category.

In the vertex removal approach [155] the vertices are ordered according to an error metric and the vertex whose removal causes the least error is removed first. Also the polygons attached to the vertex are removed and replaced by a fewer number of new polygons. The process of choosing the vertex that causes the least error is repeated as many times as necessary to achieve the desired level of reduction. When the error metric is recomputed from the reduced model after each vertex removal, the errors tend to accumulate and thus deteriorate the model. There have been attempts to improve the situation by storing error values at the vertices neighboring the removed one [154, 7] or storing the removed vertices themselves [157]. Some variants of the algorithm use the distance from the original surface as an error metric [23]. In other algorithms it is explicitly required that the original surface is within a certain distance from the reduced surface [89]. Since the requirement is set for the original surface, it is still possible the some parts of the reduced surface to be further of the original surface. A similar requirement is that reduced surface is within a certain distance from the original surface [27, 182], which allows some parts of the original surface to be further from the reduced surface, but not the other way round. Material data or curvature can also be taken into account by the error metric [7, 176]. There are also various approaches to the triangulation of the gap caused by removed polygons. The task can be simplified into a 2-dimensional problem [89, 137] or a greedy algorithm can be used [176]. Some algorithms spend more effort to find a nearly optimal triangulation [157, 23]. A com-

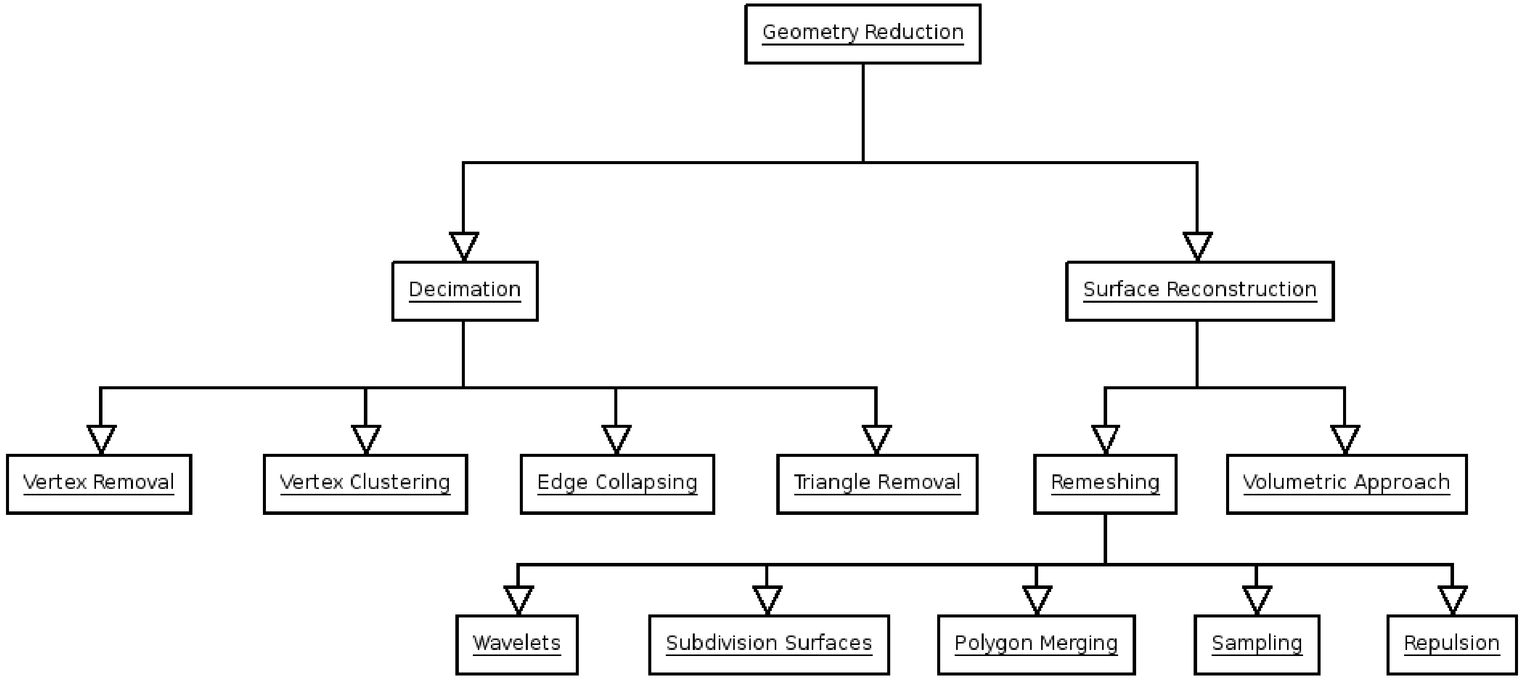

Figure 2.1: Taxonomy of geometry simplification techniques.

mon limitation of the vertex removal algorithms is that they cannot change the topology of the object and thus the reduction rate is sometimes limited.

Vertex clustering algorithms cluster together nearby vertices and replace them with one new vertex. The simplest algorithm uses a regular grid on the geometry and clusters vertices inside each cell [141]. The vertices can be graded according to their importance and the grade can be used to choose the placement of the new vertex, which can either be a weighted average of the clustered vertices or the vertex with the highest grade. A more flexible approach is to use floating boxes around the highest graded vertices as clustering volumes [116]. A hierarchical approach is also possible where the clusters are clustered into larger clusters [118]. At the finest level clusters contain only one vertex and at the coarsest level the whole object. Also in this case, there are several heuristics to chose the new vertices representing the clusters [150, 151]. Vertex clustering techniques can simplify the topology, but the quality of the reduced models is usually not high. Cracks and gaps can be created in the models which can be problematic in room acoustics modeling.

Edge collapsing is perhaps the simplest approach to geometry simplification since no retriangulation is needed. An edge is collapsed into a vertex by pulling the end points of the edge together. Thus, two triangles are effectively removed. The simplest algorithm assigns each edge the cost of collapsing it and orders the edges according to that cost. The edges are then collapsed in order until the desired error tolerance is reached. The cost metric can be the squared distance of the new vertex to the planes of the neighboring triangles and the vertex is placed so that it minimizes that cost [140]. The squared distances to the planes of the adjacent triangles can be written by using a symmetrical matrix. The sum of the squared distances can simply be found by summing the matrices. Using this matrix formulation is called the quadric error metric [57, 55] and, with some modifications, it can also be used with sharp edges, boundaries, and edges at the boundaries of different materials [56, 44, 45, 72]. There are edge collapsing algorithms which try to preserve the volume of the objects either explicitly [62, 63] or as a part of the chosen error metric [106, 158]. Also, several other error metrics have been suggested [75, 70, 71, 43, 134, 142, 108, 92, 181]. There are different approaches to placing the collapsed vertex. Most of the algorithms use the position that minimizes the error metric. Some algorithms use one of the end points of the edge [92]. The placement that preserves the volume [62, 63, 106, 108] has also be shown to be effective compared to other approaches [107]. In addition to the quadric approach, there have been other algorithms which try to handle the material properties through an edge collapse [25, 26, 136, 24]. Since, in its basic form, the edge collapsing algorithms cannot change the topology of the model, which can limit the reduction, using controlled topology simplification with it has been researched [41, 42]. Yet, there is a version of the edge collapsing algorithm that tries to recognize which features of the model are important and control the reduction accordingly [1]. Edge collapsing works well also when the original models are very large [104, 105]. Although the edge collapsing algorithms work very well with models with smooth surfaces, they are not as well suited for reducing models with rectangular features, such as rooms.

There are also triangle removal algorithms [59, 64], but since removing a triangle can be presented as a series of edge collapses [167] and requires an additional examination of the surface prior to the removal, the technique has little independent value.

Surface Reconstruction Algorithms

The surface reconstruction methods can either work on the surface or on the volume. One of the surface remeshing methods is wavelet-based reduction [113, 114], where the detailed surface is decomposed into a coarser surface and a wavelet base and coefficients which represent the difference to the detailed surface. The detailed surface can be reconstructed by subdivision of the coarse surface and perturbing the vertices according to the wavelet representation. The level of details can be controlled by filtering the coefficients. The original algorithm required that the surface can be formed by subdivision, but this constraint can be removed by re-meshing the object [40]. The wavelet decomposition can also be done for material properties [20] and a hierarchical level-of-detail structure can be created by using a quadtree [61].

If there are point samples available from the surface of an object, the surface can be reconstructed by fitting bicubic surface patches to the data while assuring continuity at the patch boudaries [152]. The level of details can be controlled by beginning with a coarse approximation and subdividing the patches until the fitting error is small enough. However, it is possible to use the same framework for simplifying objects whose surface is known. This can be done by sampling the surface and reconstructing it with a desired level of detail with the bicubic patches as if only the sampled data were available [80]. Instead of the bicubic patch approach, direct mesh optimization techniques can be used [74, 73]. They can reconstruct the surface from an unorganized set of points. A somewhat similar approach is to place the desired number of vertices on the surface to be simplified, add repulsion forces between the vertices, and then let the vertices move on the surface until the repulsion energy is minimized [173]. The simplified triangle mesh can be obtained by creating triangles between the vertices by utilizing the connectivity information of the original mesh.

Yet another important group of surface reconstruction algorithms consists of the polygon merging techniques. Especially when there is a large number of nearly coplanar polygons, this approach can be efficient. A simple approach is to choose a polygon randomly and merge adjacent polygons to it until the error becomes too large [80]. A more sophisticated algorithm groups the polygons into coplanar sets, for which there are efficient algorithms [143], and then removes the edges shared by more than one polygon in the sets [68]. If the polygons in a set are allowed to deviate slightly from the coplanarity criterion and if the edges of the sets are also simplified, even more impressive results can be achieved [84]. There are also algorithms that guarantee a certain distance between the original and the simplified surfaces [175].

In volumetric topology simplification algorithms the geometry is inserted into a volumetric structure. In the most simple case, the volumetric structure is a regular 3-dimensional grid [66, 130]. The grid cells do not contain information about the original surface, but some other information, such as "density" values [66, 67]. The density values can be obtained by applying a filter kernel over the surface so that the density depends on the distance from the surface. There are standard algorithms for creating isosurfaces for 3-dimensional density data, such as the marching cubes [112], which can be used for reconstructing a surface which lacks details smaller than the order of the cell size. The cell size thus determines the level of simplification. It should be noted that coarser simplifications require less computation time and space, while detailed reconstructions can cost too much for pratical purposes [5]. The resulting surface might still need an additional simplification by using another algorithm, but that should be easier since the volumetric approach has removed complex topological structures. The reconstructed surface can be simplified further by using, e.g., standard edge collapsing algorithms [130, 6]. There are some variations of the basic algorithm. The volumetric structure can be hierarchical [67, 5, 6]. Some algorithms do not use density values but record whether a volumetric cell is inside or outside the simplified object [5, 6, 130]. Then, it is possible to reconstruct a new surface from a hierarchical structure directly and not as an isosurface [5, 6]. In general, the output of the volumetric techniques is not depedent on the detail level or the topology of the original model, and thus they can be used for radical reductions.

2.2 Room Acoustics Modeling

The goal of room acoustics modeling is to compute the sound field in a room when the sound sources and the room geometry and material properties are known. In theory, the sound field can be described by wave equations and appropriate boundary conditions for the walls, but finding analytical solutions for all but the most trivial cases is very difficult. Thus, there are two basic approaches for constructing the field. One is to model propagating wavefronts with rays. Thus, the sound is assumed to travel along a straight path and reflect when encountering obstables. The other approach is to solve an appropriate wave equation numerically, which requires that the space is discretized into elements. The first approach is referred to as geometrical room acoustics and the latter one as wave-based room acoustics. Figure 2.2 shows a taxonomy of room acoustics modeling techniques. There exist perceptual models also [171], but since the goal of the thesis is physics-based modeling, they are not covered here.

Geometrical Room Acoustics

The assumption is that in a homogenous medium the sound travels along straight paths or rays. When the modeled wavelength is small compared to the obstacles in the room, this assumption is fairly safe. But when longer wavelengths are modeled, a diffraction model is required to compensate for the errors caused by the ray assumption. The different geometrical room acoustics modeling methods are disscussed first, then some diffraction models, followed by general reflection modeling, and finally the methods used in room acoustics prediction software.

There is a simple technique for modeling ideal specular reflections. The

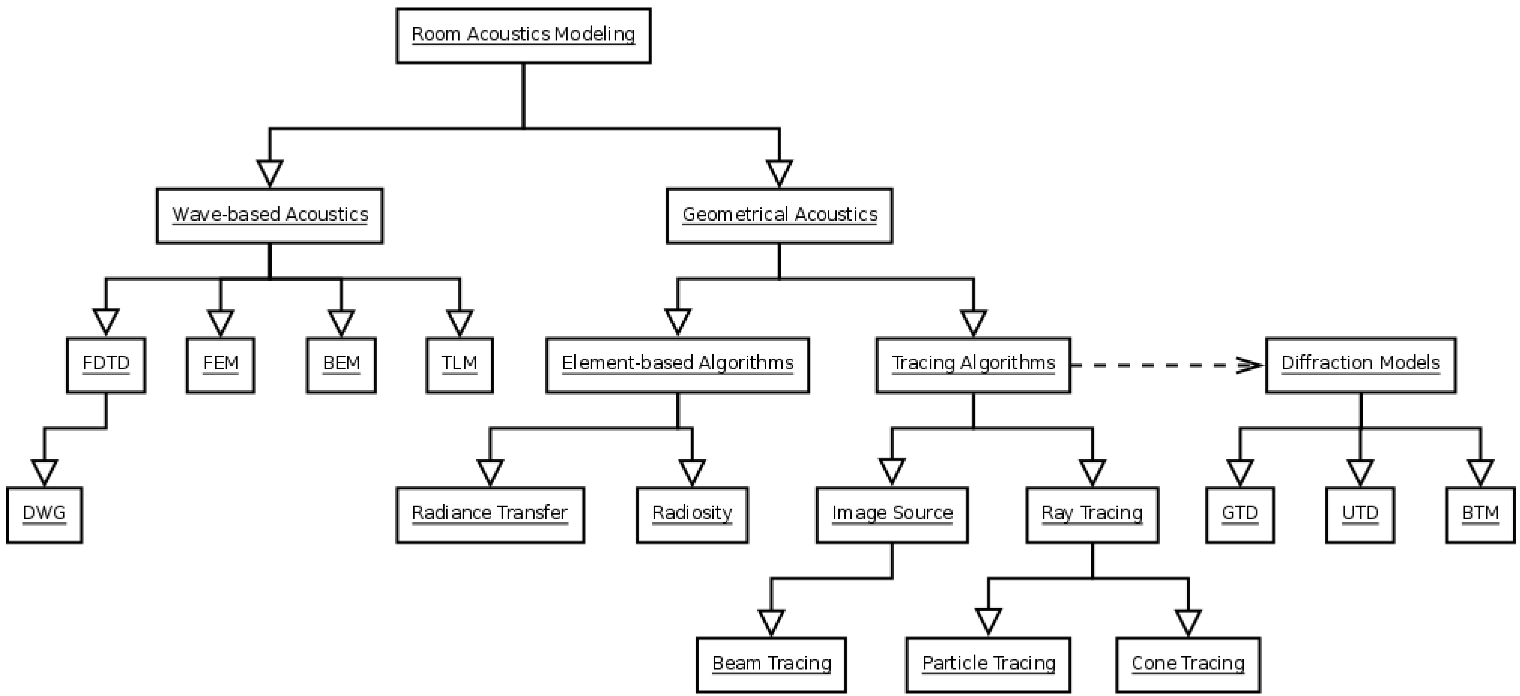

Figure 2.2: Taxonomy of room acoustics modeling techniques.

wave front emitted by a source and reflected at a plane is the same as the wave front emitted by the source mirrored at the plane. Thus, reflections can be modeled by creating a mirrored source, i.e. an image source, for each plane in the room geometry. Then the receiver gets contributions both from the real source and from the image sources which correspond to the direct field and the once-reflected fields, respectively. In addition, new image sources can be created for the first image sources, thus modeling the twice reflected field also. Higher order reflections can be modeled by creating new image sources of the image sources up to the desired order. It should be noted that the number of image sources grows exponentially in relation to the reflection order for an arbitrary geometry. However, in a general case, all the image sources are not valid since some computed paths from the image sources to the receiver do not actually hit the reflecting walls or the paths are occluded by other walls. Thus, explicit validity and visibility tests are required. In practice, the image source method is useful for computing only the first few reflection orders. An impulse response can be constructed from the reflection paths by computing their lengths and applying appropriate distance attenuation and delay. Also, when the paths hit the walls, their strength must by multiplied by the wall reflection coefficient to account for absorption. The quantity used in the computation is typically pressure and thus the phase information can also be computed from the path lengths. [3, 13]

It is possible to optimize the image source method by never creating the invalid image sources. This can be done by beam tracing [36, 47, 48, 49, 51, 50, 125, 124]. Given a polygonal room model, a cone or beam is created for each polygon so that its apex is the sound source and its base is the polygon. These beams form the first level of a beam tree. The next level is created by mirroring each beam at the plane of the base polygon and splitting the mirrored beam for each polygon it intersects. The apex of each new beam is the image source and the base is the intersecting polygon. The beams are further mirrored at the planes of the intersecting polygons until the desired depth is reached. All the beams are stored in the beam tree in a hierarchy where the original beams are parents of the reflected and split beams. The beam tree can be created as a precomputation step. Then it is relatively quick to find the reflection paths. This can be done by examining whether the receiver is inside a beam. Then the path can be found by traversing the beam tree upwards from the beam containing the listener. There are variants of the beam tree algorithm. If the intersecting polygons are accurately cut against the beams and the beams are split if they encounter occluding polygons, no additional path validity checks are required. Otherwise the same kind of test is required as in the image source method. Although the computational requirements are remarkably reduced compared to the image source method, the problem of exponential growth still exists. Fortunately, the beams tend to get narrower for higher reflection orders, and the number of beams remains low enough for interactive computation of the early part of the impulse response.

Another approach to geometrical room acoustics modeling is ray tracing [97, 133]. Instead of explicitly computing the specular reflection paths, a large number of rays is sent from the sound source and they are traced to find where they hit. If a ray hits a wall, it is reflected and traced further. The receiver is modeled as an object having a volume, typically a sphere, and an intersection test is performed between the rays and the volume. If there is an intersection, the ray contributes to the response at the receiver. The rays are thought to carry a portion of the energy of the sound sources. The energy is attenuated by the distance and some of it is absorbed at the reflections. The tracing can be stopped when the energy of a ray is neglible. The lengths of the reflection paths determine the delays. It should be noted that the quantity accumulated at the response is energy, not pressure. The energy response can be easily used for computing room acoustics parameters and thus ray tracing is often used in the room acoustics prediction software. It is not as common a technique in auralization systems. Since a limited number of rays is sent from the source, some reflection paths might be missed or their contributions might be miscalculated [101]. To improve the accuracy the rays are sometimes split at reflections or a hybrid system is used where the early reflections are computed by an image source method. Still, ray tracing can be more flexible than the image source or beam tracing methods, since the reflections do not have to be ideally specular. An arbitrary reflection pattern could be used and the direction of the reflected ray could be randomly chosen by using the pattern as a probability distribution of the reflection direction. Ray tracing can also be used for validating the reflection paths in the image source method.

One problem with the ray tracing approach is that the receiver must have a volume, and the shape and size of the volume thus change the results. The ray tracing technique can be modified by sending cones or beams from the source instead of rays [35, 179]. Then a point-like receiver can be used. This technique resembles the beam tracing technique, but accurate clipping of the cones or beams is not performed and the initial beams may overlap. Typically a radius which increases with distance is used with a ray to determine the cone. Although some sources also refer to this technique also as beam tracing, a more appropriate name would be cone tracing to distinguish between the two slightly different approaches.

Ideal specular reflection is the most appropriate reflection model for smooth, hard surfaces. But often surfaces are rough or have a detailed structure. It is infeasable to model the detailed geometry with millions of polygons. Instead, larger flat surfaces with diffuse reflection properties can be used as a decent approximation. Although ray tracing could be used to model the diffusion [98], there is a room acoustics modeling method which is based specifically on diffuse reflections: acoustic radiosity. The radiosity method has been used extensively and successfully in computer graphics [29, 28, 60] and its acoustics counterpart has also been shown to perform well in certain cases [172, 132, 131, 69]. In the method, the geometry is split into surface elements or patches. It is then possible to compute the energy transfer between each pair of patches, when the diffuse or Lambertian reflection is assumed. The transfer coefficients, i.e. form factors [153], can be gathered in a large matrix where the elements describe the portion of the energy leaving one patch and arriving at another. In the acoustic radiosity, the time delay must be also taken into account. The energy response can be computed by sending the energy from the sound source to all the patches which are visible to it and storing the received impulse at the patches. The impulses, scaled by the form factors and appropriately delayed, are then sent from each patch to all the other patches visible to that patch, and the arriving impulses are again stored at the receiving patches. Then the patches will have energetic responses stored at them. The responses are sent further to other patches, scaled and delayed as before, and the process is repeated until the transferred energy in neglible. Eventually each patch will have a time-dependent energy response corresponding to the effects of the room geometry to the energy sent by the source. The last step is to gather the energies from the patches to the receiver to construct the final energy responses. The computation time per energy propagation step is proportional to the square of the number of elements, but, on the other hand, the number of elements is constant, which leads to a constant computation time for each reflection order. Thus, exponential growth is not a problem as it is with the approaches using the specular reflection. In addition, by using the progressive radiosity approach, which prioritize the transfer according to the energy on a patch [178], the computation time could be further reduced. Still, the diffuse reflection model limits the use of this technique.

Regardless of the geometric room acoustics modeling technique, the constructed field has one defect. There is a sudden jump in pressure or energy levels when the receiver moves from a region where the sound source is directly visible to a region where it is occluded. According to the laws of physics such a jump cannot occur, but the transition should be smooth. To correct this defect, edge diffraction models have been developed to be attached to the geometrical acoustics solutions. In auralization, the effect of diffraction can be significant [166, 17]. Biot and Tolstoy derived an accurate analytical solution for an infinite length wedge of perfectly specularly reflecting material [11]. Later Medwin et al. interpreted the presented solution according to the Huygen's principle as a total contribution from infinitely many point sources along the edge and used that view as a basis for deriving the corresponding solution for a finite length wedge [122]. Svensson et al. wrote the solution in the form of a line integral which can readily be used with the geometrical room acoustics modeling methods [161]. The numerical stability of this method has further been improved [160]. The usual approach to compute the diffraction with this model is to place several point sources along the modeled edge and using the diffraction model as an emittance pattern for those secondary sources [18]. The computational load can be significant if a large number of sources has to be used. It is also interesting to note that the Biot-Tolstoy-Medwin diffraction model could be used as a geometric acoustics modeling method as such [19]. This diffraction model has also been compared to measured data in a real room, and it shows quite a good agreement [115, 110].

Other, more approximate solutions to the edge diffraction problem exist. One of them is the Kirchhoff diffraction approximation which, however, has been shown to fail in certain cases [78, 87]. Other simplified formulations include the geometric theory of diffraction [85] and the uniform theory of diffraction [93] which are high-frequency asymptotic solutions. For auralization purposes the more approximate solutions might be sufficient [170, 168], since fast computation is more important than the accurate sound levels, but for physics-based room acoustics prediction the more accurate solution would be more appropriate.

Specular and diffuse reflections are not the only kind of reflection models that can be used in acoustic modeling. In reality, the acoustic reflections can have a more complex directional dependence, which is often simplified to absorption and diffusion or scattering coefficients [31, 177]. In computer graphics there has been research on applying more general reflection models, especially in radiosity-like systems [16, 77, 156, 99]. The concept of bidirectional radiance distribution functions (BRDFs) [128] is used to describe a general reflection. This work has inspired some of the research in this thesis.

Typically, hybrid models are used in room acoustics prediction software. This is because both specular and diffuse reflections must be modeled. One technique might be effective for one type of reflection while an other algorithm is more appropriate for another type of reflection. Some systems use cone tracing and radiosity-like algorithms [103] and others use image source and ray tracing models [127]. On the other hand, some techniques, such as cone tracing, can model both specular and diffuse reflections efficiently and it has also been used in room acoustics prediction software [35]. The different room acoustics modeling programs have been compared to measured data and it has been noted that while their modeling is quite accurate in general, some wave-based phenomena are not properly taken into account [14, 15].

Wave-based Room Acoustics

In the wave-based room acoustics modeling methods the goal is to numerically solve the appropriate wave equation with the given boundary conditions. Either the space is divided into volume elements or the surface is divided into boundary elements. The size of the elements determines how high frequencies can be modeled. Typically 6-10 elements per wavelength are required. In time domain methods, the sampling frequency must also fulfill the requirements of the samping theorem. Then the computational requirements can be proportional to the fourth power of the modeled frequency. This limits the use of wave-based methods to low frequencies, where the number of elements is computationally affordable. On the other hand, at the higher frequencies the benefits of the wave-based methods are practically lost, since if the source signal has sufficient variation, the high-frequency wave phenomena cannot be heard in a typical auralization set up. In most of the wave-based methods, modeling the boundary conditions requires more effort than in the geometrical acoustics, especially for a fully absorbing boundary. The fully absorbing boundary is required when modeling infinite spaces, since the computation must be limited to a finite space. Then to avoid reflections from the virtual boundary modeling infinity, it must absorb everything. [162]

One approach with volume elements is to estimate the gradients and derivatives in the wave equations by finite differences and evaluate the field only at one point inside the element. This leads to finite-difference time-domain (FDTD) techniques. One of them is the digital waveguide mesh [147, 126]. Each element affects its neighboring elements, and the relationships between the elements are derived by using finite differencies instead of differential operators in the wave equation. The computation is performed by updating the element values step-by-step in the time domain. The cubical elements used in the technique cause some problems. Modeling non-axis aligned walls is difficult without staircase-like approximations. Undesired dispersion also occurs since the propagation speed is different in different directions. Several improvements have been suggested, such as using interpolated meshes [145], different stencils for calculating the updated values [174] or tetrahedral elements [37, 38]. There is also frequency-related dispersion which can be decreased by using frequency-warping techniques [148, 149]. Modeling different boundary conditions also requires special techniques [94, 95]. The benefits of this type of algorithm is that it is relatively simple to implement and efficient compared to the other wave-based methods. Another very similar modeling technique is the transmission line method [81, 82, 96].

A more accurate solution can be achieved by the finite element method (FEM), where the volume elements can have arbitrary shapes and size, but the pressure is forced to be continuous across the boundaries and the field satisfies the wave equation inside the elements [33, 39]. Given the initial state, a system of equations can be solved to solve the whole sound field. The solving process is typically more expensive computationally than in the FDTD methods, but since the modeling is more accurate the elements can be larger. However, if the system of equations is written is a matrix form, the resulting relationship matrices are usually sparse, since the volume elements are mostly affected by their neighboring elements. This observation allows optimizations of the computational process.

Using boundary elements instead of volume elements reduces one dimension. In the boundary element methods (BEM) the open air is not explicitly modeled, but only the surfaces. The Kirchhoff-Helmholtz integral equation allows writing the sound field inside or outside a surface as a distribution of particle velocity and pressure on the surface. Thus, the surface can be discretized into elements and given the field quantities at the surface elements, the whole field can be constructed. By using the same kind of an integral formulation, relationships between the surface elements can be computed. Given the initial state, a similar system of equations can be solved as in the FEM. However, now the relationship matrix is typically full and non-symmetrical, which prevents using most optimization methods. The boundary elements can be effectively thought of as secondary sources and the solving process as finding the source signals that satisfy the boundary conditions. Again, it is assumed that there is a sufficient number of elements per modeled wavelength. It is possible to solve the field for point frequencies in the frequency domain or the impulse response in the time domain [91]. The iterative time domain computation can more be efficient than the frequency domain computation if the full-band impulse response must be modeled. The results of the previous time step are utilized when computing the next time step. The process is somewhat similar to the radiosity method in geometrical acoustics althought the quantities to be solved are different. [32, 165, 163]

2.3 Real-time Auralization

In real-time auralization, the acoustics modeling technique is an important part of the system, but sound rendering [164] is required to make the results audible. There are basically two approaches: convolving the modeled impulse response with a dry sound or rendering each reflection path separately [90]. Then there are several issues which must be taken into account. Especially, the chosen reproduction system determines which kind of signal processing is required. Interactivity sets some constraints to the system and special care must be taken when handling a moving listener. On the other hand, the applicability of the GPU to acoustic modeling have been noted by a few authors.

Reproduction

The spatial sound can be reproduced by using two or more channels. With two channels, headphones or two loudspeakers with cross-talk cancellation are used [10, 100]. There are multichannel reproduction systems such as vector base amplitude panning (VBAP) [135] or Ambisonics [120]. Yet another option is to use wave field synthesis [159] to reproduce the field.

When using headphones, it is necessary to model the effects of the listener's head, shoulders, and pinnae, since they modify the signal arriving at the eardrums [12]. These effects are modeled as a head-related transfer function (HRTF) which can be based on an artificially constructed model [9, 58, 21] or on measured data [53]. Simply interpolating the measured data is often impractical since the memory available is limited, and thus compressed models utilizing some kind of basis functions have been suggested. Examples are models that use principal component analysis [121, 88, 123], eigentransfer functions [22], and spherical harmonics [46]. When using these techniques, the input sound signal is typically convolved with a head-related impulse response. Another approach is to construct an appropriate finite impulse response (FIR) or infinite impulse response (IIR) filter [119] and apply it to the signal.

Signal Processing for Auralization

If the real-time auralization system handles the reflections separately, one common approach is to design filters which account for reflection from materials, distance, and air absorption [76]. Such an approach can be implemented efficiently, but explicit reflection paths are required.

When the results of the room acoustics modeling process are given as a room impulse response the convolution approach must be used. Then, it is necessary to handle the sound in blocks, since convolving the whole sound is not possible if the sound is streaming and in theory the length of the sound is infinite. The convolution can be efficiently performed in the frequency domain as a multiplication which requires applying the discrete Fourier transform to the source signal [2]. The blocks can be combined back to a continuous stream by using the overlap-add technique [4]. In addition, it should be noted that an appropriate window must be applied to avoid artifacts [65].

Interactive Auralization Systems

There are a few complete real-time auralization systems. The DIVA system utilizes two room acoustics modeling techniques [111, 109, 146]. The early reflections are modeled by the image source method and the late reverberation by an artificial reverberation algorithm. The diffraction phenomenon is also modeled by using the Biot-Tolstoy-Medwin model. A similar approach has been used in other auralization systems [83, 129].

A statistical model has been used for the late reverberation since the computation times using the methods modeling specular reflections increase rapidly as the reflection order increases. Another approach that has been suggested is to model the early part of the response with the image source method, but to use ray tracing for modeling diffuse reflections for the late part of the room impulse response [102]. In addition, the image source method can be improved by using a scene graph to prune the unnecessary computation in complex models.

GPU in Acoustic Computations

The use of the graphics processing unit (GPU) in acoustical computations has been researched. Gallo and Tsingos implemented two typical signal processing tasks required in acoustics modeling on the GPU [52]. The tasks were applying a variable delay-line and simple filtering. Although the given test cases did not improve the speed, they concluded that the GPU could speed up the computation compared to the CPU implementation in future graphics cards with more flexible architectures.

The GPU has also been used in a system computing the first order scattering and diffraction effects by utilizing the Helmholtz-Kirchhoff integral theorem and the Kirchhoff approximation [169]. This formulation allows using common computer graphics techniques, such as shadow mapping [34], displacement mapping [30], and mip-mapping [180], which can be effectively implemented on the GPU. The system allows computing solutions much faster than only with the CPU and also modeling very complex geometries offline.

The ray tracing technique has been implemented on the GPU and the modeled reflection paths have been rendered audible [79, 139]. Recent research in computer graphics on ray tracing has been utilized in the GPU implementation. Also the computation of convolutions and sound synthesis is done on the GPU by shaders. Even the sound data, geometry, and material definitions are stored as texture on the GPU memory. Only the final results are written back to the CPU for playback.

In addition, the GPU has been used for efficiently computing room acoustics with the waveguide mesh technique [138]. The entire computation is implemented as a fragment shader and it is shown that the GPU implementation is 1.5–12 times faster than the corresponding CPU implementation.

2.4 Summary of the Related Work

The related work on geometry reduction, room acoustics modeling, and real-time auralizarion was reviewed. Regarding geometry reduction, it was noted that there are both decimation and surface reconstruction algorithms. The decimation algorithms were based on removing elements of the existing surface while the surface reconstruction algorithms build a new surface. Since radical reduction is required for removing acoustically insignificant details, the surface reconstruction seems better suited for reduction of room geometry. Especially, volumetric approaches offer means of control over the detail level in reduction.

The room acoustic modeling methods in general were divided into wave-based methods and geometrical room acoustics. Although the wave-based methods are usually more accurate they are also computationally more demanding, and thus geometrical acoustics is more often used in auralization systems. In particular, for modeling specular reflections, the beam tracing approach seems to be the most efficient. There is also work on acoustic radiosity. There the computational requirements do not grow exponentially as the length of the modeled responses is increased, which is a problem in the image source based methods. However, the reflection model in the radiosity method limits its applicability. Overcoming this limitation would provide an efficient room acoustics modeling method.

Finally, the issues related to the real-time auralization were discussed. Reproduction systems can even affect the final stages of the room acoustics modeling. The signal processing must be efficient and the graphics processing unit could be utilized in the computation. There are a few existing real-time auralization systems, but they are usually capable of modeling accurately only the early part of the room impulse response.

Pressure is one of the most common quantities which are used to describe the sound field excited by a source. Since the field is often assumed time harmonic, modeling the frequency-dependent phase information becomes meaningful. Both magnitude and phase can be represented with a complex number pressure values. Other commonly modeled quantities are particle velocity and velocity potential. Velocity potential is not a physical quantity that can be measured, but is related to both pressure and particle velocity in a time harmonic field in a static homogenous medium as follows

$$\begin{aligned} \mathbf{u} &= -\nabla\phi \\ p &= \rho_0 \frac{\partial\phi}{\partial t} = j\omega\rho_0\phi, \end{aligned} \tag{3.1}$$

where $\mathbf{u}$ is the particle velocity, p is the pressure, ϕ is the velocity potential, t is time, ρ_0 is the density of air, ω is the angular velocity, and j is the imaginary unit.

However, modeling the phase information is not always necessary. Especially at high frequencies and in complex geometries where numerous propagating waves with different phases are present, the phase-related effects are less significant. Thus, it is often sufficient to model the intensity, which is related to the pressure and particle velocity as follows

$$\mathbf{I} = p\mathbf{u}^*, \tag{3.2}$$

where $\mathbf{u}^*$ stands for the complex conjugate of the particle velocity. Intensity is a vector quantity, as is also particle velocity, that can be represented by a direction and a magnitude. The direction corresponds to the direction of an acoustic energy flux and the magnitude corresponds to intensity of the energy flux.

3.1 Room Acoustic Rendering Equation

Publication [P3] presents an energy-based model for sound propagation in a room. Since intensity is the modeled quantity, phase information cannot be captured. The model is still strong since most geometrical room acoustics modeling methods can be seen as special cases of the presented model.

An important part of the model is the reflection model, which is defined by means of radiance and irradiance. Irradiance $E(x, \Omega)$ is the incident power on point x on a surface from direction Ω, and radiance $L(x, \Omega)$ is the radiant power per unit projected area per unit solid angle. These are differential quantities related to the intensity, which can also be defined as the incident power per surface area. The reflection model is represented by a bidirectional reflectance distribution function

$$\rho(\Omega_i, \Omega_e; \mathbf{x}) = \frac{\mathrm{d}L(\mathbf{x}, \Omega_e)}{\mathrm{d}E(\mathbf{x}, \Omega_i)}, \tag{3.3}$$

where Ω_i and Ω_e are the incident and outgoing directions of the radiant power. The model is energy-based and allows only local reaction, but is otherwise very flexible, since the distribution can be arbitrary.

On the other hand, the propagation of sound can be described by an operator which takes into account the medium absorption and the propagation delay:

$$\hat{S_r} I(t) = e^{\alpha r} S_r I(t) = e^{\alpha r} I\left(t - \frac{r}{c}\right), \tag{3.4}$$

where $\hat{S_r}$ is the propagation operator, S_r is the corresponding operator for non-absorptive medium, α is the absorption coefficient, r is the distance from the source, and c the speed of sound. The operator is linear and the medium absorption can be separated, so that when the operator is applied several times, the absorption and delay part can be calculated separately and combined in the end. Thus, in the following discussion the medium absorption is ignored.

The effects of geometry on the sound propagation can be derived by examining the differential quantities in the case of radiation between two surface points. Let there be an infitesimal area $\mathrm{d}A_{\mathbf{x}}$ around point $\mathbf{x}$. The differential solid angle covered by that area as seen from point $\mathbf{y}$ is

$$\mathrm{d}\Omega = \frac{\mathrm{d}A_{\mathbf{x}} \cos\theta_o}{|\mathbf{x} - \mathbf{y}|^2}, \tag{3.5}$$

where θ_o is the angle between the surface normal at point $\mathbf{x}$ and the outgoing radiance. Let the direction of the corresponding irradiance arriving at $\mathbf{y}$ be denoted Ω_i. The differential radiance reflected from $\mathbf{y}$ to direction Ω_e is according to the definition of the bidirectional reflectance distribution function

$$\mathrm{d}L(\mathbf{y}, \Omega_e) = \rho(\Omega_i, \Omega_e; \mathbf{y})\mathrm{d}E(\mathbf{y}, \Omega_i). \tag{3.6}$$

The differential irradiance arriving at point $\mathbf{y}$ can be expressed as

$$\mathrm{d}E(\mathbf{y}, \Omega_i) = S_{|x-y|} L(\mathbf{x}, -\Omega_i) \cos\theta_i \mathrm{d}\Omega. \tag{3.7}$$

Combining the three equations above gives

$$\mathrm{d}L(\mathbf{y}, \Omega_e) = \rho(\Omega_i, \Omega_e; \mathbf{y}) S_{|\mathbf{x}-\mathbf{y}|} L(\mathbf{x}, -\Omega_i) \cos\theta_i \cos\theta_o \frac{\mathrm{d}A_{\mathbf{x}}}{|\mathbf{x} - \mathbf{y}|^2}. \tag{3.8}$$

The effects of geometry can be collected into a geometry term

$$g(\mathbf{x}, \mathbf{y}) = \cos\theta_i \cos\theta_o \frac{S_{|\mathbf{x}-\mathbf{y}|}}{|\mathbf{x} - \mathbf{y}|^2} \tag{3.9}$$

and the reflected differential radiance is written as

$$\mathrm{d}L(\mathbf{y}, \Omega_e) = \rho(\Omega_i, \Omega_e; \mathbf{y}) g(\mathbf{x}, \mathbf{y}) L(\mathbf{x}, -\Omega_i) \mathrm{d}A_{\mathbf{x}}. \tag{3.10}$$

Let $V(\mathbf{x}, \mathbf{y})$ be a visibility function which is one when the path from point $\mathbf{x}$ to point $\mathbf{y}$ is unobstructed and zero otherwise. Then one can write a reflection kernel

$$R(\mathbf{x}, \mathbf{y}, \Omega_e) = V(\mathbf{x}, \mathbf{y}) \rho\, (\Omega_i, \Omega_e, \mathbf{y})\, g(\mathbf{x}, \mathbf{y}). \tag{3.11}$$

Integrating over all the surfaces G the reflected differential radiance and utilizing the reflection kernel, the room acoustic rendering equation can be written:

$$L(\mathbf{y}, \Omega_e) = L_0(\mathbf{y}, \Omega_e) + \oint_G R(\mathbf{x}, \mathbf{y}, \Omega_e) L(\mathbf{x}, -\Omega_i) \mathrm{d}\mathbf{x}, \tag{3.12}$$

where $L_0(\mathbf{y}, \Omega_e)$ is the radiance irradiated by the surface itself at point $\mathbf{y}$ in direction Ω_e. The integral notation refers to integrating over whole surface area by using some two-dimensional parametrization. A Neumann series solution to this equation can be written as

$$\begin{aligned} L_{n+1}(\mathbf{y}, \Omega_e) &= \oint_G R(\mathbf{x}, \mathbf{y}, \Omega_e) L_n(\mathbf{x}, -\Omega_e) \mathrm{d}\mathbf{x} \\ L(\mathbf{y}, \Omega_e) &= \sum_{n=0}^{\infty} L_n(\mathbf{y}, \Omega_e). \end{aligned} \tag{3.13}$$

This formulation can be seen as a reflection-iterative solution to the room acoustics rendering equation.

Most geometrical room acoustics modeling methods can be seen as special cases of the presented model. In image source-based methods, the bidirectional reflectance distribution function is such that it is non-zero only in the specular direction. In the acoustic radiosity methods, the function is a constant. In ray-tracing techniques, the bidirectional reflectance distribution function corresponds to the probability distribution of the direction of the reflected rays at the surfaces.

3.2 Acoustic Radiance Transfer Technique

The room acoustic rendering equation can be discretized to derive an element-based modeling algorithm called acoustic radiance transfer technique. The surface of the geometric model is divided into N patches. Then the Neumann series terms can be written as a sum of integrals over the patches:

$$L_{n+1}(\mathbf{y}, \Omega_e) = \sum_{i=1}^{N} \int_{A_i} R(\mathbf{x}, \mathbf{y}, \Omega_e) L_n(\mathbf{x}, -\Omega_e) \mathrm{d}\mathbf{x}. \tag{3.14}$$

The left side of the equation can also be expressed for patches by using the average reflected radiance,

$$L_{n,j}(\Omega_e) = \frac{1}{A_j} \int_{A_j} L_n(\mathbf{y}, \Omega_y) \mathrm{d}\mathbf{y}, \tag{3.15}$$

which yields an approximation where the radiation is assumed invariant over a patch

$$L_{n+1,j}(\Omega_e) = \frac{1}{A_j} \sum_{i=1}^{N} \int_{A_j} \int_{A_i} R(\mathbf{x}, \mathbf{y}, \Omega_e) L_{n,i}(-\Omega_e) \mathrm{d}\mathbf{x} \mathrm{d}\mathbf{y}, \tag{3.16}$$

where x is on patch i and y is on patch j. The direction can also be discretized

$$L_{n,k,j} = \frac{1}{\int_{\Phi_k} \mathrm{d}\Omega} \int\limits_{\Phi_k} L_{n,j} \mathrm{d}\Omega, \tag{3.17}$$

where Φ_k is the solid angle covered by the directional segment k. Then an approximation where the radiation is assumed invariant over the directional segment is

$$L_{n+1,j,k} = \frac{1}{A_j \int_{\Phi_k} \mathrm{d}\Omega} \sum_{i=1}^{N} \int\limits_{\Phi_k} \int\limits_{A_j} \int\limits_{A_i} R(\mathbf{x}, \mathbf{y}, \Omega_e) L_{n,i,\Gamma_i(-\Omega_e)} \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\mathrm{d}\boldsymbol{\Omega}_\mathbf{e}, \tag{3.18}$$

where operator $\Gamma_i(-\Omega_e)$ maps the direction $-\Omega_e$ on patch i into a directional segment index. This can be written in a clearer form by introducing the discretized reflection kernel

$$R_{i,j,k} = \frac{\int_{\Phi_k} \int_{A_j} \int_{A_i} R(\mathbf{x}, \mathbf{y}, \Omega_e) \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y}\mathrm{d}\boldsymbol{\Omega}_\mathbf{e}}{A_j \int_{\Phi_k} \mathrm{d}\Omega}, \tag{3.19}$$

which gives

$$L_{n+1,j,k} = \sum_{i=1}^{N} \sum_{l \in \Psi} R_{i,j,k} L_{n,i,l}, \tag{3.20}$$

where indices $l \in \Psi$ correspond to a set of directional segment indices produced by the operator $\Gamma_i(-\Omega_e)$ over the surface integrals. This operator is piecewise constant over the patches, and since the integration is a linear operator, the surface integrals can thus be expressed as a sum of constant values over the regions where the integrand is constant.

This reflection iterative formulation allows the whole solution to be written as

$$L_{j,k} = L_{0,j,k} + \sum_{n=0}^{\infty} \sum_{i=1}^{N} \sum_{l \in \Psi} R_{i,j,k} L_{n,i,l}. \tag{3.21}$$

The time dependence is implicitly modeled with the time delay operator in the reflection kernel.

The acoustic radiance transfer method evaluates this sum directly. The $n = 0$ values are determined by the direct radiance from the sound source, which is reflected at the patches. Then the propagation of the radiation is iterated for increasing values of n until the transfered radiance has attenuated below a desired threshold. Finally, the radiance $L_{j,k}$ can be collected from the patches to a listener.

Publication [P1] describes a method for geometry reduction. It is based on volumetric decomposition and remeshing of the surface geometry. The goal of the presented algoritms is to preserve the most important acoustical properties, such as volume and absorption area, while significantly simplifying the geometry. It is shown that the reduction process produces reasonably good results even with relatively high reduction rates.

4.1 Algorithm

The room geometry is inserted in a hierarchical volumetric structure called octree [144]. The smallest cell size in the octree is chosen according to the size of the details to be preserved. The structure, but not the original geometry, is used for reconstruction of the surface. This approach allows topology simplification, which is necessary for more radical reduction rates.

For the reconstruction, a variant of the marching cubes algorithm [112] is used. The marching cubes algorithm constructs an isosurface density from values contained in a regular grid. Thus, the algorithm has to be modified so that it can access the hierarchical structure of the octree and skip empty cells. In addition, since there are no density values, the surface is placed according to the occupancy information of the octree cells.

The result of the surface reconstruction phase is a surface consisting of a large number of small polygons. The planar polygons are merged by using the algorithm by Hinker and Hansen [68]. Thus, the result is a surface consisting of large polygons, and there are no details smaller than the dimensions of the smallest octree cells.

4.2 Evaluation and Results

The quality of the results produced by the geometry reduction algorithm was tested in the case of the concert hall model pictured in Fig. 4.1. Room acoustic parameters such as early decay time (EDT), clarity (C_{80}), definition (D_{50}), sound pressure level (SPL), center time (TS), and lateral fraction (LF) were extracted using the room acoustics prediction software ODEON. The results, found in publication [P1], show that the acoustic properties are preserved in up to 80 % reduction rates.

The computation times of the reduction were tested with a larger model and the results for the different reduction rates are shown in Table 4.1. It can be seen that the reduction is efficient enough to be used interactively, e.g., when designing concert halls. In addition, coarser models can be achieved more quickly.

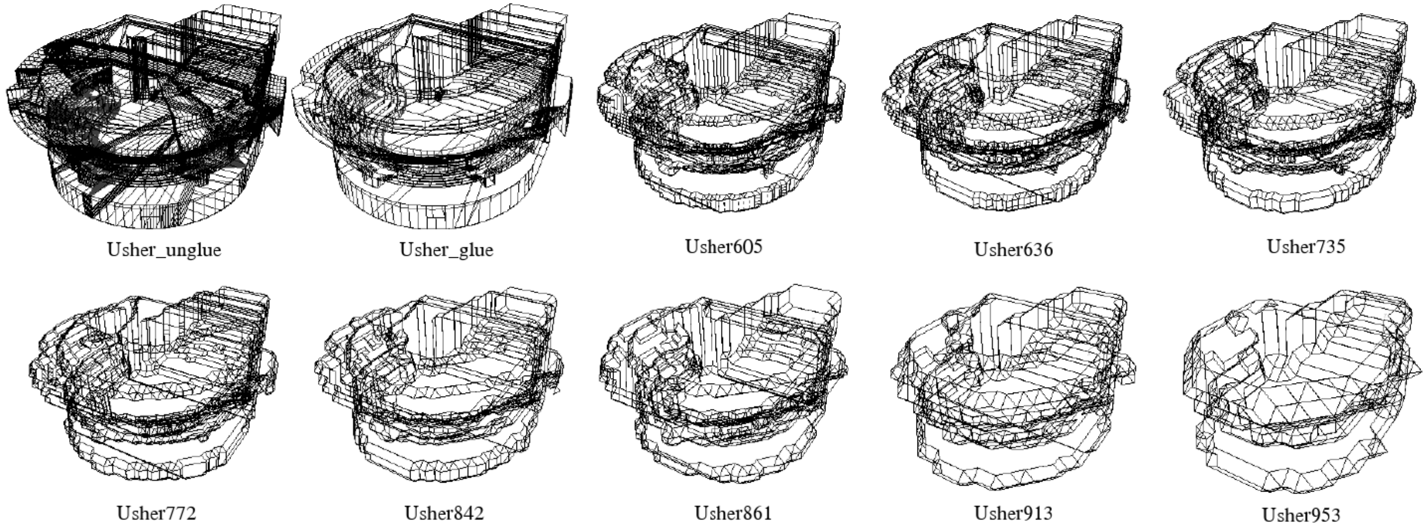

Figure 4.1: Reduction results for different reduction rates are shown.

Table 4.1: Reduction times with different volume raster resolutions. The reduction algorithm was run on a PC with a 2.8 GHz Pentium IV processor and 1 GB of RAM. The original model consisted of 119434 polygons.

Reduction (%)	Topology simpl. (s)	Surface simpl. (s)	Total (s)
70.7	25	82	107
79.0	23	58	81
79.6	21	37	58
83.3	20	24	44
87.2	15	14	29
91.5	6	5	11
94.2	5	3	8
97.0	4	1	5
99.2	2	1	3

There are two different room acoustics modeling algorithms, each of which can be used for different purposes, discussed as a result of this thesis. The outputs of the algorithms are either specular reflection paths or time-dependent energy responses. The final desired output is either room acoustic parameters or auralized sound.

Publication [P2] describes an optimized beam tracing algorithm. The output of the algorithm consists of specular reflection paths. In the precomputation phase, the goal is to construct a volumetric structure which significantly reduces the time required for run-time computations. This structure is called a beam tree.

Publication [P3] presents the acoustic radiance transfer method. The output of the algorithm is a time-dependent energy response from a source to a receiver. The majority of the computation can be done without knowing the receiver position, which allows very fast run-time computation for a moving listener. The precomputed data consists of time- and angle-dependent energy responses on surface patches.

In theory, it would be possible to precompute only the acoustic transfer matrix between the surface patches in acoustic radiance transfer technique. When an interaction matrix F for one reflection is computed, the total element-wise transfer matrix A could be computed as

$$A = \sum_{n=1}^{\infty} F^n. \tag{5.1}$$

Then the total response b from a source to a receiver could be computed as

$$b = \mathbf{s}A\mathbf{r}^{\mathbf{T}}, \tag{5.2}$$

where vector $\mathbf{s}$ contains the responses from the source to the elements and vector $\mathbf{r}$ contains the responses from the elements to the receiver. Here, the element-wise multiplication should be understood as convolution between the responses. Thus, the sound source could also be moving since sending an impulse from the source to the elements would be a fast operation. However, memory requirements set limits to the matrix size and, in practice, the sound source has to be fixed, resulting in a precomputed vector

$$c = \mathbf{s}A, \tag{5.3}$$

which takes less memory than the matrix.

5.1 Beam Tree Computation

In the beam tracing algorithm, the visibility information is precomputed and stored in a beam tree. First, the geometry is inserted in a binary space partitioning (BSP) tree for efficient queries to the geometry. Then, the first level of the tree is constructed by creating cones such that the apex is the

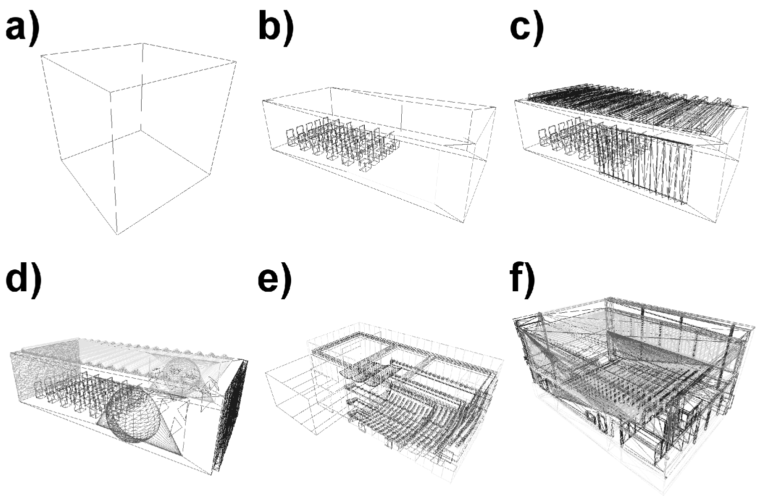

Figure 5.1: Models used in the performance tests: a) box, b) simple room, c) regular room, d) complex room, e) concert hall, and f) auditorium.

source and bases are the polygons. The beams are defined by the planes that intersect both the source and each edge of the base polygon. To construct the next level, the source is mirrored at the planes of each polygon to create an image source. Then the original beams are also reflected at those planes and split so that a new beam is created for each polygon intersecting the reflected beam. These new beams are defined by the planes intersecting the image source and the base polygon edges. The base polygons are clipped against the beam for the beam plane construction, but the clipped polygons are not saved. The beams are further reflected level by level until the desired reflection depth is reached. Only the polygon indices are saved for each beam, which leads to a compact presentation of the beam tree.

For simple models, the precomputation can be done very efficiently. The precomputation times were tested with the six different models shown in Fig. 5.1. The corresponding polygon counts and precomputation times are listed in Table 5.1.

5.2 Acoustic Energy Transfer

The acoustic radiance transfer method is an element-based technique which traces the energy sent by a sound source. The elements are patches of the surface of the room. The dimensions of an element are assumed to be small compared to the distances to the other elements. Thus, the interaction between two elements can approximately be described by an attenuation factor and a time delay. The attenuation factor takes into account the air absorption [8], attenuation by distance, as well as geometric considerations such as occlusion. The time delay is the time it takes for sound to travel from one patch to another. Similar computations can be performed for source-to-patch and patch-to-receiver interactions. Some error results from the fact that the aforementioned assumption does not hold between nearby patches.

Table 5.1: For each of the six test models, the precalculation was performed for the 1st–6th reflection orders. The beam tree took too much memory with higher-order reflections in some models, which is why the results are not given in such cases.

Model	Polyg.	1st (s)	2nd (s)	3rd (s)	4th (s)	5th (s)	6th (s)
Cube	6	0.0001	0.0005	0.0022	0.010	0.030	0.0525
Simple R.	438	0.0033	0.130	1.57	12.41	85.99	449
Regular R.	1190	0.0094	0.950	20.93	294	-	-
Complex R.	5635	0.047	4.96	77.0	549	2824	-
Concert H.	12115	0.130	12.5	325	-	-	-
Auditorium	14472	0.143	91.2	-	-	-	-

However, if the patches are small enough, the temporal spreading is neglible. In addition, more acccurate formulas or Monte-Carlo sampling can be used for the form factor computation to minimize the error.

The result of shooting energy from a sound source to the elements according to the attenuation coefficient and the delay is that every element has an energy value stored for one moment in time. In addition, this energy value should be multiplied by the reflection coefficient of the material of the element. The reflection coefficient might be direction-dependent and thus different values must be stored for different directions.

Then, gathering the direction-dependent energy values from the elements to a receiver according to the attenuation coefficient and the delay results in time-dependent energy responses at the receiver. This would represent the first-order reflections from the room surfaces.

However, it is possible to shoot the energy responses from one element to another. Shooting the energy once from all elements to all other elements and storing also the new reflectance-weighted energy contribution would result in storing the first two orders of reflections at the elements. Similarly, higher order reflections can be modeled by repeating the shooting process and accumulating the results at the elements. Eventually, the elements contain time-dependent responses which correspond to the energy reflected by the elements. Gathering the responses from the elements to a receiver yields a full energy response from the source to the receiver.

The acoustics energy transfer matrix elements are the attenuation factor – time delay pairs for each pair of model elements. In the precomputation, the power series of the acoustics transfer matrix can be computed and multiplied by the source projection vector, the result being an acoustic transfer vector. The elements of the acoustic transfer vector correspond to the energy responses on the elements.

Table 5.2 contains computation times on a PC with a Pentium IV 2.8 GHz processor and 1024 MB RAM for the geometries presented in [15], see Fig. 5.2. The extraction of the final responses is included in the times and in this case the goal was to compute acoustic parameters.

On the other hand, the precomputation can be done in the frequency

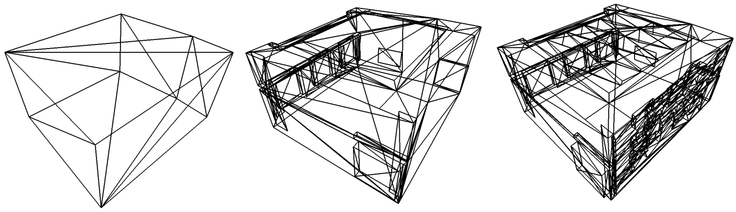

Figure 5.2: Performance of the computation of the acoustics radiance transfer evaluated in three different room geometries with 90, 378, and 986 patches.

domain, and then the goal is real-time auralization. The precomputation is practically the same as before, but discrete Fourier transforms of the responses are used instead of time-domain responses. The room models used as test cases in the frequency domain computation are shown in Fig. 5.3. The precomputation times are shown in Table 5.3.

Table 5.2: Computation times for each octave band with phase 1, phase 2, and phase 3 geometries consisting of 90, 378, and 986 patches, respectively, using angular resolution of 30 degrees for both azimuth and elevation angles, and a time resolution of 1 ms. The length of the response is 1.0 s, and the energy level for convergence is set to 10^{-4} of the initial energy, which corresponds to 40 dB attenuation. Double precision is used for floating points.

Octave Band	Phase 1	Phase 2	Phase 3
125 Hz	5 min 33 s	24 min 42 s	1 h 41 min 43 s
250 Hz	5 min 28 s	31 min 13 s	2 h 13 min 5 s
500 Hz	5 min 28 s	28 min 57 s	2 h 26 min 25 s
1 kHz	5 min 27 s	24 min 7 s	2 h 10 min 25 s
2 kHz	5 min 14 s	27 min 32 s	2 h 26 min 56 s
4 kHz	4 min 36 s	23 min 50 s	2 h 7 min 58 s

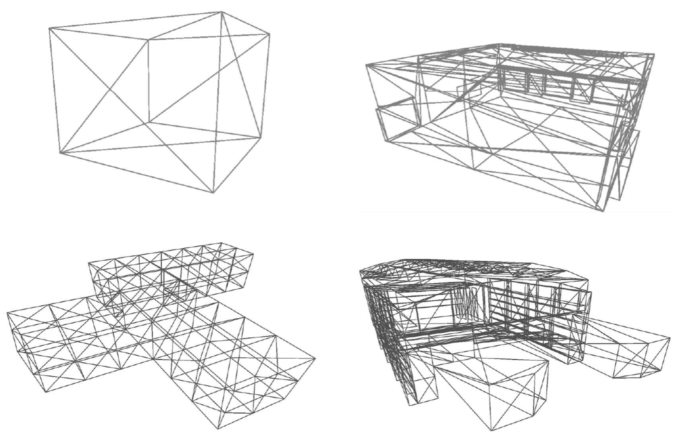

Figure 5.3: Performance of the precomputation of the frequency domain acoustics radiance transfer evaluated in four different room geometries: a cube (12 patches), a corridor model (256 patches), a small hall model (243 patches) and a concert hall model (1176 patches).

Table 5.3: Pre-computation times for the frequency domain acoustic radiance transfer. The number of iterations was set to 100 times the number of patches.

Model	Patches	Initial Shooting (s)	Iter. Propagation
Cube	12	0.03	14 min 56 s
Corridor	256	0.99	10 h 7 min 17 s
Small hall	243	1.02	3 h 7 min 5 s
Concert Hall	1176	7.15	99 h 15 min 32 s

Publication [P2] presents path validation optimizations for the beam tracing algorithm which speed up the run-time computation. The utilization of the transfer matrices for acoustics radiance transfer technique can be found in publication [P3]. Finally, implementation of a real-time auralization system utilizing the GPU is described in publication [P4].

6.1 Path Validation Optimization

Since occlusion was not taken into account, the reflection paths produced with the beam tracing algorithm must be validated. The beams could have been clipped against the occluding objects to produce only valid paths in the first place. Unfortunately, that would have caused the beam tree to branch more, which slows down the calculation, and the accurate clipping is algorithmically more complex. In addition, in the chosen approach most of the path validation computation can eventually be avoided.

An important part of the validation is checking that the path, defined by a start point, an end point, and reflection points in between, is inside the beams. The beams are cones defined by a set of limiting planes. Thus the computation involves checking on which side of a plane a reflection point is. This can be done by computing the distance to each plane, and a negative distance indicates that the point is on the backside of the plane and thus out of the beam. Because the beam structure is hierarchical, if a point falls outside a beam, the rest of the path in the rest of the branch below that beam is invalid. Since the point is outside a beam, there must be a plane against which the test failed. The idea of the optimization is to propagate that plane down the beam tree by mirroring it at the surfaces one by one like beams or image sources. Then, at each lower level the mirrored plane corresponds to the plane against which the path would eventually be tested at a higher level and testing against it at a lower level corresponds reordering the tests so that the higher level test is performed earlier. Then, when testing other paths in the same branch, the plane that has caused a path to be invalid earlier, is likely to cause the path to be detected invalid again, and this can be done early in the process without the need to traverse the tree further. This optimization saves a lot of futile computation.

Another related optimization for a moving listener is that the beams are grouped in “buckets”, and if all paths through the beams in a bucket prove to be invalid, then the distances to all the planes that caused the path to be invalid are computed. The shortest distance determines the radius of a sphere and while the listener is inside this sphere, all the paths in the bucket are invalid and no additional computation is required.

The same models were tested as in the precomputation phase. The performance results are listed in Table 6.1. It can be seen that the optimizations speed up the computation significantly and real-time modeling of a moving listener is possible in relatively complex room models.

Table 6.1: For each test model reflections up to the third order were calculated, except for the auditorium model for which only second order reflections could be calculated, while the listener was moving along a predefined path around the source. The average number of reflections (Refl.) as well as the calculation time for the algorithm are given without optimizations (Unopt.), with fail plane (FP) optimization only, and with both the fail plane optimization and skip spheres (FP + SS). For the optimizations, the speed-up factors are given in parenthesis.

Model	Polyg.	Refl.	Unopt. (s)	FP (s)	FP + SS (s)
Cube	6	63.0	0.0014	0.0008 (1.75)	0.0008 (1.00)
Simple R.	438	45.5	0.358	0.009 (39.8)	0.004 (1.50)
Regular R.	1190	34.1	3.03	0.057 (53.2)	0.019 (3.00)
Complex R.	5635	28.5	5.29	0.114 (46.4)	0.048 (2.38)
Concert H.	12115	35.8	38.2	0.943 (40.5)	0.489 (1.93)
Auditorium	14472	8.78	21.1	0.475 (44.4)	0.229 (2.07)

6.2 Utilizing Transfer Matrices

The transfer matrices contain time-dependent responses for the interaction between each pair of surface elements. Often the transfer matrices are too large to fit in the memory and thus they are multiplied with the vector of the source signals which have been propagated from the source to the elements. Thus, the elements contain responses for the acoustics transfer from the source to the elements.

The final phase involves gathering the responses from the elements to the receiver. In addition to the regular distance attenuation and delay, there might be a receiver detection pattern, such as HRTF, with which the incoming responses must be convolved.

The responses can also be gathered for acoustic parameter computation by using microphone directivities as detection patterns. This is demostrated in a test case of a music recording studio used in the third round robin on room acoustical computer simulation. This room was chosen since both measured data [14] and predictions by other room acoustics algorithms [15] exist. The comparison was done by using widely used room acoustical parameters which are the same as in the evaluation of the geometry reduction method, but SPL is replaced by strength (G). Figures 6.1-6.3 show the results for phase 2 geometry in [15] (the corresponding figures in that publication are given in paranthesis).

6.3 Real-time Implementation on the GPU

The graphics processing units in modern computers are very efficient parallel computers. Thus executing the same instruction for every element of a large data vector can be performed quickly by giving each parallel processor a different element to process. The final phase of the acoustic radiance transfer method is thus ideal for GPU processing, since it involves gathering

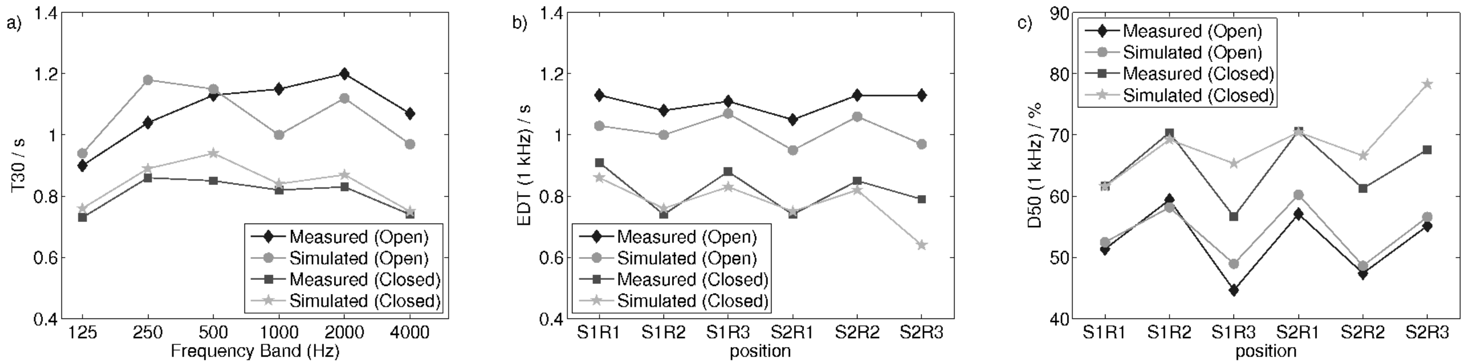

Figure 6.1: Charts show (a) T_{30} values at position S2R2, (b) EDT values in the 1 kHz octave band, and (c) D_{50} values in the 1 kHz octave band. (continued in the next figures)

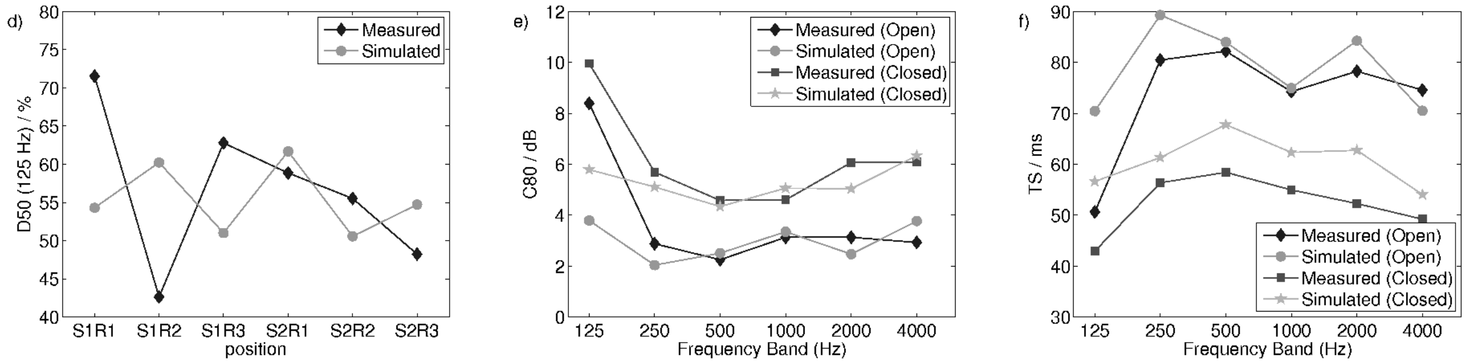

Figure 6.2: Charts show (d) D_{50} values in the 125 Hz octave band, (e) C_{80} and (f) TS values at position S1R1.

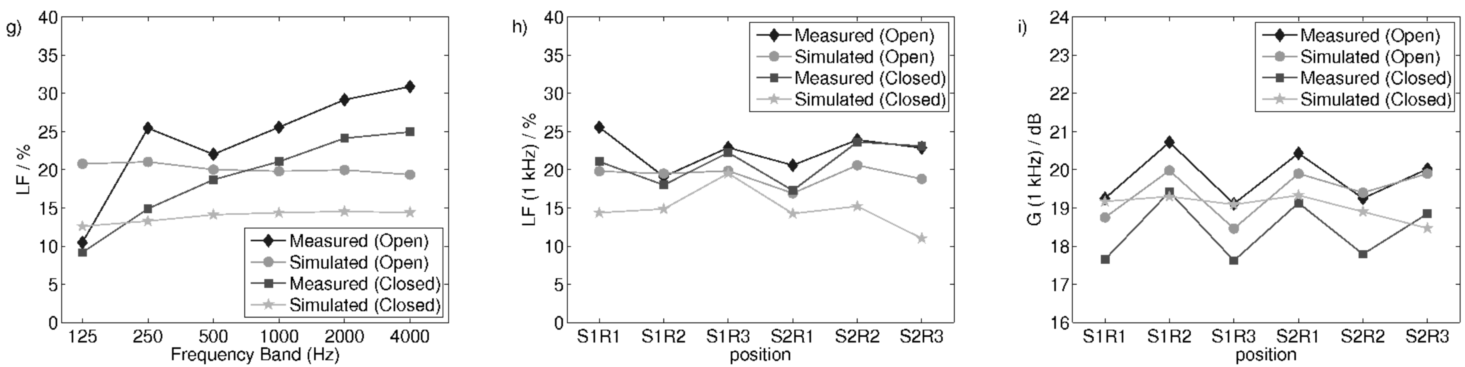

Figure 6.3: Charts show (g) LF values at position S1R1, and (h) LF and (i) G in the 1 kHz octave band.

Table 6.2: Run-time computation times for the test models illustrated in Fig 5.3. The time used for accumulating the responses, the total time used by the GPU including memory transfers, and the frame rate are shown.

Model	Accumulation (ms)	GPU total (ms)	Update rate (Hz)
Cube	4.31	26.3	30.1
Corridor	11.2	40.4	19.8
Small Hall	12.7	43.0	19.5
Concert Hall	61.5	138	6.10

the room responses from a large number of patches and applying an HRTF model to each of them.

In the precomputation phase a time-dependent room energy response for a fixed sound source was stored at each patch. These responses are uploaded into the GPU memory. If the memory required by the responses exceeds the GPU memory capabilities, the GPU memory is used as a cache for the responses and it is updated by always keeping the most recently used responses in the memory.

The listener can move during the run-time which should be reflected in the auralized sound. Thus, the parameters for each patch are re-computed regularly. These parameters include the distance to the listener, the air absorption, and the index of the direction-dependent patch response to use in the computation. These can be computed quickly, and the processing is performed on the CPU. The parameters are also uploaded to the GPU memory.

The GPU processing is then invoked while the CPU can continue to process new input. In the GPU kernel, each processor takes a small block of the responses and performs the required operations on it. Then the blocks are accumulated in two output responses, one for each ear. The operations include upsampling the responses since the output sample rate is higher than the sample rate of the patch responses, scaling and delaying the responses according to the parameters, and finally convolving them with the HRTF responses. Since the processing is performed with frequency domain responses, the convolution is transformed into a multiplication which can be done efficiently. The output responses are written back to the CPU memory. Figure 6.4 illustrates the process.

Finally, there is a signal processing thread that takes one block of the input sound stream, applies the discrete Fourier transform (DFT) to it, multiplies it with the output responses of the GPU computation, and takes an inverse DFT of the two responses. This yields auralized blocks of sound for both ears. The continuous output stream is constructed by using the overlap-add approach to the blocks.

The run-time computation times for the same models as used in the precomputation phase are listed in Table 6.2, showing that the performance is within the limits of an interactive rate even with the most complex model.

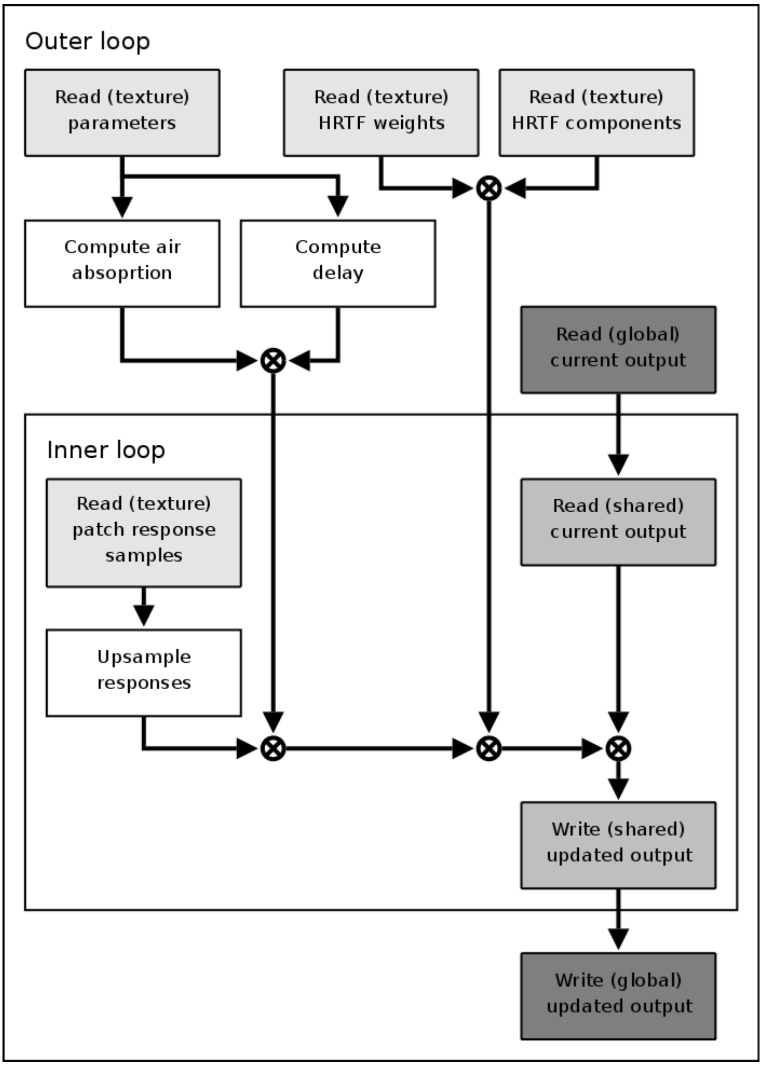

Figure 6.4: Computation pipeline in the GPU kernel. The operations in the outer loop are performed once every 256 samples, while the operations in the inner loop are performed for every sample. Note that the shared memory is used as a cache for global data so that slow global memory reads and writes can be avoided in the inner loop. Texture memory is automatically cached which makes texture reads fast.

7 SUMMARY

The goal of the thesis has been to show that real-time auralization utilizing physics-based room acoustics modeling techniques is possible in relatively complex rooms with arbitrary material properties. A method for reducing the room geometries has been introduced in addition to novel room acoustic modeling techniques. It has been shown that most of the modeling can be done as precomputation, which allows fast run-time computation, possibly utilizing the GPU. However, there is still room for future research.

7.1 Main Results of the Thesis

The main results of the thesis can be summarized as follows:

- A geometry simplification algorithm for producing lower detail level room geometry for acoustics modeling
- Analysis of the optimized beam tracing algorithm showing its applicability to real-time room acoustics modeling
- The acoustic radiance transfer method, based on the room acoustics rendering equation, which is capable of modeling arbitrary reflections and late reverberation
- A real-time auralization system for interactive walk-throughs in complex room models utilizing the acoustic radiance transfer method implemented partially on the GPU

7.2 Future Work

There remains plenty of research related to the topics presented in this thesis. Some of the possible topics are:

- Improving the geometry reduction by taking into account the original surface orientations
- Improving the bucket construction in the beam tracing algorithm
- Finding additional optimizations to the beam tracing technique based on the path coherence
- Physics-based construction of new reflection models for the acoustic radiance transfer method
- Diffraction modeling with the acoustic radiance transfer method
- Modeling multiple sound sources in the acoustic radiance transfer method
- Applying compression to the precomputed responses in the acoustic radiance transfer method

8 SUMMARY OF PUBLICATIONS AND CONTRIBUTIONS OF THE AUTHOR

The publications included in this thesis have not previously formed part of another thesis. The main results of the thesis and the author's contributions can be summarized as follows.

Publication [P1]

A method for reducing the geometry of room models used in acoustics modeling is presented. The topology of the room is simplified in a process where the model is first decomposed into a volumetric structure. The surfaces are reconstructed by utilizing this structure and subsequently simplified by merging coplanar regions. The results of the method are verified by extracting room acoustical attributes from the original and reduced models with the ODEON room acoustics modeling software. It is shown that the most important acoustic properties have been preserved, even with relatively high reduction rates.

The author designed and implemented the algorithm, ran the reductions for the test models, and wrote 75 % of the paper. Dr. Tapio Lokki was responsible for computing the acoustic attributes of the test models with the ODEON software and writing the results section of the paper.

Publication [P2]

An optimized beam tracing algorithm for finding and efficiently updating specular reflection paths for a moving listener is presented. This algorithm performs well even with complex, lightly occluded room models. It is shown that even a moving sound source can be modeled at interactive rates with moderate model complexity. The proposed optimizations are based on the concept of a propagated fail plane. Utilizing the optimization techniques allows returning the negative results for validated paths quickly. Because most of the results are negative, this speeds up the computation significantly. The bucketing optimization utilizes the local coherence in the validation results. Together these optimizations can give a two-orders-of-magnitude speed-up compared to the unoptimized algorithm.

The optimizations were invented by Dr. Samuli Laine. He also implemented the original optimized algorithm. The author modified the algorithm so that it could be better used with acoustic spaces. He also did the performance analysis and wrote 90 % of the paper.

Publication [P3]

An integral equation generalizing a variety of known geometrical room acoustics modeling algorithms is presented. Based on the room acoustic rendering equation, an acoustic radiance transfer method, which can handle both diffuse and nondiffuse reflections, is derived. In a case study, the method is used to predict several acoustic parameters of a room model. The results are compared to measured data of the actual room and to the results given by other room acoustic prediction software. It is concluded that the method can predict most room acoustic parameters reliably and provides results as accurate as current commercial room acoustic prediction software.

The room acoustics rendering equation and its special cases were first presented in Sami Kiminki's M.Sc. thesis [86]. The author designed and implemented the acoustic radiance transfer method, ran the test cases, did the analysis, and wrote 90 % of the paper.

Publication [P4]

A method for modeling room acoustics and auralizing the results in real time with a moving listener is presented. The acoustics modeling is based on the acoustic radiance transfer technique. The novel idea of implementing this technique in the frequency domain allows modeling all frequencies at once as the time domain technique required separate runs for each frequency band. Since the auralization of the results requires scaling, adding, and delaying responses as well as convolving them with head-related impulse responses, the massive parallel computation capacity of modern graphics hardware is utilized. Thus, realistic interactive walk-throughs are possible in typical room models with a stationary source.

The author has done all the technical work and 90 % of the writing for this publication. The technical work includes deriving the theoretical basis, designing and implementing the algorithm, and running the test cases and analysing the results.

[1] María-Elena Algorri and Francis Schmitt. Mesh simplification. *Computer Graphics Forum*, 15(3):77–86, 1996.

[2] Jont B. Allen. Short term spectral analysis, synthesis, and modification by discrete fourier transform. *IEEE Transactions on Acoustics, Speech, Signal Processing*, 25(3):235–238, June 1977.

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

[4] Jont B. Allen and Lawrence R. Rabiner. A unified approach to short-time Fourier analysis and synthesis. *Proceedings of the IEEE*, 65(11):1558–1564, November 1977.

[5] Carlos Andújar and Pere Brunet. The discrete polyhedra simplification (dps): A framework for polyhedra simplification based on decomposition schemes. Technical Report LSI-99-3-R, Universitat Politécnica de Catalunya, 1999.

[6] Carlos Andújar, Pere Brunet, and Dolors Ayala. Topology-reducing surface simplification using a discrete solid representation. *ACM Transactions on Graphics*, 21(2):88–105, 2002.

[7] Chandrajit L. Bajaj and Daniel S. Schikore. Error-bounded reduction of triangle meshes with multivariate data. In *Proceedings of SPIE*, volume 2656, pages 34–45, 1996.

[8] H. Bass, H.-J. Bauer, and L. B. Evans. Athospheric absorption of sound: Analytical expressions. *Journal of the Acoustical Society of America*, 52(3B):821–825, 1972.

[9] Dwight W. Batteau. The role of the pinna in human localization. In *Proceedings of the Royal Society of London Series*, pages 158–160, 1968.

[10] Durand R. Begault. *3-D Sound for Virtual Reality and Multimedia*. Academic Press, Cambridge, MA, 1994.

[11] Maurice Anthony Biot and Ivan Tolstoy. Formulation of wave propagation in infinite media by normal coordinates with an application to diffraction. *Journal of the Acoustical Society of America*, 29(3):381–391, 1957.

[12] Jens Blauert. *Spatial Hearing. The Psychophysics of Human Sound Localization*. MIT Press, Cambridge, MA, 2nd edition, 1997.

[13] Jeffrey Borish. Extension of the image model to arbitrary polyhedra. *Journal of the Acoustical Society of America*, 75(6):1827–1836, 1984.

[14] Ingolf Bork. Report on the 3rd round robin on room acoustical computer simulation - Part I: Measurements. *Acta Acustica united with Acustica*, 91(4):740–752, 2005.

[15] Ingolf Bork. Report on the 3rd round robin on room acoustical computer simulation - Part II: Calculations. *Acta Acustica united with Acustica*, 91(4):753–763, 2005.

[16] Chris Buckalew and Donald Fussell. Illumination networks: Fast realistic rendering with general reflectance functions. *Computer Graphics*, 23(3):89–89, 1986.

[17] Paul Calamia. *Advances in Edge-Diffraction Modeling for Virtual-Acoustic Simulations*. PhD thesis, Computer Science, Princeton University, June 2009.

[18] Paul Calamia and U. Peter Svensson. Edge subdivion for fast diffraction calculations. In *Proc. 2005 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, October 2005.

[19] Paul Calamia, U. Peter Svensson, and Thomas Funkhouser. Integration of edge-diffraction calculus and geometrical-acoustics modeling. In *Proc. Forum Acusticum*, August 2005.

[20] Andrew Certain, Jovan Popovic, Tom DuChamp, David Salesin, Werner Stutzle, and Tony DeRose. Interactive multiresolution surface viewing. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 91–98. ACM Press, 1996.

[21] Jiashu Chen, Barry D. Van Veen, and Kurt E. Hecox. External ear transfer function modeling – a beamforming approach. *Journal of the Acoustical Society of America*, 92(4):1933–1944, 1992.

[22] Jiashu Chen, Barry D. Van Veen, and Kurt E. Hecox. A spatial feature extraction and regularization model for the head-related transfer function. *Journal of the Acoustical Society of America*, 97:439–450, 1995.

[23] Andrea Ciampalini, Paolo Cignoni, Carlo Montani, and Roberto Scopigno. Multiresolution decimation based on global error. *Visual Computer*, 13(5):228–246, 1997.

[24] Paolo Cignoni, Carlo Montani, and Roberto Scopigno. A comparison of mesh simplification algorithms. *Computer and Graphics*, 22(1):37–54, 1998.

[25] Jonathan Cohen, Dinesh Manocha, and Marc Olano. Simplifying polygonal models using successive mappings. In *Proceedings of the 8th conference on Visualization '97*, pages 395–402. IEEE Computer Society Press, 1997.

[26] Jonathan Cohen, Marc Olano, and Dinesh Manocha. Appearance-preserving simplification. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 115–122. ACM Press, 1998.

[27] Jonathan Cohen, Amitabh Varshney, Dinesh Manocha, Greg Turk, Hans Weber, Pankaj Agarwal, Brooks Frederik, and William Wright. Simplification envelopes. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 119–126. ACM Press, 1996.

[28] Michael F. Cohen, Shenchang Eric Chen, John R. Wallace, and Donald P. Greenberg. A progressive refinement approach to fast radiosity image creation. *Computer Graphics*, 22(3):75–84, 1988.

[29] Michael F. Cohen and Donald P. Greenberg. A radiosity solution for complex environments. *Computer Graphics*, 19(3):31–40, 1985.

[30] Robert L. Cook. Shade trees. *ACM SIGGRAPH Computer Graphics*, 18(3):223–231, 1984.

[31] Trevor J. Cox, Bengt-Inge L. Dalenbäck, Peter D'Antonio, Jean-Jacques Embrechts, Jin Yong Jeon, Eckart Mommertz, and Michael Vorländer. A tutorial on scattering and diffusion coefficients for room acoustic surfaces. *Acta Acustica united with Acustica*, 92(1):1–15, 2006.

[32] Jean-Pierre Coyette. An enhanced boundary element model for 3-D transient acoustics. In *3rd International Congress on Air- and Structure-Borne Sound and Vibration*, pages 823–830, June 1994.

[33] A. Craggs. A finite element method for the free vibration of air in ducts and rooms with absorbing walls. *Journal of Sound and Vibration*, 173(4):568–576, June 1994.

[34] Carsten Dachsbacher and Marc Stamminger. Reflective shadow maps. In *Proceedings of the 2005 symposium on Interactive 3D graphics and games*, pages 203–231, 2005.

[35] Bengt-Inge Dalenbäck. Room acoustic prediction based on a unified treatment of diffuse and specular reflection. *Journal of the Acoustical Society of America*, 100(2):899–909, 1996.

[36] Ian A. Drumm and Yiu W. Lam. The adaptive beam-tracing algorithm. *Journal of the Acoustical Society of America*, 107(3):1405–1412, 2000.

[37] Scott Van Duyne and Julius O. Smith. The tetrahedral digital waveguide mesh. In *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, October 1995.

[38] Scott Van Duyne and Julius O. Smith. The 3D tetrahedral digital waveguide mesh with musical applications. In *Proc. Int. Computer Music Conf.*, pages 9–16, August 1996.

[39] V. Easwaran and A. Craggs. Transient response of lightly damped rooms: A finite element approach. *Journal of the Acoustical Society of America*, 99(1):108–113, 1996.

[40] Matthias Eck, Tony DeRose, Tom DuChamp, Hugues Hoppe, Micael Lounsbery, and Werner Stuetzle. Multiresolution analysis fo arbitrary meshes. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 173–182. ACM Press, 1995.

[41] Jihad El-Sana and Amitabh Varshney. Controlled simplification of genus for polygonal models. In *IEEE Visualization '97 Conference Proceedings*, pages 403–410, 1997.

[42] Jihad El-Sana and Amitabh Varshney. Topology simplification for polygonal virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 4(2):133–144, 1998.

[43] Jihad El-Sana and Amitabh Varshney. Generalized view-independent simplification. *Computer Graphics Forum*, 18(3):83–94, 1999.

[44] Carl Erikson and Dinesh Manocha. Simplification culling of static and dynamic scene graphs. Technical Report TR98-009, University of North Carolina at Chapel Hill, March 1998.

[45] Carl Erikson and Dinesh Manocha. GAPS: General and automatic polygonal simplification. In S. N. Spencer, editor, *Proceedings of the Conference on the 1999 Symposium on Interactive 3D Graphics*. ACM Press, 1999.

[46] Michael J. Evans, James A. S. Angus, and Anthony I. Tew. Analyzing head-related transfer function measurements using surface spherical harmonics. *Journal of the Acoustical Society of America*, 104:2400–2411, 1998.

[47] Thomas A. Funkhouser, Ingrid Carlbom, Gary Elko, Gopal Pingali, Mohan Sondhi, and Jim West. A beam tracing approach to acoustic modeling for interactive virtual environments. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 21–32, 1998.

[48] Thomas A. Funkhouser, Patrick Min, and Ingrid Carlbom. Real-time acoustic modeling for distributed virtual environments. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 365–374, 1999.

[49] Thomas A. Funkhouser, Nicolas Tsingos, Ingrid Carlbom, Gary Elko, Mohan Sondhi, and Jim West. Modeling sound reflection and diffraction in architectural environments with beam tracing. In *Proceedings of Forum Acusticum 2002*, Sevilla, Spain, September 2002.

[50] Thomas A. Funkhouser, Nicolas Tsingos, Ingrid Carlbom, Gary Elko, Mohan Sondhi, Jim E. West, Gopal Pingali, Patrick Min, and Addy Ngan. A beam tracing method for interactive architectural acoustics. *Journal of the Acoustical Society of America*, 115(2):739–756, February 2004.

[51] Thomas A. Funkhouser, Nicolas Tsingos, and Jean-Marc Jot. Sounds good to me! Computational sound for graphics, VR, and interactive systems. SIGGRAPH 2002 Course Notes, 2002.

[52] Emmanuel Gallo and Nicolas Tsingos. Efficient 3D audio processing on the GPU. In *Proceedings of the ACM Workshop on General Purpose Computing on Graphics Processor*, August 2004.

[53] William G. Gardner and Keith D. Martin. HRTF measurements of a KEMAR. *Journal of the Acoustical Society of America*, 97(6):3907–3908, 1995.

[54] Michael Garland. Multiresolution modeling: Survey & future opportunities. In *State of the art report, Eurographics '99*, pages 111–131, 1999. Available as http://mgarland.org/files/papers/STAR99.pdf.

[55] Michael Garland. *Quadratic-based polygonal surface simplification*. PhD thesis, Computer Science Department, Carnigie Mellon University, 1999. Also available as http://mgarland.org/research/thesis.html.

[56] Michael Garland and Paul S. Heckbert. Multiresolution modeling: simplifying surfaces with color and texture using quadric error metrics. In *Proceedings of IEEE Visualization '98*, 1998. Also available as http://mgarland.org/files/papers/quadric2.pdf.

[57] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 209–216. ACM Press, 1998. Also available as http://mgarland.org/files/papers/quadrics.pdf.

[58] Klaus Genuit. A description of the human outer ear transfer function by elements of communication theory. In *Proceedings of the 12th International Congress on Acoustics*, 1986.

[59] Tran S. Gieng, Bernd Hamann, Kenneth I. Joy, Gregory Schussman, and Isaac J. Trotts. Smooth hierarchical surface triangulations. In *Proceedings of IEEE Visualization '97*, pages 379–386. IEEE Computer Society Press, 1997.

[60] Cindy M. Goral, Kenneth E. Torrence, Donald P. Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. *ACM SIGGRAPH Computer Graphics*, 18(3):212–222, 1984.

[61] Markus H. Gross, Oliver G. Staadt, and Roger Gatti. Efficient triangular surface approximation using wavelets and quadtree data structure. *IEEE Transactions on Visualization and Computer Graphics*, 2(2):130–143, 1996.

[62] André Guéziec. Surface simplification inside a tolerable volume. Technical Report RC 20440, IBM Research, Yorktown Heights, NY 10598, March 1996.

[63] André Guéziec. Locally toleranced surface simplification. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):168–189, 1999.

[64] Bernd Hamann. A data reduction scheme for triangulated surfaces. *Computer Aided Geometric Design*, 11(2):197–214, April 1994.

[65] Fredric J. Harris. On the use of windows for harmonic analysis with discrete Fourier transform. *Proceedings of the IEEE*, 66(1), January 1978.

[66] Taosong He, Lichan Hong, Amitabh Varshney, and Sidney W. Wang. Voxel based object simplification. In *Proceeding of Visualization '95*, pages 296–303, 1995.

[67] Taosong He, Lichan Hong, Amitabh Varshney, and Sidney W. Wang. Controlled topology simplification. *IEEE Transactions on visualization and computer graphics*, 2(2):171–184, 1996.

[68] Paul Hinker and Charles Hansen. Geometric optimization. In Gregory M. Nielson and Dan Bergeron, editors, *Proceedings of Visualization '93*, pages 189–195, 1993.

[69] Murray Hodgson and Eva-Marie Nosal. Experimental evaluation of radiosity for room sound-field prediction. *Journal of the Acoustical Society of America*, 120(2):808–819, 2006.

[70] Hugues Hoppe. Progressive meshes. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 99–108. ACM Press, 1996.

[71] Hugues Hoppe. Efficient implementation of progressive meshes. *Computers & Graphics*, 22(1):27–36, February 1998.

[72] Hugues Hoppe. New quadric metric for simplifying meshes with appearance attributes. In *Proceedings of the conference on Visualization '99: celebrating ten years*, pages 59–66. IEEE Computer Society Press, 1999.

[73] Hugues Hoppe, Tony DeRose, Tom DuChamp, Mark Halstead, Hubert Jin, John McDonald, Jean Schweitzer, and Werner Stuetzle. Piecewise smooth surface reconstruction. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 295–302. ACM Press, 1994.

[74] Hugues Hoppe, Tony DeRose, Tom DuChamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. *ACM SIGGRAPH Computer Graphics*, 26(2):71–78, 1992.

[75] Hugues Hoppe, Tony DeRose, Tom DuChamp, John McDonald, and Werner Stuetzle. Mesh optimization. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. ACM Press, 1993.

[76] Jyri Huopaniemi, Lauri Savioja, and Matti Karjalainen. Modeling of reflections and air absorption in acoustical spaces — a digital filter design approach. In *Proc. IEEE Workshop on Application of Signal Processing to Audio and Acoustics (WASPAA'97)*, Mohonk, New Paltz, New York, October 1997.

[77] David S. Immel, Michael F. Cohen, and Donald P. Greenberg. A radiosity solution for non-diffuse environments. *Computer Graphics*, 20(4):133–142, 1986.

[78] Gary M. Jebsen and Herman Medwin. On the failure of the Kirchhoff assumption in backscatter. *Journal of the Acoustical Society of America*, 72(5):1607–1611, 1982.

[79] Marcin Jedrzejewski and Krzysztof Marasek. Computation of room acoustics using programmable video hardware. *Computer Vision and Graphics*, 32:587–592, 2006.

[80] Michael J. DeHaemer Jr. and Michael Zyda. Simplifiaction of objects rendered by polygonal approximations. *Computer and Graphics*, 15(2):175–184, 1991.

[81] Yukio Kagawa. Computational acoustics – theories of numerical analysis in acoustics with emphasis on transmission-line matrix modelling. In *Proc. ASWA 97*, pages 19–26, 1997.

[82] Yukio Kagawa, Takao Tsuchiya, B. Fujii, and K. Fujioka. Discrete Huygens' model approach to sound wave propagation. *Journal of Sound and Vibration*, 218:419–444, 1998.

[83] Raine Kajastila, Samuel Siltanen, Peter Lundén, Tapio Lokki, and Lauri Savioja. A distributed real-time virtual acoustic rendering system for dynamic geometries. In *The AES 122th International Convention*, May 2007.

[84] Alan D. Kalvin and Russell H. Taylor. Superfaces: Polygonal mesh simplification with bounded error. *IEEE Computer Graphics and Applications*, 16(3):64–77, 1996.

[85] T. Kawai. Sound diffraction by a many-sided barrier or pilar. *Journal of Sound and Vibration*, 79(2):229–242, 1981.

[86] Sami Kiminki. Sound propagation theory for linear ray acoustic modelling. Master's thesis, Helsinki University of Technology, 2005.

[87] Wayne A. Kinney, C. S. Clay, and Gerald A. Sandness. Scattering from a corrugated surface: Comparison between experiment, Helmholtz-Kirchhoff theory, and the facet-ensemble method. *Journal of the Acoustical Society of America*, 73(1):183–194, 1983.

[88] Doris J. Kistler and Frederic L. Wightmann. A model of head-related transfer functions based on principal component analysis and minimum-phase reconstruction. *Journal of the Acoustical Society of America*, 91(3):1637–1647, 1992.

[89] Reinhard Klein, Gunther Lieblich, and Wolfgang Straßer. Mesh reduction with error control. In *Proceedings of Visualization '96*, pages 311–318, 1996.

[90] Mendel Kleiner, Bengt-Inge Dalebäck, and U. Peter Svensson. Auralization – an overview. *Journal of the Audio Engineering Society*, 41(11):861–875, November 1993.

[91] A. Kludszuweit. Time iterative boundary element method (TIBEM) – ein neues numerisches Verfahren der 4-dimensionalen Systemanalyse von Wellenvorgängen zur Berechnung der Raumimpulsantwort. *Acustica*, 75:17–27, 1991.

[92] Leif Kobbelt, Swen Campagna, and Hans-Peter Seidel. A general framework for mesh simplification. In *Proceedings of Graphics Interface '98*, pages 43–50, 1998.

[93] Robert G. Kouyoumjian and Prabhakar H. Pathak. A uniform theory of diffraction for an edge in a perfectly conducting surface. *Proceedings of the IEEE*, 62(11):1448–1461, 1974.

[94] Konrad Kowalczyk and Maarten van Walstijn. Modelling frequency-dependent boundaries as digital impedance filters in FDTD and K-DWM room acoustics simulations. *Journal of the Audio Engineering Society*, 56(7/8):569–583, 2008.

[95] Konrad Kowalczyk and Maarten van Walstijn. Virtual room acoustics using finite difference methods. How to model and analyse frequency-dependent boundaries? In *3rd International Symposium on Communications, Control and Signal Processing (ISCCSP 2008)*, pages 1504–1509, February 2008.

[96] Ulf Kristiansen and Karim Jezzine. TLM model for sound propagation above ground. In *29th International Congress and Exhibition on Noise Control Engineering*, 2000.

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

[98] Yiu W. Lam. A comparison of three diffuse reflection modeling methods used in room acoustics computer models. *Journal of the Acoustical Society of America*, 100(4):2181–2192, 1996.

[99] Bertrand Le Seac and Schlick Christophe. A progressive ray-tracing based radiosity with general reflectance functions. In *Proceedings of the Eurographics Workshop on Photosimulation, Realism and Physics in Computer Graphics, 1990*, pages 103–116, 1992.

[100] Hilmar Lehnert. Principles of binaural room simulation. *Applied Acoustics*, 36(3–4):259–291, 1992.

[101] Hilmar Lehnert. Systematic errors of the ray-tracing algorithm. *Applied Acoustics*, 38(2–4, Special Issue on Computer Modelling and Auralization of Sound Field in Rooms):207–221, 1993.

[102] Tobias Lentz, Dirk Schröder, Michael Vorländer, and Ingo Assenmacher. Virtual reality system with integrated sound field simulationa and reproduction. 2007(ID 70540):19, 2007.

[103] T. Lewers. A combined beam tracing and radiant exchange computer model of room acoustics. *Applied Acoustics*, 38(2–4, Special Issue on Computer Modelling and Auralization of Sound Field in Rooms):161–178, 1993.

[104] Peter Lindstrom. Out-of-core simplification of large polygonal models. In *Proceeding of the 27th annual conference on Computer graphics and interactive techniques*, pages 259–262, 2000.

[105] Peter Lindstrom. A memory insensitive technique for large model simplification. In *Proceedings of the conference on Visualization '01*, pages 121–126, 2001.

[106] Peter Lindstrom and Greg Turk. Fast and memory efficient polygonal simplification. In *IEEE Visualization '98 Conference Proceedings*, pages 279–286,544, 1998.

[107] Peter Lindstrom and Greg Turk. Evaluation of memoryless simplification. *IEEE Transactions on Visulization and Computer Graphics*, 5(2):98–115, 1999.

[108] Peter Lindstrom and Greg Turk. Image-drive simplification. *ACM Transactions on Graphics*, 19(3):204–241, July 2000.

[109] Tapio Lokki. *Physically-based Auralization – Design, Implementation, and Evaluation*. PhD thesis, Helsinki University of Technology, Department of Computer Science and Engineering, Telecommunications Software and Multimedia Laboratory, 2002.

[110] Tapio Lokki, Ville Pulkki, and Paul T. Calamia. Measurement and modeling of diffraction from an edge of a thin panel. *Applied Acoustics*, 69(9):824–832, 2008.

[111] Tapio Lokki, Lauri Savioja, Jyri Huopaniemi, Riitta Väänänen, and Tapio Takala. Creating interactive virtual auditory environments. *IEEE Computer Graphics and Applications*, 22(4):49–57, 2002.

[112] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.

[113] Michael Lounsbery, Tony DeRose, and Joe Warren. Multiresolution analysis for surfaces of arbitrary topological type. Technical Report 93-10-05b, Department of Computer Science and Engineering, University of Washington, 1994.

[114] Michael Lounsbery, Tony DeRose, and Joe Warren. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics*, 16(1):34–73, 1997.

[115] Anders Løvstad and U. Peter Svensson. Diffracted sound field from an orchestra pit. *Acoustic Science and Technology*, 26(2), 2005.

[116] Kok-Lim Low and Tiow-Seng Tan. Model simplification using vertex-clustering. In *SI3D '97: Proceedings of the 1997 symposium on Interactive 3D graphics*, pages 75–81, New York, NY, USA. ACM Press.

[117] David Luebke. Developer's survey of polygonal simplification algorithms. *IEEE Computer Graphics & Applications*, 21(3):24–35, May 2001.

[118] David Luebke and Carl Erikson. View-independent simplification of arbitrary polygonal environments. In *Proceeding of the 24th annual conference on Computer graphics and interactive techniques*, pages 199–208, 1997.

[119] Jonathan Mackenzie, Jyri Huopaniemi, Vesa Välimäki, and Izzet Kale. Low-order modeling of head-related transfer functions using balanced model truncation. *IEEE Signal Processing Letters*, 4:39–41, 1997.

[120] David G. Malham and Anthony Myatt. 3-D sound spatialization using ambisonics techniques. *Computer Music Journal*, 19(4):58–70, 1995.

[121] William L. Martens. Principal components analysis and resynthesis of spectral cues to perceived direction. In *Proceedings of the 1987 International Computer Music Conference (ICMC'87)*, pages 274–281, 1987.

[122] Herman Medwin, Emily Childs, and Gary M. Jebsen. Impulse studies of double diffraction: A discrete Huygen's interpretation. *Journal of the Acoustical Society of America*, 72(3):1005–1013, 1982.

[123] John C. Middlebrooks and David M. Green. Observations on a principal component analysis of head-related transfer functions. *Journal of the Acoustical Society of America*, 92(1):597–599, 1992.

[124] Patrick Min and Thomas Funkhouser. Priority-driven acoustic modeling for virtual environments. *Computer Graphics Forum*, 19(3):179–188, 2000.

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

[126] Damian Murphy, Antti Kelloniemi, Jack Mullen, and Simon Shelley. Acoustic modeling using the digital waveguide mesh. *IEEE Signal Processing Magazine*, 24(2):55–66, 2007.

[127] Graham M. Naylor. ODEON–another hybrid room acoustical model. *Applied Acoustics*, 38(2–4, Special Issue on Computer Modelling and Auralization of Sound Field in Rooms):131–143, 1993.

[128] Fred E. Nicodemus, Joseph C. Richmond, Jack J. Hsia, Irving W. Gingberg, and Thomas Limperis. Geometric considerations and nomenclature for reflectance. Monograph 161, National Bureau of Standards, 1977.

[129] Markus Noisternig, Brian Katz, Lauri Savioja, and Samuel Siltanen. Framework for real-time auralization in architectural acoustics. *Acta Acustica united with Acustica*, 94(6):1000–1015, 2008.

[130] Fakir S. Nooruddin and Greg Turk. Simplification and repair of polygonal models using volumetric techniques. Technical Report GIT-GVU-99-37, GVU Center, Georgia Institute of Technology, 1999.

[131] Eva-Marie Nosal, Murray Hodgson, and Ian Ashdown. Improved algorithms and methods for room sound-field prediction by a acoustical radiosity in arbitrary polyhedral rooms. *Journal of the Acoustical Society of America*, 116(2):970–980, 2004.

[132] Eva-Marie Nosal, Murray Hodgson, and Ian Ashdown. Investigation of the validity of radiosity for sound-filed prediction in cubic rooms. *Journal of the Acoustical Society of America*, 116(6):3505–3514, 2004.

[133] Jean-Dominique Polack. Playing billiards in the concert hall: The mathematical foundations of geometrical room acoustics. *Applied Acoustics*, 38(2–4, Special Issue on Computer Modelling and Auralization of Sound Field in Rooms):235–244, 1993.

[134] Jovan Popovic and Hugues Hoppe. Progressive simplificial complexes. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 217–224, 1997.

[135] Ville Pulkki. Virtual source positioning using vector base amplitude panning. *Journal of the Audio Engineering Society*, 45(6):456–466, June 1997.

[136] Martin Reddy. Scrooge: Perceptually-driven polygon reduction. *Computer Graphics Forum*, 15(4):191–203, 1996.

[137] Kevin J. Renze. Generalized unstructured decimation. *IEEE Computer Graphics and Applications*, 16(6):24–32, 1996.

[138] Niklas Röber, Martin Spindler, and Maic Masuch. Waveguide-based room acoustics through graphics hardware. In *Proceedings of ICMC06*, November 2006.

[139] Niklas Röber, Martin Spindler, and Maic Masuch. Ray acoustics using computer graphics technology. In *Proceedings of the 10th International Conference on Digital Audio Effects*, pages 117–123, September 2007.

[140] Rémi Ronfard and Jarek Rossingac. Full-range approximation of triangulated polyhedra. *Computer Graphics Forum*, 15(3):67–76, August 1996.

[141] Jarek Rossignac and Paul Borrel. Multi-resolution 3D approximation for rendering complex scenes. In B. Falcidieno and T. L. Kunii, editors, *Geometric modeling in computer graphics*. Springer Verlag, 1993.

[142] Jörg Sahm, Ingo Soetebier, and Horst Birthelmer. Efficient representation and streaming of 3D scenes. *Computers & Graphics*, 28(1):15–24, February 2004.

[143] David Salesin and Filippo Tampieri. Grouping nearly coplanar polygons into coplanat sets. In David Kirk, editor, *Graphics Gems III*. Academic Press, 1992.

[144] Hanan Samet. *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, 1990.

[145] Lauri Savioja. Improving the three-dimensional digital waveguide mesh by interpolation. In *Proc. Nordic Acoustical Meeting*, pages 265–268, September 1998.

[146] Lauri Savioja, Jyri Huopaniemi, Tapio Lokki, and Riitta Väänänen. Creating interactive virtual acoustics environments. *Journal of the Audio Engineering Society*, 47(9):675–705, 1999.

[147] Lauri Savioja, Timo J. Rinne, and Tapio Takala. Simulation of room acoustics with a 3-D finite difference mesh. In *Proc. Int. Computer Music Conf.*, pages 463–466, September 1994.

[148] Lauri Savioja and Vesa Välimäki. Reducing the dispersion error in the digital waveguide mesh using interpolation and frequency-warping techniques. *IEEE Transactions on Speech and Audio Processing*, 8(2):184–194, March 2000.

[149] Lauri Savioja and Vesa Välimäki. Interpolated rectangular 3-D digital waveguide mesh algorithms with frequency warping. *IEEE Transactions on Speech and Audio Processing*, 11(6):783–790, November 2003.

[150] Gernot Schlaufer and Wolfgang Stürzlinger. Generating multiple levels of detail. In M. Göbel, editor, *Virtual Environments '95 (Eurographics Workshop)*, pages 33–41. Springer Verlag, January 1995.

[151] Dieter Schmalstieg and Gernot Schlaufer. Smooth levels of detail. In *Proceedings of the IEEE 1997 Virtual Reality Annual International Symposium*, pages 12–19, 1997.

[152] Francis J. M. Schmitt, Brian A. Barsky, and Wen-Hui Du. As adaptive subdivision method for surface-fitting from sampled data. *Computer Graphics*, 20(4):179–188.

[153] Peter Schröder and Pat Hanrahan. On the form factor between two polygons. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 163–164, 1993.

[154] William J. Schroeder. A topology modifying progressive decimation algorithm. In *Proceedings of the 8th conference on Visualization '97*, pages 205–212. IEEE Computer Society Press, 1997.

[155] William J. Schroeder, Jonathan A. Zarge, and William E. Lorensen. Decimation of triangle meshes. *ACM SIGGRAPH Computer Graphics*, 26(2):65–70, 1992.

[156] François X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. A global illumination solution for general reflectance distributions. *Computer Graphics*, 25(4):187–196, 1991.

[157] Marc Soucy and Denis Laurendeau. Multiresolution surface modeling based on hierarchical triangulation. *Computer Vision and Image Understanding*, 63(1):1–14, 1996.

[158] Oliver G. Staadt and Markus H. Gross. Progressive tetrahedralizations. In *Proceedings of the conference on Visualization '98*, pages 397–403. IEEE Computer Society Press, 1998.

[159] Evert W. Start. *Direct Sound Enhancement by Wave Field Synthesis*. PhD thesis, Delft University of Technology, 1997.

[160] U. Peter Svensson and Paul T. Calamia. Edge-diffraction impulse responses near specular-zone and shadow-zone boundaries. *Acta Acustica united with Acustica*, 92(4):501–512, 2006.

[161] U. Peter Svensson, Roger I. Fred, and John Vanderkooy. Analytic secondary source model of edge diffraction impulse responses. *Journal of the Acoustical Society of America*, 106(5):2331–2344, 1999.

[162] U. Peter Svensson and Ulf R. Kristianssen. Computational modelling and simulation of acoustic spaces. In *AES 22nd Int. Conf. on Virtual, Synthetic and Entertainment Audio*, pages 11–30, Espoo, Finland, June 2002.

[163] U. Peter Svensson, Mayumi Nakano, Kimihiro Sakagami, and Masayuki Morimoto. A study of the sound radiation from musical instruments in rooms using the equivalent source method. In *Proc. of the 16th International Congress on Acoustics*, pages 365–366, 1998.

[164] Tapio Takala and James Hahn. Sound rendering. *Computer Graphics*, 26(2):211–220, 1992.

[165] Manabu Tanaka, Yoshihiro Furue, Yoshinari Horinouchi, and Toshio Terai. Numerical calculations of transient sound response in rooms by Kirchhoff's integral equation. In *Proc. of the 15th International Congress on Acoustics*, pages 185–188, 1995.

[166] Rendell R. Torres, U. Peter Svensson, and Mendel Kleiner. Computation of edge diffraction for accurate room acoustics auralization. *Journal of the Acoustical Society of America*, 109(2):600–610, February 2001.

[167] Issac J. Trotts, Bernd Hamann, Kenneth I. Joy, and David F. Wiley. Simplification of tetrahedral meshes. In *Proceedings of the conference on Visualization '98*, pages 287–295. IEEE Computer Society Press, 1998.

[168] Nicolas Tsingos, Ingrid Carlbom, Gary Elko, Robert Kubli, and Thomas Funkhouser. Validating acoustical simulations in the bell labs box. *IEEE Computer Graphics and Applications*, 22(4):28–37, 2002.

[169] Nicolas Tsingos, Carsten Dachsbacher, Sylvain Lefebvre, and Matteo Dellepiane. Instant sound scattering. In *Proceedings of the Eurographics Symposium on Rendering*. Eurographics, 2007.

[170] Nicolas Tsingos, Thomas Funkhouser, Addy Ngan, and Ingrid Carlbom. Modeling acoustics in virtual environments using the uniform theory of diffraction. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 545–552, 2001.

[171] Nicolas Tsingos, Emmanuel Gallo, and George Drettakis. Perceptual audio rendering of complex virtual environments. In *Proceedings of the 31st annual conference on Computer graphics and interactive techniques*, pages 249–258, 2004.

[172] Nicolas Tsingos and Jean-Dominique Gascuel. Acoustic simulation using hierarchical time-varying radiant exchanges. http://citeseer.ist.psu.edu/721526.html [cited 2007-06-27].

[173] Greg Turk. Re-tiling of polygonal surfaces. *ACM SIGGRAPH Computer Graphics*, 26(2):55–64, 1992.

[174] Maarten van Walstijn and Konrad Kowalczyk. On the numerical solution of the 2D wave equation with compact FDTD schemes. In *Proc. of the 11th Int. Conference on Digital Audio Effects*, pages 1–8, September 2008.

[175] Amitabh Varshney. *Hierarchical geometric approximations*. PhD thesis, Department of Computer Science, Univerity of North Carolina at Chapel Hill, 1994.

[176] Philippe Véron and Jean-Claude Léon. Static polyhedron simplification using error measurements. *Computer-Aided Design*, 29(4):287–298, April 1997.

[177] Michael Vorländer and Eckard Mommertz. Definition and measurement of random-incident scattering coefficients. *Applied Acoustics*, 60(2):187–199, 2000.

[178] John R. Wallace, Kells A. Elmquist, and Eric A. Haines. A ray-tracing algorithm for progressive radiosity. *Computer Graphics*, 23(3):315–324, 1989.

[179] Andrew Wareing and Murray Hodgson. Beam-tracing model for predicting sound field in rooms with multilayer bounding surfaces. *Journal of the Acoustical Society of America*, 118(4):2321–2331, 2005.

[180] Lance Williams. Pyramidal parametrics. *ACM SIGGRAPH Computer Graphics*, 17(3):1–11, 1983.

[181] Julie C. Xia, Jihad El-Sana, and Amitabh Varshney. Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):171–183, 1997.

[182] Steve Zelinka and Michael Garland. Permission grids: practical, error-bounded simplification. *ACM Transactions on Graphics*, 21(2), 2002. Also available as http://mgarland.org/files/papers/pgrid.pdf.